

AIAND GAME THEORY

Editor: Michael Wooldridge, University of Oxford, mjw@cs.ox.ac.uk

Defeating Terrorist Networks with Game Theory

Tomasz P. Michalak, University of Oxford Talal Rahwan, Masdar Institute of Science and Technology Oskar Skibski, Kyushu University Michael Wooldridge, University of Oxford

Despite enormous efforts to tackle terrorism in the aftermath of 9/11, many terrorist networks are growing in size and in threat. Faced with this challenge, security agencies urgently require robust and efficient techniques to analyze and understand the structure and operation of terrorist organizations. One key problem is that of identifying the key members of the organization using information about the terrorist network's topology: this capability would enable security agencies to focus severely limited resources on just those key members.

To this end, many standard measures of centrality from the field of social network analysis can be used. Centrality measures aim to give a numerical characterization of a node's significance in a network. For example, IBM's Analyst Notebooka software package used worldwide by law enforcement and intelligence agencies-supports the degree, closeness, and betweenness centralities, which are probably the most widely used centrality measures. However, these measures have wellknown limitations, and the development of more sophisticated measures is currently a key research area. One important thread of current research is to apply techniques from the field of cooperative game theory to this problem. In this article, we present a survey of this work.

Classic Measures of Centrality

The networks we are interested in are undirected weighted graphs: they consist of a set of nodes (vertices) and undirected edges that connect some nodes; each edge is associated with a numerical weight. The specific interpretation of the network will depend on the application domain, but in the analysis of terrorist networks, nodes are interpreted as members of the network, edges indicate links between them, and weights indicate the strength of the link (for example, the number of meetings between the two individuals). Centrality analysis aims to create a principled ranking of the importance of nodes within such a network. Since "importance" depends on the problem at hand, researchers have considered many different centrality measures. Among them, degree centrality, closeness centrality, and betweenness centrality are the most well-known and widely applied. According to weighted degree centrality, a node's importance is equal to the weight of that node's adjacent edges (that is, the weighted degree of the node). For instance, in the network in Figure 1, nodes v_1 and v_2 are the most important according to weighted degree centrality, because they have a weighted degree of 6 each, which is greater than that of any other node in the network. In contrast, according to closeness centrality, a node's importance is based on the average distance between that node and other nodes in the network. In Figure 1, for example, this measure considers node v_8 to be the most important in the network. Finally, betweenness centrality focuses on shortest paths between any two nodes in the network; the more shortest paths a given node belongs to, the higher its ranking. In Figure 1, node v_{11} is the most important from the perspective of betweenness centrality.

All of these centralities can be defined on weighted, unweighted, directed, and undirected networks. For instance, the unweighted in-degree centrality of a node counts the total number of its incoming edges.

Clearly, all three measures expose a node's different characteristics. For example, assume that



Figure 1. Sample network of 13 nodes with weighted edges. The letter *F* represents a terrorist with the necessary funds, and the letter *B* represents a terrorist with bomb expertise.

Figure 1 represents a network of cell-phone connections made by the members of a terrorist network; the weight of any given edge represents the number of calls made between the two ends of that edge. In this context, we can think of degree centrality as a simple tool to identify the terrorist(s) with the highest frequency of direct contacts. The hypothesis here is that such individuals are the most active within the network. Conversely, we can think of closeness centrality as a tool to identify the terrorist(s) whose message would spread fastest throughout the organization. Finally, betweenness centrality reveals the nodes that play a crucial role in passing the information from one individual in the network to another.

Although the standard measures indeed deliver nontrivial insights, they cannot capture certain important aspects of centrality in a terrorist organization.

First, standard centrality measures are defined for networks in which we can consider at most a single realvalued weight for every node or edge. This assumption restricts the qualitative analysis of terrorist organizations. For instance, it might be difficult or even impossible for a single weight to reflect the fact that a terrorist possesses certain critical resources and/or skills, such as financial endowments or expertise in explosives. To illustrate this point, let us assume that node v_1 in Figure 1 represents the terrorist who provides funding to the entire organization, while nodes v_2 , v_6 , and v_{11} represent the available bomb experts. Now, let us try to express the fact that the funding and bomb expertise are both critical to carry out an attack, and that the former skill is possessed by a single individual, while the latter skill by three. To this end, we could for instance assign a weight of 10 to v_1 , a weight of 5 to every bomb expert, and a weight of 1 to every other node. Then, if an attack requires the funds, one bomb expert, and at least two other members, we can set a threshold of 17 on the total weight of the terrorists needed to carry out an attack. Unfortunately, such a solution would not work, because a group of three bomb experts and two other members (other than the one providing the funds) would cross this threshold, although they are unable to carry out any attack. This simple example illustrates how it could be hard for a quantitative analysis to replace a qualitative one.

