THE UNIVERSITY OF ADELAIDE AUSTRALIA

SUB CRUCE LUMEN

The University of Adelaide
School of Economics

Research Paper No. 2013-08
May 2013

The Small World of 9/11 and the Implications for Network Dismantlement Strategies

Virginie Masson and Kelsey Wilkins

The Small World of 9/11 and the Implications for Network Dismantlement Strategies

Virginie Masson and Kelsey Wilkins
School of Economics, University of Adelaide SA 5005

April 19, 2013

Abstract

We retraced the development of the network of those who participated in the 9/11 attacks through four stages: 1998-99, December 2000, May 2001 and August 2001. We established that throughout its development, the network had the characteristics of a small world. The implications of this result pointed towards an easily detectable but difficult to dismantle network due to its large clusters. We then assessed the performances of traditional measures of network strength and node centrality. We found that although betweenness surpasses all other measures for all stages, we could improve its performance. The new measure, termed the Jenga index, proved to perform best through all stages.

JEL Classification Codes: F52, D74, D85.
Keywords: Terrorism, Counterterrorism, Social Network Analysis.

1 Introduction

The interest in the 9/11 terrorist attacks and the hunt for and sentencing of their perpetrators remains strong, as attested by the recent release of the movie "Zero Dark Thirty". Unfortunately, the pre-trial hearings of Khalid Sheikh Mohammed, the alleged mastermind, and four of his co-accused stalled earlier this year, delaying answers about how attacks of such destruction could be executed.

In this paper, we focus on the evolution of the network of those closely involved in the 9/11 attacks. We start our analysis by looking at family ties.

We then map the relationships of the terrorists involved through four stages between 1998 to 2001. For each stage, we analyze the characteristics of the network and find that throughout its development it presents a small-world topology, as defined by Watts and Strogatz (1998). This result is intriguing as small world networks usually emanate from casual social interactions rather than carefully planed contacts. It also challenges the idea of hierarchy one would intuitively expect from such an organization and could explain why, despite numerous leads, the network of those involved remained uncovered. Furthermore, we compute and compare traditional measures of network strength and node centrality for each stage. One measure, betweenness, clearly outperforms others at all stages of the network's development. However, it can undermine agents who despite their low number of connections are essential to the functioning of the network. We thus propose a refined measure of betweenness, termed the Jenga index. Defined as the ratio of betweenness to degree, it assesses the role an agent plays globally rather than locally. Its performance surpasses all others' in particular during the intermediate stages of the network's development.

The first application of network analysis to terrorist cells, by Krebs (2002), also maps the terrorist network responsible for the 9/11 attacks. However, the network we present in the final stage differs, due to new public release of information. Jordan et al. (2008) and Azad and Gupta (2011) also attempt to analyze terrorist cells through networks. Jordan et al. (2008) use the Madrid Bombings in 2004 as a case study for analyzing the strength, weaknesses and influence of grass-root jihadist networks, while Azad and Gupta (2011) model the Mumbai terrorist cell.

Sparrow (1991) was one of the first in a wide field of disciplines to assess the applicability of network theory to the analysis of criminal networks. He discusses the merits in great detail; the ease of information conveyance through graphing techniques, the ability to identify key agents and relationships, and the fact that strategic analysis of criminal networks can save the time and resources associated with following every possible lead in an investigation. While this is the case, he also notes the possible limitations of applying network theory to criminal networks. The size of the organizations, asymmetric information, fuzzy boundaries (the frequent inability to distinguish agents' affiliation with a network) and the dynamic nature of criminal enterprises mean that capturing a complete picture of the network is a challenging task.

The successful use of network theory to analyze terrorists' cells depends heavily on the conceptualization of qualitative information into network theory terms; Helfstein and Wright (2011) offer an example of such mapping for OPSEC (operational security) networks. It also depends on whether the measures created to identify key agents withstand the evolution of the terrorist network. To that effect, Il-Chul and Carley (2007) model the evolution of terrorist networks through simulations of agent movement. They overlay four networks: an agent network (the terrorists), a task network, a location network and a knowledge network. In their model, agents have the ability to change location, acquire knowledge and communicate with others. They find that variations in the parameters of any of the four networks do not alter who the important agents and locations are. Frantz et al. (2009) also examine the impact of changes in network topology on measures of centrality, in particular the role they play in being able to identify key agents. They find that under uncertainty, the topology of a network has a measurable effect on the robustness of agent-level measures of centrality. The magnitude and direction of the effect of topologies on network measures also vary significantly.

Identifying efficient ways to dismantle terrorist networks is key. Since 9/11, the literature concerning the dismantlement and destabilization of terrorist cells through network theory has proliferated. Farley (2003) suggests the use of order theory in breaking al-Qaeda cells. His analysis however involves the assumption of a hierarchical network where no two agents have the same rank and all agents have an equal likelihood of being captured. Carley et al. (2002) propose the destabilization of networks based on the removal of leader agents, defined as those with the highest cognitive load. They describe destabilization, rather than dismantlement, as having three key indicators; a reduced rate of flow of information, the inability for those in the network to reach a consensus, or a lower degree of effectiveness. In their model, they allow the network to be dynamic by adapting to the removal of a single agent through the addition or removal of connections. They conduct their analysis on both hierarchical and decentralized networks, and find that removal of agents in decentralized networks is less likely to impact functionality than in hierarchical networks. Carley (2006) extends models of destabilization by exploring methods that inhibit the adaptability of agents. Through experiments with adaptive random and cellular network models, Carley finds that the impact of different destabilization strategies depends largely on the structure of the network. Cellular networks, whereby agents tend to cluster in cells, with few links between cells, prove to be the most difficult to destabilize. This corroborates our intuition that by adopting, purposely or not, a small world topology throughout its development, the 9/11 network was somewhat too intricate to destabilize.

The paper is organised as follows: the next section presents the 9/11 network evolution, starting with family ties. We then introduce some basic network terminology in Section 2. Section 3 explores the similarities between the 9/11 network and small world topology, while Section 4 discusses the performances of traditional measures such as degree, clustering coefficient, point strength and betweenness, and compares them to the performance of the Jenga index. Section 5 concludes.

