Articles

The Use of Structures in Communication Networks to Track Membership in Terrorist Groups

by H.A. Eiselt and J. Bhadury

Abstract

This concept paper investigates possibilities to detect terrorist cells based on communications between individuals without the need for wiretapping. The advantages of such procedure are apparent: fewer (if any) legal requirements, and, most importantly, the possibility to automate the surveillance. After a brief review of the pertinent literature, we offer three approaches that are designed to aid in the detection of not only terrorist cells, but also the command structures within the cells. The techniques are demonstrated by using a small illustration. The paper concludes by outlining limitations of the procedures described here.

Key words: Counterterrorism, terrorist cell detection, command structure, graph structures, social network analysis.

1. Introduction

Recent events surrounding Mr. Edward Snowden (Greenwald, 2013) have put some activities of the National Security Agency (NSA) into the limelight and under the microscope. In particular, questions regarding the legality of collecting (and storing indefinitely) a large number of data appear troublesome. One such example is the ruling by the European Court of Justice that mandated Google to remove websites with personal information or links to them if requested by individuals, if warranted (The Times-Picayune, 2014). Furthermore, the potential of such information gathering resulting in a very large number of "false positives" is problematic. This paper attempts to describe automatic approaches that only deal with metadata, thus falling into the category of traffic analysis (a toolkit that originated during World War II; see, e.g., Ressler, 2006), thus limiting invasion of privacy as well as avoiding cumbersome legal requirements (as described in the Sentinel Visualizer, 2013). Traffic analysis has the advantage of not having to deal with translations, steganography, cryptography or similar difficulties. Furthermore, it is able to deal with huge sets of data, leaving only fairly small numbers of individuals who have to be investigated further, a labor-intensive, legally challenging and costly undertaking. In other words, the analyses presented in this paper are of a preparatory nature that filters out many of the connections and individuals, who are most likely not related to acts of terrorism. The underlying assumption here is that once this smaller set of connections and individuals are identified, they will be analyzed using much more in-depth techniques that will require knowledge of content and therefore, legal approval.

As is the case with most (and not just automated) methods, there are possible errors. Erring on the cautious

side, the analyst will be left with many "false positive" leads that require surveillance at great cost. On the other hand, very tight selection rules may allow actual terrorists remain undetected with possibly disastrous results. This clearly calls for fine tuning by knowledgeable analysts.

Any social network analysis that deals with terrorist activities consists of three phases: first, there is the *development phase*, in which a network of relevant individuals is developed. Typically, work in this phase starts with a small number of known or suspected terrorists, and, based on their former or present contacts a "trust network" (Everton, 2012) is grown. The second phase is the *delineation phase*, in which present connections between individuals in the trust network are determined, thus forming operational networks. Furthermore, clusters (cells) and the hierarchy within these cells (if any) are also determined. Finally, phase 3 is the *deletion phase*, in which a plan is drawn up to remove the threat by eliminating individuals in the cell, so as to maximize the (short or long-term) damage done to the terrorist network. It appears reasonable that steps in this phase should preferably wait until a terrorist act appears imminent, as the destruction of even a small part of the network will signal to the terrorists that they have been detected, perhaps necessitating the work of counterterrorist operations to start afresh.

While this paper does not deal with the deletion phase, we offer a few thoughts on the subject since that is generally the step that follows detection. While the elimination of threats attributed to (mostly, but not necessarily, leading) individuals, appears an obvious task, it is interesting to note from the literature that this may actually be counterproductive in some circumstances. For instance, Medina (2012) has demonstrated that his 381-node terror network was surprisingly resilient and did not show any major deterioration after two of the key figures were removed. The author asserts that "Attacks on the network can lead to further decentralization and splintering." This leads to a general question of the effectiveness of many measures of counterterrorism: if the elimination of top individuals leads to the fragmentation of the movement, as well as its further decentralization and metastasizing (thus the creation of an enemy that is difficult to detect, let alone destroy), would the lack of counterterrorism measures (at least as they do not directly threaten a country's own population), let the movement morph from a decentralized collection of cells (such as al-Qaeda) into a hierarchically structured organization (such as Hamas) that is potentially easier to deal with? This and similar questions are posed by MacGinty (2010).

Another issue in the deletion phase is dealt with by Xu and Chen (2008). They examine optimal removal strategies in random, small-world, and scale-free networks and confirm Holme *et al.*'s (2002) assertion that small-world networks are more sensitive to attacks on "bridges," i.e., nodes that are included on many shortest paths, on which communication presumably takes place.

For the most part, this paper deals with the delineation phase. More specifically, we are interested in using social network analysis in order to detect operational networks and, if possible, determine their command structure of a terrorist network or cell.

This paper develops three different techniques that allow to automatically analyze metadata regarding (mostly, but not exclusively, electronic) communications. The first method uses two stages: in the first stage, it employs a static analysis to detect cells, while the second, more complex, stage, it is dynamic in that it provides a way to analyze temporal data in a reasonably compact way so as to determine potential command structures within cells. The second technique describes graph structures, which, if detected in temporal communication networks, may indicate terrorism related activity that warrants closer inspection. Finally, the third approach is also dynamic. It tracks the degree of separation (Milgram, 1967) of suspected individuals over time in a manner that makes evident the potentially emerging leaders.