The second deficiency of standard centrality measures is that they assess

the importance of each node by focusing only on the role that this node plays by itself in the network. However, in many applications, such an approach is inadequate as it fails to capture any positive or negative synergies that might exist among different groups of nodes. To illustrate this point, let us consider the problem of finding a group of two individuals that together have the highest influence in the network. As a possible solution to this problem, let us focus on unweighted degree centrality. With this measure, because nodes v_1 and v_2 have five neighbors each, and because no other node has that many neighbors, then v_1 and v_2 each have the highest centrality. One might then conclude that $\{v_1, v_2\}$ is the most influential pair in the network. However, it is easy to observe that their spheres of influence (that is, their neighborhoods) overlap. In particular, both v_1 and v_2 are linked to v_6 , v_7 , and v_8 . Thus, nodes v_1 and v_2 together have only 7 unique neighbors, not 5 + 5 =10 neighbors. To put it differently, in this example there is a negative synergy between v_1 and v_2 due to their overlapping spheres of influence. Taking such synergies into consideration, we find that the most influential pair is in fact $\{v_1, v_{11}\}$, not $\{v_1, v_2\}$, because v_1 and v_{11} together have eight unique neighbors. This is despite the fact that v_{11} is actually ranked lower than v_2 .

To account for such synergies, and to quantify the importance of groups of nodes more accurately, Martin Everett and Steve Borgatti introduced the notion of *group centrality*.¹ Intuitively, group centrality is derived in the same way as standard centrality, but now we consider the functioning of a group of nodes rather than an individual node. For instance, as we have already established in Figure 1, the unweighted degree centrality of group { v_1 , v_2 } is 7, whereas the unweighted degree centrality of group $\{v_1, v_{11}\}$ is 8.

Unfortunately, the notion of group centrality has the following inherent drawback: given a network of size n, there are 2^n groups of nodes to consider. Not only is this computationally prohibitive, but it is also unclear how to construct a coherent ranking of individual nodes using such an exponential number of group-wise results. For instance, we could consider the most important nodes to be those belonging to the most important group. Alternatively, we could rank individual nodes according to the average importance of the groups they belong to. In fact, one can think of an endless list of such alternatives. Things become even more complicated if, for instance, we are only interested in groups of nodes that are connected (that is, they induce a connected subgraph).

In summary, on top of evaluating the role that each node plays in a network's topology, we have identified the following desirable properties to have in a centrality measure:

- 1. It should be able to account for the individual properties of each node and/or edge (expressed either quantitatively or qualitatively).
- 2. It should be able to consider not only the functioning of individual nodes but also the positive or negative synergies that might exist in different groups of nodes.
- 3. It should be able to aggregate all the above information in an informed and meaningful way to produce a coherent ranking of individual nodes.

By adapting the standard centrality measures to account for qualitative properties of the nodes and/or edges, and by extending those measures to groups, we can address properties 1 and 2. The third property, however, remains elusive. In the next section, we argue that the game-theoretic approach to centrality measures has all the above properties, including the third property.

Game Theory for Network Analysis

This direction of research is based on the following observation: by having an exponential number of groups of nodes, each with an assigned value reflecting its centrality, we end up with exactly the same combinatorial structure as a coalitional game (see the "Coalitional Games" sidebar). As such, we can consider the network's nodes to be the players of a coalitional game, with the value of each group, or "coalition," of nodes being equal to its group centrality. With this analogy, we can capitalize on decades of research in cooperative game theory, which focuses on analyzing the values of different coalitions to determine each individual's worth. In particular, a coalitional game representing the network, along with an adopted solution concept, produce what we call a game-theoretic network centrality measure, whereby a node's centrality is the payoff of that node in the corresponding coalitional game; this payoff is computed according to the adopted solution concept, such as the Shapley value or the Myerson value. Let us consider a few examples.