2 The Development of the 9/11 Network

Using publicly available information[1], we retraced the evolution of the 9/11 network from 1998 to 2001. The criterion we set for inclusion in the network (to address the issue of fuzzy boundaries) is that agents required active participation in the planning or execution of the attacks, with close to full knowledge of the plot. There were many additional terrorists or members of the Muslim community both in the US and Middle East, who assisted agents involved in the plot, but who had no or incomplete knowledge of what was to come. These individuals quite often assisted non-extremists in Muslim communities as well, and were thus deemed inappropriate to include in the network.

As stated earlier, the static final stage of the 9/11 terrorist network was first mapped by Krebs in 2002, using publicly available data. As more information regarding the events was released, we could include further familial and financial connections. In addition, we could link agents who, on record, shared a residence or contacted one another over an extended period of time (over one month).

From Krebs' original static network of 37 agents, we removed six agents and added two. The two agents that were added, Khalid Sheikh Mohammed (KSM is the accepted acronym which will henceforth be used) and Mohammed Haydar Zammar, played important roles in both the evolution and execution of the plot. KSM is the suspected mastermind behind the 9/11 plot who initially proposed the idea to al-Qaeda and Bin Laden. Mohammed Haydar Zammar was a member of the Hamburg al-Qaeda cell from which much of the planning and finance was sourced. The six agents that were removed have since been cleared of any direct involvement in 9/11, or at best their connection with any of the plotters was deemed circumstantial.

We begin our analysis with the identification of family and tribal ties, i.e. links present before individuals became involved in terrorism. A chronological approach then sees the development of the network captured in four main stages: 1998-1999, December 2000, May 2001 and the final stage in August 2001.

2.1 Family ties

The non-familial longstanding ties between agents were formed through tribal affiliations, school relationships or battles fought together. Figure 1 captures these familial ties and longstanding relationships that were present before the 9/11 plot came to be.

While this network existed before the 9/11 attacks were planed and hence provides almost no insights as to how the plan evolved, it is still worth acknowledging the origins of those connections.

2.2 Stage one: 1998-1999

This is the period during which the 9/11 plot first began to develop. Prior to moving to Germany in the late nineties, three of the pilots: Mohamed Atta, Marwan Al-Shehhi and Ziad Jarrah did not have an affiliation with any terrorist groups or extremists. Once there, Mohammed Haydar Zammar persuaded them to join the al-Qaeda cell in Hamburg. The three then quickly formed links with other terrorist agents in the Hamburg cell, and others who would become part of the plot later on. Members of the Hamburg cell and greater Muslim extremist community associated and socialised together frequently.

During this period, Mohamed Atta, Ziad Jarrah and Said Bahaji moved into an apartment together. Many other agents, including members of the Hamburg cell and other known terrorists, also lived there or visited frequently. Figure 2 shows the network at this stage of development; the dense component on the left side of the network is the Hamburg cell, which was under German and US surveillance. In late 1999, as various agents trained in terrorist camps together in Afghanistan and Turkey, more connections were established.

2.3 Stage two: December 2000

By December 2000, the 9/11 plot is well into planning stages, with the recruitment of additional agents and the financing of the training and travel of most of the agents. Four new agents have now joined the network, two of whom will become hijackers; the number of links within the network has increased from 62 to 82. At the beginning of 2000, KSM met with high level al-Qaeda agents as well as hijackers Nawaf Alhazmi, Khalid Al-Midhar and Salem Alhazmi in Malaysia. This meeting was monitored by Malaysian Authorities who passed this information onto their US counterparts. Throughout 2000, many of the terrorist agents travelled throughout Asia and the Middle East for training, establishing relationships with other jihadists. Some even arrived in the US as early as January 2000 and began living with others. Figure 3 captures the network at this stage of development.

2.4 Stage three: May 2001

May 2001 signified the beginning of capital movement in the plot. KSM and al-Qaeda financiers had begun to facilitate the arrival of the hijackers into the United States. Some agents were already there by this time and were living and working together. Many of the terrorists flew across the country several times, often together, in surveillance or practice runs. It was also common that a third party, a financier terrorist, paid for the airfares. These created easily identifiable paper trails that lead to the establishment of more links, shown in Figure 4.

2.5 Final stage: 9/11

The final map of the network was in place by August 2001. All necessary agents were in the United States, most living together in clusters. Some had moved around during the months leading up to 9/11, living with several other agents across the country. The period between May and September 2001 is when the network became the most visibly financially active and many more links were established through large transactions. The final network is given in Figure 5.

Thus far, we mapped the relationships between those involved in the 9/11 attacks in four stages. We now proceed to analyse the network at each stage of its development; this starts with the introduction of some basic network terminology.

3 Network Terminology

We represent a terrorist network using a graph $G(\mathcal{N}, \mathcal{L})$, where the set of vertices or nodes $\mathcal{N} = \{n_1, n_2, ..., n_N\}$ represent the terrorist agents and $\mathcal{L}$ is the set of all links or edges between agents. The existence of a link in the graph represents a relationship between two agents; the relationship may be of a familial, financial or friendship nature.

3.1 Network Characteristics

A path from agent i to some other agent j in a network is an altering sequence of nodes and edges in which no node or edge is retraced throughout the journey. It captures the number of links and connections agent i must exploit in order to reach agent j (to convey information, resources, etc). A path of minimal length between two agents i and j, denoted d_{ij}, is called the shortest path or geodesic.

A graph or network is referred to as connected if there exists a path from agent i to agent j for all $i \neq j$ in the graph. Connectedness is paramount for terrorist networks, in particular in the later stages of their development; without it, precise timing and execution of attacks are likely to fail.

A subgraph of $G(\mathcal{N}, \mathcal{L})$, denoted by $G'(\mathcal{N}', \mathcal{L}')$, is defined as a graph such that $\mathcal{N}' \subseteq \mathcal{N}$ and $\mathcal{L}' \subseteq \mathcal{L}$. If G' contains all the links from G that connect any two agents in $\mathcal{N}'$, then the subgraph G' is said to be induced by $\mathcal{N}'$ (Boccaletti et al. (2006)). Furthermore, a subgraph G' is maximally connected if it cannot be extended without violating the aforementioned property. Such a subgraph is referred to as a component within the original graph $G(\mathcal{N}, \mathcal{L})$. The higher the number of components in a terrorist network, the less likely the network is able to function.