The limitations of the techniques described in this concept paper are clearly those that are related to the implementation of these techniques. In particular, extensive field tests are required to fine-tune the techniques so as to balance the occurrence of false positives and missed detections. It is also worth re-emphasizing that while we tend to put our discussion in the framework of electronic communications, such as phone, email or social media, the methods described in this work are applicable to all types of communication.

The remainder of this paper is organized as follows. Section 2 provides a short literature review, while Section 3 introduces static and temporal (dynamic) graphs, how to generate them from communication records, and Section 4 outlines some structures in dynamic graphs and what they tend to symbolize in general and in terrorist networks in particular. The last Section summarizes our analysis and provides some potential extensions of this study.

2. Literature review

As evident from the references, terrorist networks and their detection have received wide attention in the literature. However, most contributions in the literature deal with the deletion stage. As Krebs (2008) put it succinctly: "When activity reaches a certain pattern and threshold, it is time to stop monitoring the network, and time to start removing nodes." Farley's (2007) work also deals with ways to optimally render a known terrorist network inoperable, given that not all of its members can be reached, as they may reside in countries or areas the affected state has no jurisdiction over.

In order to determine approximate rankings of individuals in a terrorist cell, a number of measures of centrality have been devised. Starting with the classic work by Freeman (1979), the main measures are the degree centrality of a node, the closeness centrality, and the betweenness centrality as well as the eigenvector centrality. In most of the literature, there is the understanding that a central figure represents an important or powerful individual.

Based on the work of Cook *et al.* (1983), Bonacich (1987) was among the first to disagree with such an assertion. He demonstrates that *centrality* and *power* are two different concepts, emphasizing that power derives from connections to powerless individuals (to paraphrase Caesar, this is akin to being more powerful as the first in a village as opposed to the second in Rome). While this may not always be the case, researchers such as Takkala (2005), and Varden (2011) argue that powerful leaders in terrorist networks (as opposed to the heads of leaders of other organizations or organized crime, e.g., Mafia, Cosa Nostra, Hell's Angels) tend to be charismatic individuals, since they have to convince other members of the organization to participate in attacks that can be lethal for their own selves, something not likely achieved by force. As Gordijn and Stapel (2008) remark, charismatic leaders with a controversial message are most likely to convince (certain) people to joining their group. Sociologist Max Weber (1922) distinguishes between legal, traditional, and charismatic authority. The leaders who belong to these classes are pragmatic, ideologues, and charismatic leaders. Charismatic leaders tend to be a more of a rare commodity as opposed to technocrats, so that depriving a terrorist cell of its leader is a much harder blow to the organization as opposed to depriving an organized crime organization of their chief. This makes the determination of the position of a member of a terrorist cell valuable information to law enforcement agencies. Furthermore, this argument makes leadership decapitation much more crippling for terrorist networks than for networks of non-terrorist organized crime; see Price (2012). Finally, we mention that a thoughtful analysis of leadership decapitation in the specific case of the *sendero luminoso* in Peru can be found in Oliva (2005) and that the use of theoretical methods such as

partially ordered sets to render terrorist networks nonfunctioning are presented by Farley (2007).

Similarly, individuals in *important* specialist positions, such as bomb makers, encryption specialists, cannot be replaced as easily. Knowledge of the position an individual occupies is central to the decision whether or not to remove him. However, it appears much easier to identify the leader of a cell based on concepts of centrality than specialists, who tend to occupy peripheral staff positions and are not necessarily well connected. Some caution is advised by Kitsak *et al.* (2011), who demonstrate that the best transmitter of information is not necessarily the individual who is best connected.

When terrorists design their networks, theirs is the choice between *efficiency* and *security*. Simply speaking, the most efficient network (as defined by Qiang and Nagurney, 2008, as the inverse of the average length of the shortest path between all pairs in the network) would be a star graph with the leader being at the center of the graph, while all other cell members occupy the leaves of the tree. Such a tree is, however, very vulnerable: the removal of the leader would leave all other nodes disconnected. The fact that terrorist networks avoid such structures and lean, in general, much more towards security, was already voiced by Bin Laden (United States Department of Defense, 2001, quoted via Krebs, 2002b) about the perpetrators of the September 2011 attack, who stated that "Those who were trained to fly didn't know the others. One group of people did not know the other group." The emphasis on security in terrorist communication networks has also been pointed out by Morselli *et al.* (2007). For the choice of secrecy over efficiency in non-terrorist criminal networks, see, e.g., Baker and Faulkner (1993), Hutchins and Benham-Hutchins (2010), and Kilberg (2012).