Example 1: Detecting the Top-*k* Most Influential Terrorists

Given an unweighted and undirected network, one might represent such a setting with a coalitional game in which the value of a group of nodes, *C*, is computed as follows:

value(C) = the number of nodes in C + the number of neighbors of C

Consider coalition $\{v_1, v_9\}$ in Figure 1, for example. This coalition has a total of 7 neighbors—namely, v_2 , v_3 , v_4 , v_5 , v_6 , v_7 , and v_8 —and the size

of the coalition itself is 2. Therefore, value($\{v_1, v_9\}$) = 7 + 2 = 9. This characteristic function can be interpreted as a function that counts the nodes that are directly influenced by a given coalition; those nodes are the members themselves, as well as their neighbors. Having specified how to evaluate different coalitions in the game, we now must specify how to compute the payoff of each player in the game; that is, we need to specify a suitable solution concept. Let that solution concept be the Shapley value, for example. Recall that the Shapley value of a node v_i is a weighted average of its marginal contributions. Now, given the above characteristic function, the marginal contribution of v_i to any given coalition is the change in the coalition's influence that occurs when v_i joins that coalition. In Figure 1, for example, the marginal contribution of v_8 to coalition $\{v_1, v_9\}$ equals 1, because adding v_8 to $\{v_1, v_9\}$ brings one additional node—namely, v_{11} —under the coalition's influence. Formally, value($\{v_1, v_2\}$ v_8, v_9) - value({ v_1, v_9 }) = 1. Thus, by interpreting the Shapley value of v_i as its centrality, we end up with a measure with which the ranking of v_i is proportional to the change in group influence that v_i causes when joining different groups.

Example 2: The Myerson Value-Based Centrality

In the previous example, whether a coalition was connected or not did not have a direct bearing on the coalition's value. Conversely, in this example, we will consider a Myerson's graph-restricted game with the characteristic function defined as follows:²

value(C)

$$= \begin{cases} \frac{\text{the number of edges in } C}{\text{the weight of edges in } C} & \text{if } C \text{ is connected,} \\ \text{the sum of values of} \end{cases}$$

connected components of C otherwise.

Coalitional Games

cooperative game models scenarios where individuals, or players, benefit from cooperation, and where binding agreements are possible. This means it is possible for players to agree on which groups, or coalitions, to form, and on how the payoff, or value, of the formed coalitions shall be divided among the members. A coalition's value is often assumed to depend entirely on the identities of the coalition members. In such games, the function that assigns a value to every coalition is called the *characteristic function*, and is denoted by v.

The vector that specifies each player's share is called the *payoff vector*. The main question here is which payoff vector should we adopt? One common way of addressing this question is to carefully select a number of desirable properties, and then prove that there exists only one possible payoff vector that satisfies all of those properties. The selected properties are then called the *axioms* of this unique vector. In this context, Shapley focused on scenarios where the players form the *grand coalition*—that is, one big coalition containing all players in the game.¹ He proposed the following axioms:

- Efficiency. The grand coalition value is divided entirely among the players (that is, there are no leftovers).
- Null player. The share is zero for every "null player"—that is, a player whose membership in any coalition does not affect the value of that coalition.
- Symmetry. All "symmetric" players receive equal shares, where any two players are symmetric if replacing one with the other in any coalition makes no impact on that coalition's value.
- Additivity. Given two games, G_1 and G_2 , that involve the same set of players, and given a third game, G_3 , in which a coalition's value is the sum of its value in G_1 and in G_2 , the share of a player in G_3 should be the sum of its share in G_1 and in G_2 .

Shapley proved that there is a unique payoff vector satisfying all of the above, whereby the share of a player *i*, called "the Shapley value of *i*," is

Shapley_value^v_i =
$$\sum_{C \subseteq N \setminus \{i\}} \frac{|C|!(n-|C|-1)!}{n!} (v(C \cup \{i\}) - v(C)).$$

For instance, the value of the connected coalition { v_1 , v_5 } in Figure 1 is 2, and the value of the connected coalition { v_2 , v_3 , v_9 , v_{10} } is 2/3. In contrast, the coalition { v_1 , v_2 , v_3 , v_5 , v_9 , v_{10} } is disconnected, meaning that the terrorists involved cannot communicate, and so they will proceed with their activities as two independent groups. Consequently, our graph-restricted game assigns the value of 2 + 2/3 to { v_1 , v_2 , v_3 , v_5 , v_9 , v_{10} }.

The fraction in the above characteristic function was proposed by Lindelauf and colleagues, who provided the following rationale: "A terrorist organization will try to shield its important players by keeping the frequency and duration of their interaction with others to a minimum. However, to be able to coordinate and control the attack, an important player needs to maintain relationships with other individuals in the network."³ Although those authors focused on the Shapley value, Skibski and colleagues showed that the Myerson value has more interesting properties as a centrality measure.²

Example 3: Centrality with Skills

Let us again consider the situation when the critical skills of nodes v_1 , v_2 , v_6 , and v_{11} are known to the investigators. Michalak and colleagues proposed the Myerson value-based

The Shapley value of *i* is a weighted average of player *i*'s marginal contributions, where the marginal contribution of *i* to a coalition *C* is $v(C \cup \{i\}) - v(C)$. More precisely, for any permutation of players, let us consider the marginal contribution of v_i to the coalition consisting of all players that appear before it in the permutation. Then, the Shapley value is simply the average such marginal contribution taken over all permutations.