Another characteristic of interest is the average path length of a network. It is defined as the average number of hops needed to get from one agent to another. Formally, the average path length is the average of all the geodesics in a network:

$$L = \frac{1}{N(N-1)} \sum_{j,i \in \mathcal{N}, j \neq i} d_{ij}$$

Hence, a long average path length is based on the need to distance cells for security purposes, as it lowers the risk of detection from a single agent capture. However, this also means a lower efficiency level due to delays in transmissions of information or resources from one agent to another.

Finally, an agent's clustering coefficient provides a measure of connectivity between his neighbours. In particular, if agent i has a clustering coefficient of unity, then all of agent i's neighbours are also neighbours of each other. The clustering coefficient of an agent is formally defined by

$$C_i = \frac{2e_i}{k_i(k_i-1)}$$

where e_i is the number of edges in the subgraph G'_i (the subgraph induced by node i) and k_i is the degree of node i (Boccaletti et al. (2006)). Networks with a high clustering coefficient average, such as small world networks are quite easily detectable; they are however more robust to dismantlement.

3.2 Network Strength

The literature commonly focuses on two measures of strength that we introduce below: connectivity and point strength. At a macro level, a network's connectivity is a measure of its overall strength. Closely related to order theory, the connectivity of a network is the minimum number of agents that must be removed for the network to become disconnected. The limitation of such measure is that in some cases, it is possible that the removal of one agent induces a single agent to become disconnected from the larger component. So while the theoretical implication of a low connectivity is a weak network, its practical significance is not always as relevant.

At a micro level, the point strength of an agent in a network is the additional number of components induced by the removal of that agent (Capobianco and Molluzzo (1980)). In relatively dense networks, most agents will have a point strength of zero- the removal of a single agent does not suffice to create an additional component. However, in the rare instances where an agent does have a point strength greater than zero, there is no assurance that this agent has an important function in the network. For instance, an agent may have a point strength of one, but the additional components that are created by its removal may be two single isolated agents, leaving the rest of the network intact and functional. The odds are small that these now isolated agents were crucial to the function of the network. While point strength and connectivity are useful measures of network strength in a general context, their use alone limits the capacity to severely reduce the function and efficiency of terrorist networks in a practical sense.

3.3 Node Centrality

Carley (2003) and Memon and Larsen (2006) suggest methods of network destabilization based on network measures such as degree and betweenness. However, targeting agents with the highest degree may not result in destabilizing a terrorist network, as in fact, logic dictates that terrorist organizations would want to lower the connections of the most important agents.

Betweenness is considered one of the standard measures of node centrality. It indicates the influence or control an agent has over information or resource flows in the network. The betweenness of agent i is defined as the fraction of geodesics in the network that pass through agent i. More formally, the betweenness of agent i is

$$B_i = \sum_{j,k \in \mathcal{N}, j \neq k} \frac{d_{jk}(i)}{d_{jk}}$$

where d_{jk} is the number of geodesics connecting nodes j and k, and $d_{jk}(i)$ is the number of those geodesics that pass through node i (Boccaletti et al. (2006)). This is then summed over the entire network for all nodes j and k so as to determine exactly how many geodesics pass through node i. There can be more than one geodesic connecting any two nodes j and k. However, when there is only one geodesic connecting nodes j and k and node i lies on this path, then the above summation increases by 1 (since $d_{jk}(i) = 1$ and $d_{jk} = 1$). Hence, the measure above is of absolute betweenness.

Freeman (1977) developed a relative measure of betweenness to allow for comparisons across different networks. In a connected undirected network, there are $n(n-1)/2$ possible paths connecting pairs of nodes and of these $n-1$ are connected to node i. Therefore the number of possible geodesics on which node i could fall is

$$\frac{1}{2}n(n-1) - (n-1) = \frac{n^2-3n+2}{2}$$

Multiplying the absolute measure of betweenness by the inverse of this scalar provides a measure of relative betweenness between 0 and 1 (Freeman (1977)). Agents with a high betweenness have a relatively high number of geodesics that pass through them. On the downside, this means that an agent with a low clustering coefficient and a high degree can have high betweenness; this results in identifying agents whose roles may matter locally rather than globally.

We introduce a new measure aimed at resolving this issue. Termed the Jenga index[2], it is defined as the ratio of betweenness to degree; it targets agents who are important to the flow of information and resources, but who may possess only few connections, i.e., who are not easily detectable. The Jenga index allows the importance of an agent to be assessed in the global context of the network, and limits the weight given to this agent for his role within his immediate neighbourhood. In other words, it describes how significant an agent is in connecting large but distant components which may otherwise be unreachable.

4 A Small World

In this section, we investigate the topology of the 9/11 network throughout its development. We know from Helfstein and Wright (2011) that in its last stage, it is neither OPSEC nor scale-free [3]. As it turns out, the 9/11 network exhibits strong characteristics of a small world network, in all of its development stages.

Small world networks present a higher clustering coefficient and similar average path length relative to a random graph (Watts and Strogatz (1998)). They can be created through a random rewiring procedure. The process begins with a regular ring lattice (a graph in which each node has the same degree), with any edge being rewired at random with probability p. The intuition of the rewiring procedure is that it connects particular areas in the network which would otherwise be far apart. Hence, the structure of small world networks lies in the interval between completely regular and completely random networks' structures.

Small world networks with the same number of nodes and similar number of edges were generated for each stage as per Watts and Strogatz' procedure described above.[4] To obtain small world networks whose degrees are similar to the actual networks, we had to start with regular networks with relatively low degrees, as the rewiring procedure increases the average degree. Also, due to the random aspect of the rewiring process, simulated small world networks have additional links: 20 additional links in stage 1, but only between 2 and 10 additional links in the remaining three stages.

Table 1 below is a summary of the characteristics of the 9/11 network and simulated networks for each stage. We denote n the number of agents; random, small world and regular columns present the metrics for simulated networks.

The average path length was calculated based on the assumption of a connected network. Thus, in the early stages of evolution, the small components were excluded from the calculation, using only the largest connected component instead. More precisely, the network in stage 1 has 27 agents with 62 connections between them; it also has four components, of which three were removed (inducing the removal of 7 agents) to compute the average path length. The network is still disconnected in the second stage, but only one component consisting of two agents was removed for the calculations.