Notwithstanding the discussion presented above about leaders of terrorist networks, some terrorist networks tend to be decentralized; in fact, Sageman (2008) coined the phrase of "leaderless jihad" to describe one such situation. Similarly, Ressler (2006) refers to many terrorist networks as being "loosely connected and adaptive," while Rodríguez (2007) points out that these networks are constantly changing. In addition, as Medina and Hepner (2008) assert, terrorist cells tend to be self-sufficient as far as most decisions and their financing is concerned. However, as Gunaratna and Oreg (2010) assert, "…a network-based organization remains mostly unsuited for carrying out complex tasks that require communication, cooperation, and mostly significant professional trainings." In other words, there will have to be some structure with a leader, lieutenants, and foot soldiers in order to organize difficult tasks in secrecy. It is with this in mind we posit that identification of suspicious structures in terrorist networks should be a pertinent tool of detection regardless of the existence or lack thereof of a singular leader; hence the focus of this paper on the same.

As an aside, using the police adage of "follow the money trail" does not appear promising, as terrorist attacks tend to be relatively cheap for the attackers: the United Nations has estimated a price tag of about US$50,000 for a terrorist attack, while the deadly Madrid bombings were estimated to have cost the attackers a mere US$10,000 (The Age, 2004).

As Sparrow (1991) remarked, criminal networks tend to be incomplete, have fuzzy boundaries, and are dynamic. And it is this latter feature that we will concentrate on in the remainder of this paper.

3. *Analyses of temporal networks*

This section will present three distinct models that allow analysts to identify potential terrorist cells or the standing of cell members within a cell given massive amounts of communication data. As alluded to above, the only input that is needed with these techniques are metadata, i.e., the origin of the communication, its destination, the time of its beginning, and the time of its end. The actual content of the communication is not

required in any of these suggested techniques.

The first method starts with a number of network structures that may be expected to indicate the existence of network cells. The second method is a technique that allows the visualization and analysis of temporal data. It avoids the massive sizes of networks that typically arise in these analyses. At first glance, it may appear that both of these models assume that all orders are made by the leader and propagate through all chains of commands until the last of the foot soldiers. In fact, regardless of who originates in order, the implementation of that order is always operationalized by lower-level commanders and foot soldiers. Thus, regardless of whether a global leader or a local superior gives the orders, the propagation of the information generates certain standard patterns/structures within the communication network and the focus of these two models is to detect them, thereby making these models independent of who generates orders in a terrorist network. Finally, the third technique uses static slices of temporal data to show changes of the degree of separation of individual members of a suspected cells and it allows an analysis of their standing within the organization.

3.1. Model 1: Suspicious graph structures in temporal networks

The general approach presented below deals with the identification of graph structures that potentially point to terrorist networks. We first outline a number of basic types of communications and how they would appear in the rooted trees introduced in the previous section. The listing below is by no means comprehensive, the idea is to simply demonstrate certain structures that may be observed. For example, Figure 1 shows unrelated calls between members of the observed set. Concerning the task of identifying terrorist activities, this is essentially noise.


Figure 1

In contrast, Figure 2 shows the typical flow of information in organizations from a superior through the different levels to the foot soldiers


Figure 2

Figure 3 captures the information flow in case of one-to-many reporting, such as an individual in charge of sending out or calling members of a group for a meeting similar to what might be seen in companies, where the information officer sends out mass mailings.


Figure 3

Figure 4 shows a combination of the path in Figure 2 and the one-to-many graph in Figure 3. It is typically for forwarded email messages, such as jokes, Ponzi schemes, etc., which are forwarded from one originator to many recipients.


Figure 4

Consider now the likely communication in security-conscious terrorist cells regardless of the organization, leaders will evolve and chains of command will develop. These are the chains we attempt to find automatically. As alluded to above, there are two types of information flow in an organization, *viz.*, the reporting flow and the flow of orders. There are also flows of general information, which are disseminated by whomever first obtains it. These flows are likely to follow similar paths as the aforementioned to flows. Apart from some redundancies, the reporting flow will be a path similar to that in companies as shown in Figure 2. With the command flow, however, the leader will give a command, which is handed down the chain of command, until at some point, it is distributed by some member to the remaining foot soldiers. The shape of such flow is shown as the *broom graph* in Figure 5.

Figure 5

The idea is to detect subgraphs such as these in the dynamic procedure described in the previous section. Given the evidence from existing literature on the average size of terrorist cells, the length of such chains will be somewhere between four and seven or eight nodes. Clearly, such shapes, if they exist in a communication network, will not appear as shown in the Figures above, but will be hidden, i.e., the structures that we are looking for are almost certainly overlaid with noise. The separation of the main structure from the surrounding noise is fairly simple. As any individual arc that enters or leaves a node in the "handle" of the broom in Figure 5 could be communication by or two an unrelated outside source, the best procedure would be one of periodic overlay. In other words, communications can be monitored at multiple times, then overlaid, and only connections above a preset frequency of occurrence are used in the evaluation of the structures.