Myerson followed a similar approach, but for games defined on a graph G in which nodes represent players and edges represent communication channels.¹ Thus, every coalition can be thought of as a subgraph of G. Myerson assumed that if a coalition, or a subgraph, C was not connected in G, then the value of C is simply the sum of the values of the connected components in C. The intuition is that disconnected components cannot communicate with one another, and so there is no value added when considering their union. Formally, the value of C is

$$w(C) = \begin{cases} \sum_{\text{connected component S in C}}^{v(C)} \text{ if } C \text{ is connected in } G \\ \text{otherwise.} \end{cases}$$

For such scenarios, Myerson proposed the following axioms:

- Component efficiency. The payoff of every connected component is divided entirely among its members (that is, there are no leftovers).
- Fairness. Any two connected players benefit equally from the bilateral connection between them.

Myerson proved that there is a unique payoff vector satisfying the above axioms, whereby the share of a player i, called the Myerson value of i, is

 $Myerson_value_i^v = Shapley_value_i^w$.

Reference

1. M. Maschler, E. Solan, and S. Zamir, *Game Theory*, Cambridge Univ. Press, 2013.

Flexibility of Game-Theoretic Centrality

n interesting property of game-theoretic centrality is its flexibility. Specifically, it can be adapted to a particular application, along three general dimensions.

- The form of the game. The game is usually assumed to be in characteristic function form—that is, a coalition's value depends solely on the identities of its members. However, other forms may better model the application at hand. For instance, partition function games allow for a coalition's value to be influenced by the way nonmembers are partitioned. Alternatively, generalized characteristic function games assume that the members of a coalition are ordered, and allow for a coalition's value to be influenced by the order of its members.
- The coalition-value function. Usually, an arbitrary coalition-value function is assumed. Thus, one can pick the function that best represents the worth of a coalition given the application at hand. For instance, Michalak and colleagues consider alternative such functions when studying the problem of information diffusion in networks.¹
- The solution concept. This concept specifies the criteria the players use to decide which coalitions to form and which payoff vector to adopt. Different solution concepts are based on different prescriptive and normative considerations. Examples include the Shapley value and Myerson value (see the "Coalitional Games" sidebar), which focus on fairness, and the core and the nucleolus, which focus on stability.²

A recent overview on the literature on game-theoretic centrality can be found in Tarkowski and colleagues. $^{\rm 3}$

References

- P.T. Michalak et al., "Efficient Computation of the Shapley Value for Game-Theoretic Network Centrality," J. Artificial Intelligence Research, vol. 46, 2013, pp. 607–650.
- 2. M. Maschler, E. Solan, and S. Zamir, *Game Theory*, Cambridge Univ. Press, 2013.
- M. Tarkowski et al., "Game-Theoretic Network Centrality Measures: A Critical Survey," mimeo, Univ. of Oxford, 2014.

centrality measure that accounts for such qualitative data.⁴ Specifically, they proposed the following characteristic function, where *F* denotes a terrorist capable of funding multiple attacks, *B* denotes a terrorist capable of handling bombs, and *M* denotes a terrorist skilled in martial arts:

value(C)

$$= \begin{cases} \min\left\{\text{the number of } B, \frac{M}{2} \text{ in } C\right\} \\ \times (\text{the number of } F \text{ in } C) & \text{if } C \text{ is connected}, \\ \text{the sum of values of the} \\ & \text{connected components of } C & \text{otherwise.} \end{cases}$$

Intuitively, this qualitative characteristic function represents scenarios in which every attack requires a terrorist with the necessary funding, in addition to one bomb expert and two terrorists skilled in martial arts. This can be generalized to more than just a single attackfor example, the terrorist providing the funds, combined with two bomb experts and four martial artists, can carry out two attacks. Naturally, the numbers and skills specified in the above function can be adjusted to better reflect the scenario under investigation. Note that, for each attack, the group must be connected to be able to coordinate its activities.

These examples demonstrate that game-theoretic centrality measures account for individual properties of nodes and edges, expressed both quantitatively and qualitatively (property 1). They also analyze all groups (that is, subsets) of nodes (property 2). Finally, they use all those pieces of information to construct a coherent ranking of individual nodes (property 3). On top of this, by drawing parallels between networks and coalitional games, we can build on a large body of literature on game-theoretic solution concepts. Importantly, the game-theoretic approach to centrality places no assumptions on the coalition-evaluation function, nor the adapted solution concept, nor even the form of the game (see the "Flexibility of Game-Theoretic Centrality" sidebar); these can be adjusted as necessary, thus providing a desirable degree of flexibility to capture a wide range of scenarios.