At each stage of its development, the 9/11 network has a higher clustering coefficient than the simulated small world and random networks; it also has an average path length roughly on par with what is expected in a small world. This means that the 9/11 network presented the characteristics of a small-world throughout its development, making it easy to detect but difficult to destabilize or dismantle.

5 Measure Performances

This part of our analysis addresses the issue of the 9/11 network's dismantlement by assessing the performances of the measures presented earlier. Our definition of performance includes two aspects: First, the number of components after the removal of one or more agents. Second, the size of each components. Agents are removed based on their rankings from the measure we evaluate the performance of.

We first consider the performance of point strength. If we were to rely on point strength alone, a maximum of two agents would be targeted for removal at any stage of the network's development, as most agents have a point strength of zero. In particular, in stages 2 to 4, the removal of Hani Hanjour suffices to create an additional component, constituted by a pair of agents. It is thus obvious that the large component of 9/11 will continue to function relatively unimpaired. A similar observation holds regarding the removal of Hamza Alghamdi in stage 2, and the removal of Nawaf Alzhami and Waleed Alshehri in stage 1. Thus, even though the use of point strength does help disconnect the network, it does not impact much the network's overall functionality.

Point strength is not the only measure whose performance is unconvincing; degree and clustering coefficient are equally disappointing. As a high number of agents share a similar degree or clustering coefficient, it is impossible to sequentially remove those agents and assess the effect of each removal on the network.

In contrast, betweenness and the Jenga index offer a clear ranking, with ties only appearing at rank 13 in the worst case. Both measures therefore provide a specific order of removal. Table 2 shows, throughout the network's development, the ranking of the 15 agents with the highest Jenga index in the final stage.

It is interesting to note that KSM, the suspected mastermind, frequents the top five Jenga Index rankings through all first three stages. However, by August 2001, as the plan is already finalised, his significance is reduced.

The partial assessment of relative performance of betweenness and the Jenga index is summarized in Table 3; the size of each component remains to be evaluated to fully assess both performances.

Each cell in Table 3 represents the number of agents needed to be removed to obtain the number of components indicated in the left-hand side of the Table. For example, in December 2000, it is possible to disconnect the network into five components by either removing the three agents with the highest Jenga index ranking, or the five agents with the highest betweenness. As a convention, we write zero to indicate the number of components in the network before any removal.

At first glance, the Jenga index performs better than betweenness; fewer agents need to be removed to obtain the same number of components. We still need nonetheless to assess the size of each component, especially during the intermediate stages. Indeed, if a component consists of a pair of agents or an isolated agent, then the network remains largely unaffected.

In May 2001, the removal of the five agents with the highest Jenga index results in the network having three components; two large (more than 10 agents) and one pair of agents. Similar size components are obtained by removing the six agents with the highest betweenness. Hence, the Jenga index performs better, as it requires the removal of fewer agents to obtain similar size components.

In December 2000, the difference between both measures is the greatest. Table 4 presents a summary of the components' size obtained by using either measure. Large, Medium, Pair and Isolated refer to the size of each component. Since no small components with between 3 and 5 agents were created, this category has been omitted.

The first line provides the size of each of the four components that have resulted once agents have been removed. In this case, the agents removed were the same, whether one used the Jenga index or betweenness; hence, the (size of the) four components are the same for both measures.

When looking at the case with five components (2nd line), betweenness seems to perform better as the network has one large and one medium size components. However, it should be reminded that disconnecting the network into five components requires the removal of five agents with betweenness, and only three with the Jenga index. If five agents were to be removed following the Jenga index, the network would then have six components, including one large and one medium.[5] Hence, we conclude that the Jenga index performs better than betweenness in the second stage.

Overall, the Jenga index consistently performs better than betweenness, and in particular during the second stage of the network's development. In that stage, Mohammed Atta's betweenness measure is overinflated by a large degree and relatively small clustering coefficient. Hence, the role he plays locally is given too much weight, while Hani Hanjour's small degree allows him to remain relatively unnoticed.

6 Conlusion

This paper retraces the development of the network of those involved in the 9/11 attacks. It presents a dynamic view of the network through four stages and shows that the network's structure is close to a small world topology. This is a rather surprising result as those networks often emerge through casual social interactions.

We also highlight the unwanted effect a high degree and a low clustering coefficient combined have on betweenness, and propose a solution to this issue. Defined as the ratio of betweenness to degree, the Jenga index records the best performance in all stages of the network's development. Its ability to identify agents who are key in connecting distant parts of the network is really noteworthy.

Notes

[1]Please refer to the References section for a detailed list.

[2]The term "Jenga index" comes from the game of the same name in which blocks must be iteratively removed from a tower until the tower finally collapses.

[3]Although they consider the network presented by Krebs (2002), we find similar results for the network we presented in the final stage.

[4]The probability of an edge being randomly rewired was set to 0.1.

[5]The removal of only four agents would also lead to this result.

References

The complete 9/11 timeline. The History Commons.

(2002). Joint Inquiry into Intelligence Community Activities Before and After the Terrorist Attacks of September 11. Technical report, U.S. Senate Select Committee on Intelligence and U.S. House Permanent Select Committee on Intelligence.

(2003). 11 September: The Plot and the Plotters. Technical report, Central Intelligence Agency.

(2003). *United States V. Zacarias Moussaoui Indictment*, United States District Court for the Eastern District of Virginia.

(2004). The 9/11 Commission Report. Technical report, 9/11 Commission.

Azad, S. and A. Gupta (2011). A quantitative assessment on 26/11 Mumbai attack using social network analysis. *Journal of Terrorism Research 2*, 4–11.

Boccaletti, S., V. Latora, Y. Moreno, M. Chavez, and D. U. Hwang (2006). Complex networks: Structure and dynamics. *Physics Reports 424*, 175–308.

Capobianco, M. F. and J. C. Molluzzo (1980). The strength of a graph and its application to organizational structure. *Social Networks 2*, 275.

Carley, K. M. (2003). Destabilizing terrorist networks. In *Proceedings of the 8th International Command and Control Research and Technology Symposium: Evidence Based Research*, National Defense War College, Washington DC.

Carley, K. M. (2006). Destabilization of covert networks. *Computational & Mathematical Organization Theory 12*, 51–66.