3.2: Model 2: Studying directed rooted networks

Temporal networks were extensively used by Carley (2003); for a recent discussion, see, e.g., Holme and Saramäki (2012). In our work, we used the main concepts introduced in that work. In order to perform a dynamic analysis on terrorist networks, we adopt a distinction proposed by Everton (2012), who differentiates between "trust networks" and "operational networks." In this parlance, the network in Figure 6 is a trust network (the 9/11 trust network, first published in Krebs, 2002b, reprinted with permission), while operational networks are those that show communication between members of the trust network (and potentially also communication from/to members of the trust network to/from others on the outside). As an aside, we wish to note that this specific network was constructed in retrospect, after the 9/11 attack; as a result, they were constructed using approaches different than those being outlined in this paper. To reiterate: we assume that a few potential terrorists are known, and their communications are observed and used for analysis.

Figure 6

Since the focus of this paper is the communication between individuals, we will use the term "communication" for observable contacts, such as (directed) email messages, (directed) phone calls, (directed) Facebook contacts, and similar communications as well as non-electronic observations resulting from police work.

For the analysis of the command structure, we suggest a four-step procedure. The first two steps are designed to identify potential terrorist cells. Additional information, including information regarding internal command structures can be obtained by continuing the procedure with the last two steps.

Step 1: Set up the trust network based on prior contacts.

Step 2: Based on the communications between nodes in the network (and possibly to/from nodes beyond), set up a static operational network for the detection of terrorist cells.

Step 3: Given the contact times observed in the metadata of the communication, set up a dynamic operational network.

Step 4: Search for prescribed command structures in dynamic operational network for terrorist cells and command structures.

The usual way to "build" a network is to start with a small number of known terrorists. Given their contacts to individuals outside of the known network (but presumably within their trust network), we can gradually expand the small network by including these contacts as well as the connections between them. This procedure continues for a small number of steps, until a static graph can identify a potential cell.

In order to illustrate the suggested procedure, suppose we have the communications data shown in Table 1. For simplicity, we only consider the origin and destination of a contact as well as the time the call was made. For simplicity, we do not consider the length of the call and the location the calls were made from and to.

Time	Contact from	Contact to
8:01	A	B
8:03	A	C
8:05	D	E
8:06	B	E
8:06	F	A
8:07	E	F
8:09	B	C
8:10	D	C
8:11	A	D

Table 1: Contacts between individuals in our illustration

The first two steps of the procedure result in the usual *static graph*, which ignores the timing of the calls. The static graph for this example is shown in Figure 7.

Figure 7

The interpretation of this network structure is not clear: we discern the circuit *A – D – E – F – A*, as well as the fact that *C* only receives, but does not send, messages. If the messages concern the flow of information, *C* may be a leader, whereas in case the messages carry orders, *C* is either a foot soldier or occupies a specialty function. In order to obtain more information (without attempting to obtain information of the actual contacts), we will have to introduce the timing of the contacts.

First, however, we need to introduce a subtle, yet important distinction. When analyzing the flow of information, it would technically be irrelevant which node initiates a call. For instance, if in the above example node *B* has a certain information, then the directed call from *A* to *B* at 8:01 will result in both, *A* and *B* having this information, assuming that node *B* has actually forwarded it. The situation would be the same if *B* had initiated the call instead, in which case *A* and *B* would also both have this information. In case it is unknown or deemed irrelevant who contacted whom, we refer to *undirected information*. On the other hand, if *B* has a certain information such as an order or a report, then *B* will actually initiate the call rather than wait for other to call him. Such *directed information* can indicate the standing of the caller in relation to the called. This is what we analyze below.

In order to do so, we consider each node, one at a time, as the potential source of information & determine which other nodes in *N* can be reached by the observed contacts. Suppose we start with some node n_i. We then could consider only arcs to nodes n_j, $j \neq i$, if node n_i contacts n_j within a prespecified amount of time, say, Δ.

One possibility here is to consider only arcs, whose origin starts at a time that is not later than Δ (a predefined time frame) than its predecessors. The idea is that an order will be forwarded from one cell member to the next within a fairly short time, especially if it is important. However, if a node made contact with another node within Δ minutes after being contacted by another individual, we will set back the time and allow up to another Δ minutes to make another call. This will permit an individual to transmit a message to multiple recipients. The actual value of Δ will have to be fine-tuned by the analyst, our recommendation being to use a value that is sufficiently large so as not to be restrictive.

We first describe an efficient procedure that determines |*N*| *temporal graphs* for directed information. In essence, we start with one node at a time, assuming that there is some information (e.g., an order) available

at the beginning of the observation period. We then determine the set of nodes that will have obtained this order at the end of the period. This concept is similar to that of the "information set" in game theory. We refer to the resulting graphs as rooted trees (even though, strictly speaking, they are branchings, as the links are directed).

Procedure for the spread of directed information from a node s

Initialize: All (directed) contacts are listed in chronological order as (*i, j*). Define a set *I*, which is originally set to *I*:= {*s*}.