But alas, every rose has its thorns. All the advantages of game-theoretic centrality measures might never materialize if we cannot compute those measures in the first place. Next, we discuss recent works that address this issue.

The Computational Challenge

Unfortunately, due to their inherent combinatorial nature, computing solution concepts for coalitional games is often computationally challenging. In fact, already the size of the input (that is, the number of elements of the characteristic function) grows exponentially with the number of players. Luckily, in many cases the coalition values in the coalitional game have some underlying structure, which makes it possible to represent the characteristic function compactly. This has led to an important research direction in algorithmic game theory, aimed at identifying classes of games that admit such a compact representation, and determining whether and how certain solution concepts are computable in polynomial time.

In the context of game-theoretic centrality, the underlying coalitional game is defined over the network, and is somehow influenced by the network topology. This suggests that, in some cases, such coalitional games can be (but do not have to be) represented compactly. For instance, this is the case for the games in Examples 1, 2, and 3 given earlier. However, having a compact representation does not necessarily imply that

JANUARY/FEBRUARY 2015

the solution concept of interest is computable in polynomial time. In this respect, we can summarize the available results in the literature as follows:

- On one hand, negative results were obtained for game-theoretic centralities in which, in the spirit of Myerson's graph-restricted games, the value of a coalition explicitly depends on whether this coalition is connected or not. For instance, both the Myerson value-based centrality measures in Examples 2 and 3 are hard to compute.²
- On the other hand, positive results were obtained for various gametheoretic centralities in which, in the spirit of Everett and Borgatti, the value of a coalition of nodes is simply equal to its group centrality, meaning that whether the coalition is connected or not does not have an explicit bearing on its value.⁵

Next, we will analyze the Shapley value-based measure from Example 1, which is an instance of the above positive results. After that, we will briefly discuss the state-of-the-art algorithm to compute the Myerson value-based centrality, which has exponential worst-case complexity. In practice, however, it solves problems of sparse networks with about 50 nodes in a matter of minutes on a modern workstation. Furthermore, for bigger problems, an efficient albeit inexact Monte Carlo sampling algorithm can be used.⁵

The Shapley Value-Based Centrality Computable in Polynomial Time

Let us now see how the game-theoretic centrality from Example 1 can be computed in polynomial time.⁵ Recall that the Shapley value can be interpreted as the average marginal contribution taken over all possible permutations of players (nodes) (see the "Coalitional Games" sidebar). In particular, to compute the Shapley value of v_i , we first identify all possible permutations of nodes in which v_i makes a positive marginal contribution to the coalition C consisting of all the nodes that appear before v_i in the permutation. To this end, if v_i is a neighbor of v_i , we must ask what is the necessary and sufficient condition for node v_i to "marginally contribute" its neighbor v_i to C? This clearly happens if and only if neither v_i nor any of its neighbors are present in C (in this case, when v_i joins C, its neighbor v_i becomes under the influence of C). In Figure 1, for example, v_2 marginally contributes its neighbor v_9 to a coalition C if and only if neither v_9 nor any of its neighbors (that is, v_2 and v_3) are members of C. Now, let us count the number of possible permutations in which v_2 marginally contributes its neighbor v_9 to the coalition C consisting of all the nodes that appear before v_2 in the permutations. To do so,

- We randomly choose |{v₂, v₃, v₉}| = 3 positions in a permutation consisting of all 13 nodes from Figure 1.
- This can be done in $\binom{13}{3}$ ways. We

place v_2 in the first position, and place v_3 and v_9 randomly in the remaining two positions. This can be done in 2! ways.

• The remaining elements can be arranged in (13 - 3)! = 10! ways.

Thus, there are in total
$$\binom{13}{3}$$
 2!10! =

2075673600 permutations in which v_2 marginally contributes v_9 . Given that there are in total 13! = 6227020800 permutations, the fraction of the Shapley value of v_2 that is due to contributing v_9 is 1/3. To compute the total value of the Shapley value of v_2 , we should perform the above procedure for v_2 and every neighbor of v_2 . The general formula is as follows:

ShapleyValue_i =
$$\frac{1}{1 + \text{degree of } v_i}$$

+ $\sum_{v_i \in \text{Neighbors of } v_i} \frac{1}{1 + \text{degree of } v_j}$

The algorithm that implements this expression for all nodes in the network iterates exactly once through all nodes and their neighbors. Thus, its running time is O(|V| + |E|), where |V| is the number of nodes and |E| is the number of edges in the network.