Carley, K. M., J.-S. Lee, and D. Krackhardt (2002). Destabilizing networks. *Connections 24*, 79–92.

Farley, J. D. (2003). Breaking Al Qaeda cells: A mathematical analysis of counterterrorism operations (a guide for risk assessment and decision making). *in Conflict & Terrorism 26*, 399–411.

Frantz, T. L., M. Cataldo, and K. M. Carley (2009). Robustness of centrality measures under uncertainty: Examining the role of network topology. *Computational and Mathematical Organization Theory 15*, 303–328.

Freeman, L. C. (1977). A set of measure of centrality based on betweenness. *Sociometry 40*, 35–41.

George, M. N. (2007). B-A scale-free network generation and visualization (Matlab).

Gleich, D. (2006, April). MatlabBGL.

Helfstein, S. and D. Wright (2011). Covert or convenient? Evolution of terror attack networks. *Journal of Conflict Resolution 55*, 785–813.

Il-Chul, M. and K. M. Carley (2007). Modeling and simulating terrorist networks in social and geospatial dimensions. *IEEE Intelligent Systems 22*, 40–49.

Jordan, J., F. M. Manas, and N. Horsburgh (2008). Strengths and weaknesses of grassroot jihadist networks: The Madrid bombings. *Studies in Conflict & Terrorism 31*, 17–39.

Krebs, V. E. (2002). Mapping networks of terrorist cells. *Connections 24*, 43–52.

Memon, N. and H. L. Larsen (2006). Practical algorithms for destabilizing terrorist networks. In *Visualising Network Information. Meeting Proceedings*, Copenhagen, Denmark. NATO Research & Technology Agency.

Scheinerman, E. R. (2007). Matgraph: A Matlab toolbox for graph theory.

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

Taylor, A. and D. J. Higham (2008, May). CONTEST: A controllable test matrix toolbox for Matlab.

Thompson, P. (2004). *The Terror Timeline: Year by Year, Day by Day, Minute by Minute: A Comprehensive Chronicle of the Road to 9/11 - and America's Response*. ReganBooks.

Watts, D. J. and S. H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature 393*, 440–442.

7 Figures

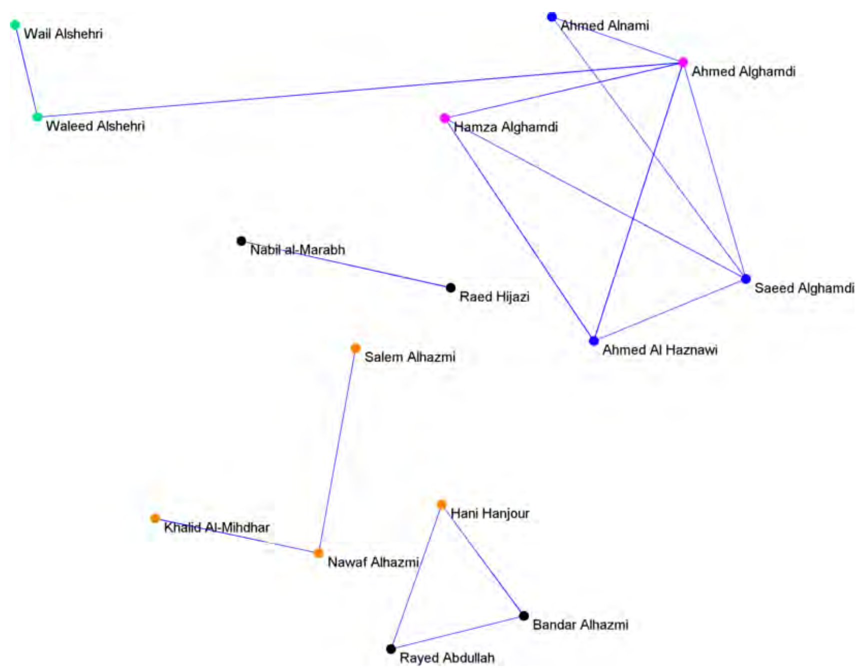

Figure 1: Family ties and longstanding relationships

The green nodes are the hijackers that were aboard flight AA#11 that crashed into the North Tower of the World Trade Centre; the pink nodes represent those terrorists aboard flight AA#175 which crashed into the South Tower; orange nodes are the hijackers who crashed flight AA#77 into the Pentagon; and the blue nodes represent the terrorists aboard United Airlines flight #93 that went down in Pennsylvania. The remaining nodes are associates of the hijackers who were responsible for money, planning and resource procurement within the network. This colour system will be used consistently throughout the paper.

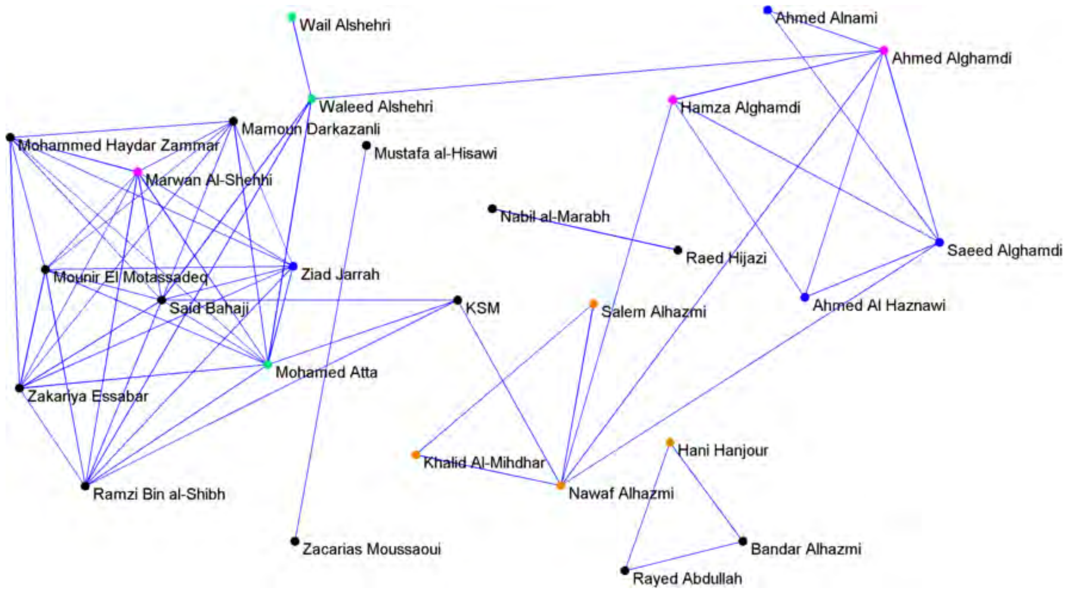

Figure 2: Stage one (1998-1999)