Procedure: *While* the list is not empty,

> *do*
>
> > scan the list from the top until a contact (*i, j*) with
> >
> > $i \in I$ if found, insert (*i, j*) in the graph, set
> >
> > $I := I \cup \{j\}$ & delete this contact & all entries above it.
>
> *end*

The computational complexity is linear in the length of the list, i.e., the number of contacts, for each graph. Given *p* contacts, the complexity of the procedure is then $O(|N|p)$.

The spread of (undirected) information is determined in a similar fashion. We only have to replace "$i \in I$" in the first line of the *do* loop by "*i* or $j \in I$." It is interesting to note that the idea of this procedure is the same used in ego networks described by Liljeros *et al.* (2003) for the spread of sexually transmitted infections.

Given the example in Table 1, the above procedure determines the rooted trees shown in Figures 8a – 8f , given that our observation commences at 8 a.m.

(a) Tree rooted at *A*   (b) Tree rooted at *B*   (c) Tree rooted at *C*

(d) Tree rooted at *D*   (e) Tree rooted at *E*   (f) Tree rooted at *F*

Figure 8

The tree rooted at *A* in Figure 8 has a broken line from node *B* to node *C*, as the message that originated at *A*, was already received by node *C*. The graphs in Figure 8 allow a number of possible conclusions. Among them we have either

- *A* is the leader and nodes *A* through *E* are involved in the cell, or
- *B* is the leader and nodes *A* and *D* are unrelated, or
- *D* is the leader and nodes *A* and *B* are unrelated, or
- *F* is the leader and nodes *B*, *C*, & *E* are unrelated.

Before closing, we should also indicate the sensitivity of the above approach. In order to do so, suppose that we work with the same set of data, but now let our observations commence at 8:06 rather than six minutes earlier. The resulting rooted trees are shown in Figure 9a – 9f.


(a) Tree rooted at *A* (b) Tree rooted at *B* (c) Tree rooted at *C*

(d) Tree rooted at *D* (e) Tree rooted at *E* (f) Tree rooted at *F*

Figure 9

Note that the conclusions, and particularly the emergence of node *A* as a leader, has changed: at this time, we may conclude that

- either *B* is the leader with nodes *A* and *D* unrelated, or
- *F* is the leader with nodes *B*, *C*, and *E* unrelated.

We hypothesize that the true leader will emerge after repeating this model a few times on the communications between these suspects – either *F* or *B* (or both, if they are heading different parts of a planned operation) will appear as the root of most of these directed trees and can be reasonably assumed to

be the leader.

3.3 Model 3: Changes of the degree of separation over time

Our final approach to detection is based on a dynamic analysis of static communication data in terrorist networks. In order to illustrate the same, recall that from a formal point of view, social networks are represented as a set of nodes $N = \{1, 2, ..., n\}$ with each node representing one actor (suspected terrorist or other individual), while binary connections are either represented by the set of (undirected) edges $E = \{e_{ij}:, i, j \in N\}$ or the set of (directed) arcs $A = \{a_{ij}: i, j \in N\}$. Many of the papers in the open literature (see, e.g., Krebs, 2002, Qin, 2005, Yang, 2006, or Rodriguez, 2007) deal with networks that are based on what the authors refer to as "connections" between the actors. Clearly, such connections can denote a variety of things that should be dealt with differently. Among the most popular networks are Krebs's 9/11 networks (Krebs, 2002b, reprinted with permission) which are reproduced here in Figures 10a and 10b. While the network in Figure 10a represents the terrorists' "trusted prior contacts," i.e., their common experiences and backgrounds (growing up in the same village, attending the same school or mosque, participation in the same flight instruction program—the importance of such factors has been emphasized by Ericson, 1981 and Abrahms, 2008), Figure 10b also includes additional edges that indicate contacts that were made during meetings shortly prior to the 9/11 attacks. While each of the networks is static, the two networks, viewed together, present two slices of a temporal continuum, thus providing a dynamic dimension.

Using Milgram's (1967) concept of degrees of separation (where we define the degree of separation between two nodes i and j as the number of arcs or edges between them on the shortest path that connects them), we can determine the degrees of separation between each node in the network, one at a time, and all the other nodes. In other words, we can determine how many nodes are within one, two three, etc., steps from each of the given nodes. Similar to the concept of a Lorenz curve (Lorenz, 1905), we can graph the standings of individuals in a cell at two points in time and can record any dramatic changes. For any given node n_i we plot on the abscissa the proportion of the degree of separation (i.e., the length of the shortest path from n_i to other nodes in relation to the longest path from n_i to any other node in the network), while on the ordinate we plot the percentage of degree of coverage, i.e., the proportion of nodes that are within a given number of degrees of separation from n_i. In one extreme case, all nodes are within one degree of separation from n_i, which means that the curve will start at the origin, but then jump up to 100% and stay there. This indicates that the individual that is represented by n_i is very well connected and extremely close to everybody in the group. In the other extreme case, the curve, starting again at the origin, stays on the abscissa until it reaches 100%, at which point it moves up. This indicates a peripheral individual who is very poorly connected to other members of the group. In order to compare the two snapshots, we can—similar to the Gini index (Gini, 1909)—compute the area that measures the differences in the separation for each individual comparing the two snapshots, i.e., the areas between the two curves. Any individual, whose curve has moved dramatically up and to the left (such as that of M. Atta in this example) indicates the emergence of a leader, and also possible signals an impending attack.