Recall that the Shapley value is used as a centrality measure in our example. Thus, by looking at the above formula, one can see where a node's centrality comes from. In particular, if a node has a high degree, the number of terms in its Shapley value summation is high. However, the terms themselves will be inversely related to the degree of neighboring nodes. Thus, a node has high centrality not only when its degree is high, but also whenever the degree of its neighboring nodes is low. To put it differently, the power comes from being connected to many who are powerlessa phenomenon well-recognized in the centrality literature.

An Algorithm for the Myerson Value-Based Centrality

When computing game-theoretic centralities based on Myerson's graphrestricted games, the most critical operations are the enumeration of induced connected subgraphs, and the identification of cut vertices. Although researchers studied both operations independently, until recently there was no dedicated algorithm that performs both at the same time. Moreover, the state-of-the-art algorithm for each operation was different than the other algorithm in the way it traversed the graph; enumerating all induced connected subgraphs required breadth-first search (BFS), whereas identifying cut vertices required depth-first search (DFS).

To address this issue, Skibski and colleagues proposed the first algorithm that performs both operations simultaneously.2 The main trick is to perform the first operation (enumerating all induced connected subgraphs) using DFS rather than BFS, thus making it compatible with the other operation's performance (identifying cut vertices). Broadly speaking, the DFS enumeration algorithm starts with a subgraph consisting of a single node, and gradually expands the subgraph to (some of) its neighboring nodes in a DFS manner. This is done using a divide-and-conquer technique, which works as follows: whenever a neighbor is found that has not yet been visited by the algorithm, the calculations are split into two parts, the first with the neighbor added to the subgraph, and the second with the neighbor marked as forbidden. This entire process is repeated recursively such that eventually all induced subgraphs are enumerated without any redundancies (that is, without performing an operation twice). For a detailed description of the algorithm, as well as the pseudocode, see Oskar Skibski and colleagues²

Case Study: World Trade Center, 9/11

In the already classic work, Valdis E. Krebs constructed the World Trade Center 9/11 terrorist network from publicly available sources and used standard centrality measures to determine the key terrorists.⁶ Let us focus on the core of the 9/11 network composed of 19 hijackers with 32 relationships, as illustrated in Figure 2. To hijack the four planes on 11 September 2001, the terrorists needed





Figure 2. World Trade Center 9/11 terrorist network.^{3,6} The letter *P* stands for a pilot, and the letter *M* indicates a terrorist skilled in martial arts.

to split into four groups. Each group had to contain at least one terrorist with pilot training (P) and some terrorists skilled in martial arts (M) in order to take control over the crew and the cabin. To fulfill these requirements, terrorists S. Suqami and A.A. Al-Omari took part in martial arts training in the weeks before the attack; it is unknown who else was skilled in martial arts. Therefore, we assigned this skill to seven other terrorists chosen uniformly at random, and assumed that two terrorists with such skills were needed per flight.

Let us now consider a qualitative characteristic function similar to that of Example 3, but now modified such that a *P*-terrorist (pilot) and two *M*-terrorists (martial arts) are needed to carry out a single attack. Table 1 compares the (qualitative) Myersonvalue-based centrality with the other standard centralities for Figure 2. We observe that, as expected, the new centrality measure promotes the terrorist