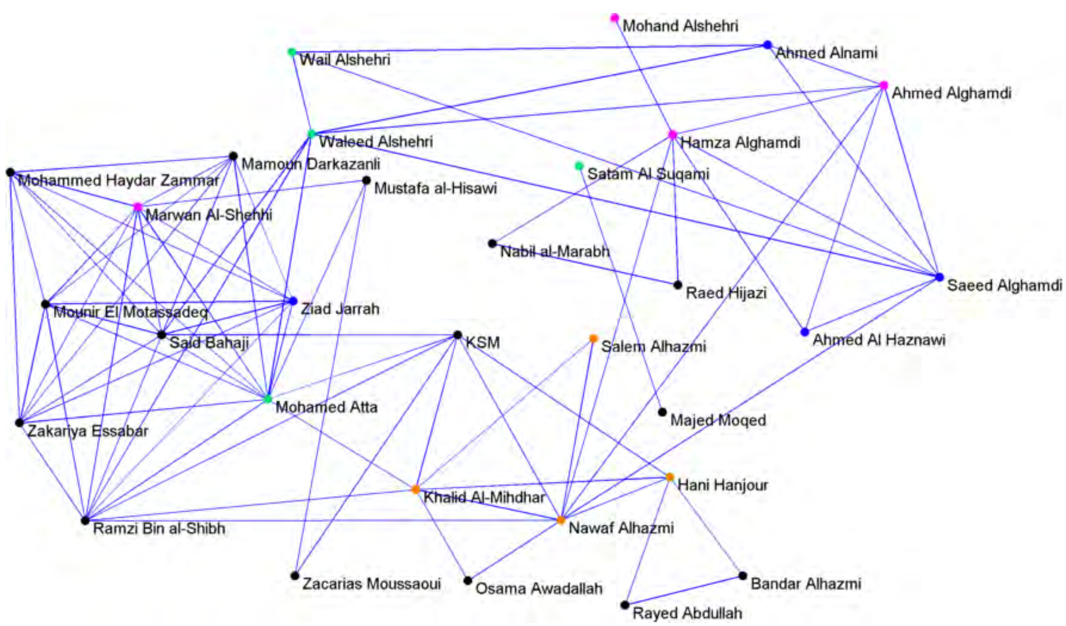

Figure 3: Stage two (December 2000)

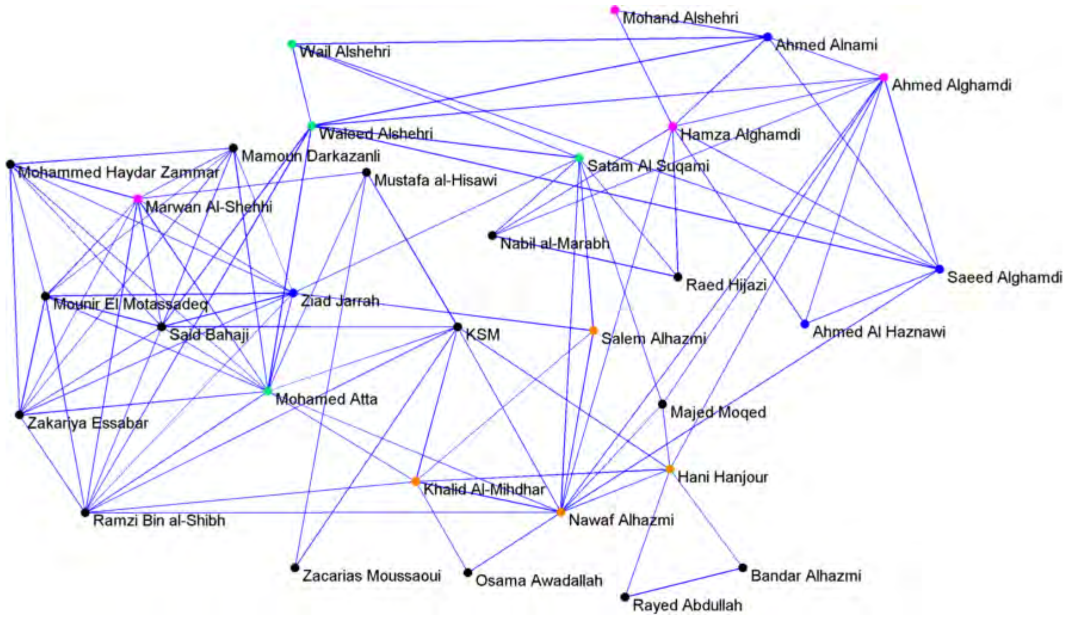

Figure 4: Stage three (May 2001)