Figure 10a

Figure 10b

We have computed this change for three of the members, *viz.*, M. Atta, S. Alghamdi, & H. Alghamdi for both graphs and plotting the results before the meeting as solid lines and those after the meeting as broken lines in Figure 1, it is very much apparent that M. Atta's status had changed a great deal. In particular, he had clearly emerged as a very well connected individual within the observed group—an indication that he might emerge as a leader, which he ultimately did. We posit that by conducting a similar analysis of static communication data, any change of this nature should put up a red flag of detection. As Krebs remarked in Bohannon (2009), Atta's status as the ringleader was apparent by investigating any of the usual graph-theoretical measures. While this is undoubtedly true, a dynamic analysis such as that suggested in this paper, allows counterterrorism officials to see his status emerging over time, thus allowing the possibility of appropriate interventions during this emergence period.
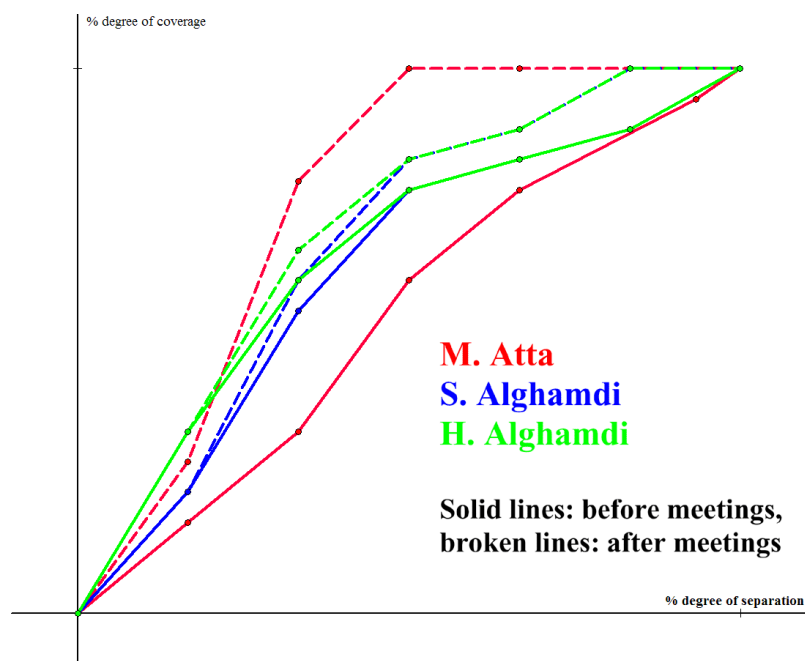
% degree of coverage

M. Atta
S. Alghamdi
H. Alghamdi

Solid lines: before meetings,
broken lines: after meetings

% degree of separation

Figure 11

4. Summary, Extensions and Limitations

The contribution of this paper consists of three techniques that can be used to automatically detect potential terrorist cells by using metadata of communications. The dynamic analysis allows not only to detect cells, but also the standings of individuals within the cell, the second technique outlines some of the graph structures that may indicate terrorist activities (a technique that can be coupled with the dynamic analysis presented earlier) while the third method compares two static snapshots taken over time (preferably before and after an important event) and compare these. The results of the latter can indicate emerging leaders or other dramatic changes.

Some extensions of the analysis present themselves. For one, recent events have already caused the "terrorist community" to make some changes. In particular (Fishel, 2013), terrorists have already started to move away from fast, but not secure, electronic communications. Some even have—as Bin Laden before them—started using trusted couriers they know personally (even though, in Bin Laden's case, this was what ultimately caused his demise). While multimode communication can be incorporated in the framework presented in this paper, it must be automatically observable, which personal courier services are not.

With regards to the limitations of the approaches described in our paper, an important issue is that automatic methods actually identify phone numbers or computers and the people they are registered to, but not the individuals who actually send the message. Even though it is unlikely that terrorists use public computers, e.g., in libraries, this may present a problem. A similar issue occurs with single-use cell phones.

As an aside, another implementation feature of our model arises from the fact that although the models proposed in this paper do not specifically require it, they will work most efficiently in cases where counterterrorism officials have already determined that there is a high probability of the existence of

terrorists and further, that the communication being examined relates primarily to terrorism related activities. While our models will work in all other cases also where this is not true, the collection of metadata from all communication technologies of even one country, or city, may impose an infeasible resource requirement.

There are, of course, problems beyond the detection phase discussed in this contribution. The use of anonymity systems such as TOR (TOR, 2014), the use of coded language, or the employment of different communication methods may obliterate tracks and make detection more difficult. These issues are, however, beyond the scope of this paper.