Rank	Degree		Closeness		Betweeness		Qualitative Myerson	
1	v ₁₆ —N. Alhazmi	N/A	v ₁₆ —N. Alhazmi	N/A	v ₁₆ —N. Alhazmi	N/A	v ₉ —H. Hanjour	Р
2	v ₇ —M. Al-Shehhi	Р	v ₄ —A.A. Al-Omari	М	v_5 —M. Atta	Р	v ₇ —M. Al-Shehhi	Р
3	v ₁₇ —H. Alghamdi	М	v_5 —M. Atta	Р	v ₇ —M. Al-Shehhi	Р	v ₁₂ —Z. Jarrah	Р
4	v₀—H. Hanjour	Р	v ₇ —M. Al-Shehhi	Ρ	v_9 —H. Hanjour	Р	v ₄ —A.A. Al-Omari	М
5	v ₅ —M. Atta	Р	v_1 —Wail Alshehri	N/A	v ₁₂ —Z. Jarrah	Р	<i>v</i> ₅—M. Atta	Р
6	v ₁₂ —Z. Jarrah	Р	v_{17} —H. Alghamdi	М	v_{17} —H. Alghamdi	М	v ₁₃ —A. Al-Haznawi	М
7	v ₁₄ —S. Alghamdi	N/A	v_9 —H. Hanjour	Р	v_{11} —S. Alhazmi	М	v_{11} —S. Alhazmi	М
8	v ₄ —A.A. Al-Omari	М	v ₁₂ —Z. Jarrah	Р	v ₄ —A.A. Al-Omari	М	v ₁₇ —H. Alghamdi	М
9	v ₁ —Wail Alshehri	N/A	v_8 —F. Ahmed	М	v_{14} —S. Alghamdi	N/A	v ₁₆ —N. Alhazmi	N/A
10	v ₁₃ —A. Al-Haznawi	М	v ₁₅ —M. Alshehri	N/A	v ₁₃ —A. Al-Haznawi	М	v ₈ —F. Ahmed	М
11	v_{11} —S. Alhazmi	М	v ₁₃ —A. Al-Haznawi	М	v_8 —F. Ahmed	М	v_{10} —K. Al-Mihdhar	М
12	v ₁₈ —A. Alnami	М	v_{11} —S. Alhazmi	М	v ₁₈ —A. Alnami	М	v_6 —M. Moqed	М
13	v ₈ —F. Ahmed	М	v_{14} —S. Alghamdi	N/A	v_{10} —K. Al-Mihdhar	М	v ₁₈ —A. Alnami	М
14	v ₁₅ —M. Alshehri	N/A	v ₁₈ —A. Alnami	М	v_{15} —M. Alshehri	N/A	v_3 —Waleed Alshehri	N/A
15	v ₁₀ —K. Al-Mihdhar	М	v_{10} —K. Al-Mihdhar	М	v_6 —M. Moqed	М	v ₂ —S. Suqami	М
16	v ₂ —S. Suqami	М	v_2 —S. Suqami	М	v ₁ —Wail Alshehri	N/A	v ₁₅ —M. Alshehri	N/A
17	v ₃ —Waleed Alshehri	N/A	v_3 —Waleed Alshehri	N/A	v_{19} —A. Alghamdi	N/A	v ₁₄ —S. Alghamdi	N/A
18	v ₁₉ —A. Alghamdi	N/A	v_{19} —A. Alghamdi	N/A	v_3 —Waleed Alshehri	N/A	v ₁₉ —A. Alghamdi	N/A
19	v ₆ —M. Moqed	М	v_6 —M. Moqed	М	v₂—S. Suqami	М	v ₁ —Wail Alshehri	N/A

 Table 1. Comparison of rankings for the World Trade Center 9/11 network from Figure 2, given different centrality measures.

 (The Degree, Closeness, and Betweenness columns are from Lindelauf and colleagues.³)

with the necessary skills. Interestingly, the M-terrorist A.A. Al-Omari is ranked higher than the P-terrorist M. Atta. In other words, when considering the positions (in the network) and the skills of different terrorists, it turns out that the martial arts skill of A.A. Al-Omari is more often indispensable to increasing the attack capabilities of connected coalitions of terrorists than the pilot skills of M. Atta. Conversely, N. Alhazmi-the terrorist with no skills, who is ranked at the top by all three standard centrality measures (unable to recognize skills)-is now ranked ninth.

he problem of identifying key members of a terrorist network is an important topic in the literature; it has attracted significant interest within the social network analysis community and well beyond. In this article, we argued that this problem requires a centrality measure that can account for the properties of each terrorist, compute the synergies in various groups (that is, subsets) of those terrorists, and aggregate all this information into a coherent ranking of individual terrorists. This was not possible until recently, when gametheoretic centrality measures were proposed. We showed how these new measures capitalize on decades of research in cooperative game theory, thus providing a rich pool of solution concepts to evaluate individuals on the basis of their performance in different groups. Unfortunately, these new measures are inherently complex and raise various computational challenges, even for relatively small networks. However, we showed how to overcome some of these computational challenges, following recent results from algorithmic game theory. Hopefully, the further development of these measures will shed more light on the dark nature of covert terrorist organizations.

Acknowledgments

Tomasz Michalak and Michael Wooldridge were supported by the European Research Council under Advanced Grant 291528 ("RACE"). Oskar Skibski and Tomasz Michalak were supported by the Polish National Science grant DEC-2013/09/D/ ST6/03920.

References

- M.G. Everett and S.P. Borgatti, "The Centrality of Groups and Classes," *J. Mathematical Sociology*, vol. 23, no. 3, 1999, pp. 181–201.
- O. Skibski et al., "Algorithms for the Shapley and Myerson Values in Graph-Restricted Games," Proc. 13th Int'l Conf. Autonomous Agents and Multiagent Systems (AAMAS 2014), 2014, pp. 197–204.
- R. Lindelauf, H. Hamers, and B. Husslage, "Cooperative Game Theoretic Centrality Analysis of Terrorist Networks: The Cases of Jemaah Islamiyah and Al Qaeda," *European J. Operational Research*, vol. 229, no. 1, 2013, pp. 230–238.