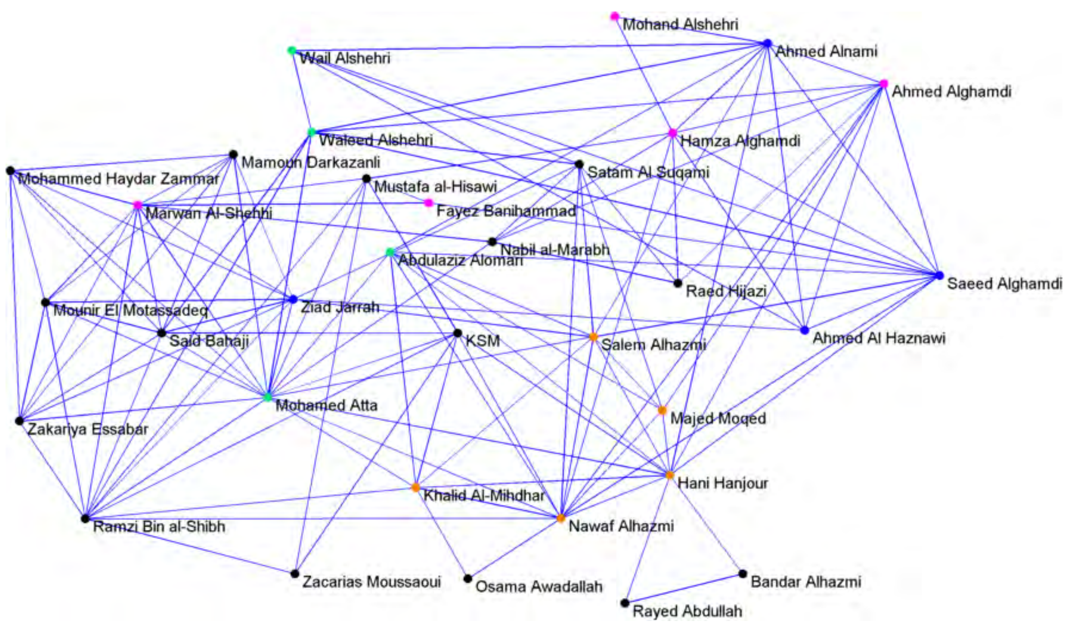

Figure 5: Final network (August 2001)

8 Tables

Table 1: Evolution metrics

Network	9/11	Random	Small World	Regular
Stage One: $n = 27$ (62 links)				
Average Degree	4.59	4.59	6	3
Average Clustering Coefficient	0.6528	0.1774	0.4738	0.6000
Average Path Length	1.6816	1.6048	1.6184	1.8462
Stage Two: $n = 31$ (82 links)				
Average Degree	5.29	5.29	6	3
Average Clustering Coefficient	0.6320	0.1771	0.4638	0.6000
Average Path Length	1.7549	1.5806	1.6793	2.0000
Stage Three: $n = 31$(99 links)				
Average Degree	6.38	6.38	6	3
Average Clustering Coefficient	0.6506	0.2009	0.4714	0.6000
Average Path Length	1.640	1.4937	1.6998	2.0000
Final Stage: $n = 33$ (127 links)				
Average Degree	7.69	7.69	8	4
Average Clustering Coefficient	0.6102	0.2406	0.5034	0.6429
Average Path Length	1.5104	1.4381	1.5350	1.7500

Table 2: Betweenness and Jenga index rankings

AGENTS	Aug. 2001		May 2001		Dec 2000		1998-99	
	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn
Hani Hanjour	1	2	2	3	3	6	19	19
Mohamed Atta	2	1	6	2	7	3	5	5
Nawaf Alhazmi	3	3	1	1	1	1	3	2
Hamza Alghamdi	4	5	8	8	2	2	9	9
Waleed Alshehri	5	9	3	4	4	5	1	1
Ziad Jarrah	6	6	9	9	15	15	13	13
Saeed Alghamdi	7	7	13	13	9	9	8	8
Ahmed Alnami	8	8	12	14	18	18	17	17
Ramzi Bin al-Shibh	9	4	10	10	6	4	7	7
Mustafa al-Hisawi	10	13	15	16	12	13	23	23
Nabil al-Marabh	11	16	18	19	19	19	24	24
Satam Al Suqami	12	12	5	6	19	19	N/A	N/A
Ahmed Alghamdi	13	10	7	7	10	11	4	3
KSM	14	15	4	5	5	7	2	4
Marwan Al-Shehhi	15	10	19	15	14	12	10	10

Table 3: Number of Components after the Removal of Highest Ranked Agents

# SCCS	Aug. 2001		May 2001		Dec 2000		1998-99	
	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn
1	0	0	0	0	N/A	N/A	N/A	N/A
2	1	-	2	2	0	0	N/A	N/A
3	8	8	5	6	-	-	N/A	N/A
4	-	-	8	8	2	2	0	0
5	12	13	-	-	3	5	1	1
6	13	-	-	-	4	6	-	2
7	14	-	-	-	7	-	2	3

Table 4: Size of Components - December 2000

# Components	Large (> 10)		Medium (> 5)		Pair (= 2)		Isolated (= 1)	
	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn	Jenga	Btwn
4	1	1	0	0	2	2	1	1
5	1	1	0	1	3	2	1	1
6	1	1	1	1	3	3	1	1

9 Supplementary Material

9.1 Network Measures

Table 5: Network metrics at stage one (1998-1999) -
Nodes 4, 6, 12, 13, 15, 17 and 23 removed for average path length and betweenness (disconnectedness)

Node	Name	Degree	Clustering Coefficient	Betweenness	Point Strength	Jenga Index
1	Ahmed Al Haznawi	3	1	0	0	0
2	Ahmed Alghamdi	6	0.4	0.4982	0	0.083
3	Ahmed Alnami	2	1	0	0	0
4	Bandar Alhazmi	2	1	0	0	0
5	Hamza Alghamdi	4	0.8333	0.0147	0	0.0037
6	Hani Hanjour	2	1	0	0	0
7	Khalid Al-Mihdhar	2	1	0	0	0
8	Mamoun Darkazanli	7	1	0	0	0
9	Marwan Al-Shehhi	8	0.9286	0.0037	0	0.0005
10	Mohamed Atta	10	0.6667	0.293	0	0.0293
11	Mounir El Motassadeq	8	0.9286	0.0037	0	0.0005
12	Mustafa al-Hisawi	1	0	0	0	0
13	Nabil al-Marabh	1	0	0	0	0
14	Nawaf Alhazmi	6	0.2667	0.5275	1	0.0879
15	Raed Hijazi	1	0	0	0	0
16	Ramzi Bin al-Shibh	8	0.6786	0.1685	0	0.0211
17	Rayed Abdullah	2	1	0	0	0
18	Saeed Alghamdi	5	0.6	0.0476	0	0.0095
19	Said Bahaji	10	0.6667	0.293	0	0.0293
20	Salem Alhazmi	2	1	0	0	0
21	Wail Alshehri	1	0	0	0	0
22	Waleed Alshehri	5	0.3	0.5934	1	0.1187
23	Zacarias Moussaoui	1	0	0	0	0
24	Zakariya Essabar	8	0.9286	0.0037	0	0.0005
25	Ziad Jarrah	8	0.9286	0.0037	0	0.0005
26	KSM	4	0.5	0.3956	0	0.0989
27	Mohammed Haydar Zammar	7	1	0	0	0
	Mean	4.5926	0.6528	0.1423		0.0179
	Variance	9.2507	0.1507	0.0438		0.00028415
	Average Path Length	1.6816				
	Edges	62				
	Density	0.1766				