Finally, as far as network structures are concerned, it can be expected that terrorist organizations, flexible as they are, can and will adapt to counterterrorist measures. One aspect in which they can adapt is to change the structure of the cell networks, so as to minimize the impact of interference from law enforcement. One step into the direction of analyzing different network structures and their performance (albeit from the point of view of counterterrorism authorities) was made by Gutfraind (2010) and, more recently, Krebs (2014).

Acknowledgments

About the authors:

H.A. Eiselt, *Faculty of Business Administration,University of New Brunswick,Fredericton, NB   E3B 5A3,CANADA,*haeiselt@unb.ca

J. Bhadury, *Bryan School of Business and Economics,University of North Carolina at Greensboro,Greensboro, NC 27402-6170,USA,*joy_bhadury@uncg.edu

References

Abrahms M (2008) What terrorists really want: terrorist motives and counterterrorism strategy? *International Security* 32/4: 78-105.

Baker WE, Faulkner RR (1993) The social organization of conspiracy: illegal networks in the heavy electrical equipment industry. *American Sociological Review* 58/6: 837-860.

Bohannon J (2009) Counterterrorism's New Tool: "Metanetwork" Analysis. Science 325 (No. 5939), July 24, 2009: 409-411.

Bonacich P (1987) Power and centrality: a family of measures. *American Journal of Sociology* 92/5: 1170-82

Carley KM (2003) Dynamic network analysis. In *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, Breiger R, Carley K, Pattison Ph (eds.) Committee on Human Factors, National Research Council, National Research Council. Pp. 133–145, Washington, DC. http://ooobgy.googlecode.com/svn-history/r208/trunk/biyesheji/docs/2009Institute_NA_Track_Carley_2003_dynamicnetwork.pdf, last accessed 1/27/2015

Cook KS, Emerson RM, Gilmore MR, Yamagishi T (1983). The distribution of power in exchange networks: theory and experimental results. *American Journal of Sociology* 89: 275-305.

Ericson BH (1981) Secret societies and social structures. *Social Forces* 60/1: 188-210.

Everton SF (2012) Network topography, key players, & terrorist networks. *Connections* 32/1: 12-19.

Farley JD (2007) Toward a mathematical theory of counterterrorism: *The Proteus Mongraph Series*, vol. 1, issue 2, http://www.csl.army.mil/usacsl/Publications/MathematicalTheoryofCT(Farley)(Web).pdf, last accessed 1/27/2015.

Fishel J (2013) Terrorists changing tactics in wake of surveillance program leaks, officials say. http://www.foxnews.com/politics/2013/06/25/terrorists-changing-tactics-in-wake-surveillance-program-leaks-officials-say/, last accessed 1/27/2015.

Freeman LC (1979) Centrality in social networks conceptual clarification. *Social Networks* 1/3: 215-239.

Gini C (1909) Concentration and dependency ratios. English translation in *Rivista di Politica* Economica 87: 769–789 (1997).

Gordijn EH, Stapel DA (2008) When controversial leaders with charisma are efective: the influence of terror on the need for vision and impact of mixed attitudinal messages. *European Journal of Social Psychology* 38: 389-411.

Greenwald G (2013) XKeyscore: NSA tool collects "nearly everything a user does on the internet." *The Guardian* July 31, http://www.theguardian.com/world/2013/jul/31/nsa-top-secret-program-online-data last accessed 1/27/2015.

Gunaratna R, Oreg A (2010) Al Quaeda's organizational structure and its evolution. *Studies in Conflict & Terrorism* 33: 1043-1078.

Gutfraind A (2010) *Mathematical terrorism*. Dissertation, Cornell University, Ithaca, NY.

Holme P, Kim BJ, Yoon CN, Han SK (2002) Attack vulnerability of complex networks. *Physical Review* E 65/5: 1-14.

Holme P, Saramäki J (2012) Temporal networks. *Physics Reports* 519: 97-125.

Hutchins CE, Benham-Hutchins M (2010) Hiding in plain sight: criminal network analysis. *Computational Mathematical & Organization Theory* 16: 89-111.

Kilberg J (2012) A basic model explaining terrorist group organizational structure. *Studies in Conflict & Terrorism* 35: 810-830.

Kitsak M, Gallos LK, Havlin S, Liljeros F, Muchnik L, Stanley HE, Makse HA (2011) Identification of influential spreaders in complex networks. *Nature Physics* 6/11: 888-893

Krebs VE (2002a) Uncloaking terrorist networks. *First Monday* 7/4. http://firstmonday.org/ojs/index.php/fm/article/view/941/863, last accessed 1/27/2015.

Krebs VE (2002b) Mapping networks of terrorist cells. *Connections* 24/3: 43-52.

Krebs VE (2008) Connecting the dots: tracking two identified terrorists. http://www.orgnet.com/prevent.html, last accessed 1/27/2015.

Krebs V (2014) Organizational dynamics…adapting old structures to new challenges. http://www.thenetworkthinkers.com/, last accessed 1/27/2015.

Lederer EM (2004) UN calculates the cost of terrorism. The Age, August 27, 2004, http://www.theage.com.au/articles/2004/08/27/1093518081060.html, last accessed 1/27/2015.