- 4. T.P. Michalak et al., *The Qualitative Game-Theoretic Centrality Measure*, Univ. of Oxford, 2014.
- P.T. Michalak et al., "Efficient Computation of the Shapley Value for Game-Theoretic Network Centrality," J. Artificial Intelligence Research, vol. 46, 2013, pp. 607–650.
- V.E. Krebs, "Mapping Networks of Terrorist Cells," *Connections*, vol. 24, 2002, pp. 43–52.

Tomasz P. Michalak is a research assistant in the Department of Computer Science at the University of Oxford and an assistant professor in the Institute of Informatics at the University of Warsaw. Contact him at tomasz.michalak@cs.ox.ac.uk.

Talal Rahwan is an assistant professor in the Computing and Information Science Department at the Masdar Institute of Science and Technology. Contact him at trahwan@gmail. com.

Oskar Skibski is an assistant professor in the Department of Computer Science at

Kyushu University. Contact him at oskibski@ inf.kyushu-u.ac.jp.

Michael Wooldridge is a professor in the Department of Computer Science at the University of Oxford. Contact him at michael.wooldridge@ cs.ox.ac.uk.

C11 Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.

IEEE (computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field. **COMPUTER SOCIETY WEBSITE:** www.computer.org

Next Board Meeting: 26-30 January 2015, Long Beach, CA, USA

EXECUTIVE COMMITTEE

President: Dejan S. Milojicic

President-Elect: Thomas M. Conte; Past President: David Alan Grier;
Secretary: David S. Ebert; Treasurer: Charlene ("Chuck") J. Walrad; VP,
Educational Activities: Phillip Laplante; VP, Member & Geographic
Activities: Elizabeth L. Burd; VP, Publications: Jean-Luc Gaudiot; VP,
Professional Activities: Donald F. Shafer; VP, Standards Activities: James
W. Moore; VP, Technical & Conference Activities: Cecilia Metra; 2014
IEEE Director & Delegate Division VIII: Roger U. Fujii; 2014 IEEE Director
& Delegate Division V: Susan K. (Kathy) Land; 2014 IEEE Director-Elect & Delegate Division VIII: John W. Walz

BOARD OF GOVERNORS

Term Expiring 2014: Jose Ignacio Castillo Velazquez, David S. Ebert, Hakan Erdogmus, Gargi Keeni, Fabrizio Lombardi, Hironori Kasahara, Arnold N. Pears

Term Expiring 2015: Ann DeMarle, Cecilia Metra, Nita Patel, Diomidis Spinellis, Phillip Laplante, Jean-Luc Gaudiot, Stefano Zanero Term Expriring 2016: David A. Bader, Pierre Bourque, Dennis Frailey, Jill I. Gostin, Atsuhiro Goto, Rob Reilly, Christina M. Schober

EXECUTIVE STAFF

Executive Director: Angela R. Burgess; Associate Executive Director & Director, Governance: Anne Marie Kelly; Director, Finance & Accounting: John Miller; Director, Information Technology & Services: Ray Kahn; Director, Membership Development: Eric Berkowitz; Director, Products & Services: Evan Butterfield; Director, Sales & Marketing: Chris Jensen

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928 Phone: +1 202 371 0101 • Fax: +1 202 728 9614 Email: hq.ofc@computer.org

Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720 Phone: +1 714 821 8380 • Email: help@computer.org

MEMBERSHIP & PUBLICATION ORDERS

Phone: +1 800 272 6657 • Fax: +1 714 821 4641 • Email: help@computer.org Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan • Phone: +81 3 3408 3118 • Fax: +81 3 3408 3553 • Email: tokyo.ofc@computer.org

IEEE BOARD OF DIRECTORS

President: J. Roberto de Marca; President-Elect: Howard E. Michel; Past President: Peter W. Staecker; Secretary: Marko Delimar; Treasurer: John T. Barr; Director & President, IEEE-USA: Gary L. Blank; Director & President, Standards Association: Karen Bartleson; Director & VP, Educational Activities: Saurabh Sinha; Director & VP, Membership and Geographic Activities: Ralph M. Ford; Director & VP, Publication Services and Products: Gianluca Setti; Director & VP, Technical Activities: Jacek M. Zurada; Director & Delegate Division V: Susan K. (Kathy) Land; Director & Delegate Division VIII: Roger U. Fujii

revised 6 November 2014