Liljeros F, Edling CR, Amaral LAN (2003) Sexual networks: implications for the transmission of sexually transmitted infections. *Microbes & Infection* 5/2: 189-196.

Lorenz MO (1905) Methods of measuring the concentration of wealth. *Publications of the American Statistical Association* 9/70: 209–219.

MacGinty (2010) Social network analysis and counterinsurgency: a counterproductive strategy? *Critical Studies on Terrorism* 3/2: 209-226.

Medina RM (2012) Social network analysis: a case study of the Islamist terrorist network. *Security Journal* 27/1: 97-121.

Medina R, Hepner G (2008) Geospatial analysis of dynamic terrorist networks. pp151 – 167 in *Values and Violence: Intangible Aspects of Terrorism*, Karawan I, McCormack W, Reynolds SE (eds.). Springer-Verlag, Berlin.

Milgram S (1967) The small world problem. *Psychology Today* 2: 60-67.

Morselli C, Giguère C, Petit K (2007) The efficiency/security trade-off in criminal networks. *Social Networks* 29: 143-153.

Oliva OI (2005) Targeting terrorist leaders: the Peruvian Untouchables experience. MSc thesis, Naval Postgraduate School, Monterey, CA. http://www.dtic.mil/dtic/tr/fulltext/u2/a443352.pdf, last accessed 1/27/2015.

Price BC (2012) Leadership decapitation and the end of terrorist groups. *International Security*. Policy Brief, Belfer Center for Science and International Affairs,

Harvard Kennedy School, May 2012. http://belfercenter.ksg.harvard.edu/files/price_policybrief-final-june-2012.pdf, last accessed 1/27/2015.

Qiang Q, Nagurney A (2008) A unified network performance measure with importance identification and the ranking of network components. *Optimization Letters* 2: 127-142.

Qin J, Xu JJ, Hu D, Sageman M, Chen H (2005) Analyzing terrorist networks: a case study of the global Salafi Jihad network. In: Kantor P, Roberts F, Wang FY, Merkle RC (eds.) *Intelligence and Security Informatics.* Proceedings of the IEEE International Conference on Intelligence and Security Informatics, Atlanta, GA, May 2005. Lecture Notes in Computer Science 3495, Springer-Verlag, Berlin-Heidelberg.

Ressler S (2006) Social network analysis as an approach to combat terrorism: past, present, and future research. *Homeland Security Affairs* II/2: 1-10, http://www.hsaj.org/?article=2.2.8, last accessed 1/27/2015.

Rodríguez JA (2007) The March 11th terrorist network: in its weakness lies its strength. https://sites.google.com/site/estudiospoderprivilegio/seguridad, last accessed 1/27/2015.

Sageman M (2008) *Leaderless Jihad*. University of Pennsylvania Press, Philadelphia, PA.

Sentinel Visualizer (2013) Telephone call data records analysis software. http://www.trackingthethreat.com/LinkAnalysis/telephone_logs/call_data_records.htm, last accessed 1/27/2015.

Sparrow MK (1991) The application of network analysis to criminal intelligence: An assessment of the prospects. *Social Networks* 13: 251-274.

Takala T (2005) Charismatic leadership and power. *Problems and Perspectives in Management* 3/2005: 45-57. Available online at http://www.businessperspectives.org/journals_free/ppm/2005/PPM_EN_2005_03_Takala.pdf, last accessed 1/27/2015.

The Times-Picayune (2014) Google must "forget" some search request info, European Court says.http://www.nola.com/science/index.ssf/2014/05/google_must_remove_some_search.html, last accessed 1/27/2015.

TOR (2014) https://www.torproject.org/, last accessed on 1/27/2015.

United States Department of Defense (2001) Transcript of bin Laden Video Tape. December 13. www.defenselink.mil/news/Dec2001/d20011213ubl.pdf, last accessed 1/27/2015.

Varden JD (2011) Targeting terrorist leaders: a case study. MSc thesis, Naval Postgraduate School, Monterey, CA. http://edocs.nps.edu/npspubs/scholarly/theses/2011/March/11Mar_Varden.pdf, last accessed 1/27/2015.

Weber M (1922) Die drei reinen Typen der legitimen Herrschaft. *Preussische Jahrbücher* 187/1-2. Translated by H Gerth 1958 in *Berkeley Publications in Society and Institution* 4/1: 1-11.

Xu J, Chen H (2008) The topology of dark networks. *Communications of the ACM* 51/10: 58-65.

Yang CC, Liu N, Sageman M (2006) Analyzing the terrorist social networks with visualization tools. Mehotra S, Zeng DD, Chen H, Thuraisingham B , Wang F-Y (eds.) *Intelligence and Security Informatics*. Proceedings of the IEEE International Conference on Intelligence and Security Informatics, ISI 2006, San Diego, CA, USA, May 23-24, 2006, *Lecture Notes in Computer Science* 3975: 331-342, Springer-Verlag, Berlin-Heidelberg.