Table 6: Network metrics at stage two (December 2000) -
Nodes 8 and 24 removed for the calulations of average path length and betweenness (disconnectedness)

Node	Name	Degree	Clustering Coefficient	Betweenness	Point Strength	Jenga Index
1	Ahmed Al Haznawi	3	1	0	0	0
2	Ahmed Alghamdi	6	0.5333	0.0909	0	0.0152
3	Ahmed Alnami	4	0.8333	0.0014	0	0.0004
4	Bandar Alhazmi	2	1	0	0	0
5	Hamza Alghamdi	7	0.2857	0.3491	2	0.0499
6	Hani Hanjour	5	0.4	0.2251	1	0.045
7	Khalid Al-Mihdhar	7	0.4286	0.1508	0	0.0215
8	Majed Moqed	1	0	0	0	0
9	Mamoun Darkazanli	7	1	0	0	0
10	Marwan Al-Shehhi	9	0.75	0.0283	0	0.0031
11	Mohamed Atta	12	0.5	0.2754	0	0.0229
12	Mohand Alshehri	1	0	0	0	0
13	Mounir El Motassadeq	8	0.9286	0.0051	0	0.0006
14	Mustafa al-Hisawi	3	0.3333	0.0187	0	0.0062
15	Nabil al-Marabh	2	1	0	0	0
16	Nawaf Alhazmi	9	0.2778	0.524	0	0.0582
17	Osama Awadallah	2	1	0	0	0
18	Raed Hijazi	2	1	0	0	0
19	Ramzi Bin al-Shibh	10	0.5111	0.2622	0	0.0262
20	Rayed Abdullah	2	1	0	0	0
21	Saeed Alghamdi	7	0.4762	0.1434	0	0.0205
22	Said Bahaji	10	0.6667	0.1097	0	0.011
23	Salem Alhazmi	2	1	0	0	0
24	Satam Al Suqami	1	0	0	0	0
25	Wail Alshehri	3	1	0	0	0
26	Waleed Alshehri	7	0.381	0.245	0	0.035
27	Zacarias Moussaoui	2	0	0.0098	0	0.0049
28	Zakariya Essabar	8	0.9286	0.0051	0	0.0006
29	Ziad Jarrah	8	0.9286	0.0051	0	0.0006
30	KSM	7	0.4286	0.2047	0	0.0292
31	Mohammed Haydar Zammar	7	1	0	0	0
	Mean	5.29	0.632	0.092		0.0113
	Variance	10.28	0.129	0.018		0.00028415
	Average Path Length	1.755				
	Edges	82				
	Density	0.176				

Table 7: Network metrics at stage three (May 2001)

Node	Name	Degree	Clustering Coefficient	Betweenness	Point Strength	Jenga Index
1	Ahmed Al Haznawi	3	1	0	0	0
2	Ahmed Alghamdi	9	0.3611	0.0931	0	0.0103
3	Ahmed Alnami	6	0.5333	0.0287	0	0.0048
4	Bandar Alhazmi	2	1	0	0	0
5	Hamza Alghamdi	8	0.3571	0.0779	0	0.0097
6	Hani Hanjour	7	0.3333	0.1421	1	0.0203
7	Khalid Al-Mihdhar	7	0.4762	0.0425	0	0.0061
8	Majed Moqed	4	0.6667	0.0077	0	0.0019
9	Mamoun Darkazanli	7	1	0	0	0
10	Marwan Al-Shehhi	9	0.75	0.012	0	0.0013
11	Mohamed Atta	13	0.4744	0.1501	0	0.0115
12	Mohand Alshehri	2	1	0	0	0
13	Mounir El Motassadeq	8	0.9286	0.0008	0	0.0001
14	Mustafa al-Hisawi	4	0.5	0.0112	0	0.0028
15	Nabil al-Marabh	4	0.5	0.0058	0	0.0014
16	Nawaf Alhazmi	12	0.2727	0.2451	0	0.0204
17	Osama Awadallah	2	1	0	0	0
18	Raed Hijazi	3	0.6667	0.0021	0	0.0007
19	Ramzi Bin al-Shibh	10	0.5333	0.0624	0	0.0062
20	Rayed Abdullah	2	1	0	0	0
21	Saeed Alghamdi	7	0.5238	0.0311	0	0.0044
22	Said Bahaji	10	0.6667	0.0437	0	0.0044
23	Salem Alhazmi	4	0.5	0.0074	0	0.0019
24	Satam Al Suqami	8	0.1786	0.1073	0	0.0134
25	Wail Alshehri	4	0.6667	0.0041	0	0.001
26	Waleed Alshehri	8	0.3214	0.1153	0	0.0144
27	Zacarias Moussaoui	2	1	0	0	0
28	Zakariya Essabar	8	0.9286	0.0008	0	0.0001
29	Ziad Jarrah	10	0.6	0.068	0	0.0068
30	KSM	8	0.4286	0.1086	0	0.0136
31	Mohammed Haydar Zammar	7	1	0	0	0
	Mean	6.3871	0.6506	0.0441		0.0051
	Variance	9.7785	0.0704	0.0036		3.79E-05
	Average Path Length	1.6398				
	Edges	99				
	Density	0.2129				

Table 8: Network metrics at final stage (9/11)

Node	Name	Degree	Clustering Coefficient	Betweenness	Point Strength	Jenga Index
1	Abdulaziz Alomari	8	0.6786	0.0117	0	0.0015
2	Ahmed Al Haznawi	6	0.5333	0.0176	0	0.0029
3	Ahmed Alghamdi	10	0.4	0.0411	0	0.0041
4	Ahmed Alnami	10	0.4222	0.0512	0	0.0051
5	Bandar Alhazmi	2	1	0	0	0
6	Fayez Banihammad	3	0.3333	0.005	0	0.0017
7	Hamza Alghamdi	9	0.3333	0.06	0	0.0067
8	Hani Hanjour	12	0.4091	0.1479	1	0.0123
9	Khalid Al-Mihdhar	8	0.6071	0.0231	0	0.0029
10	Majed Moqed	6	0.6667	0.0042	0	0.0007
11	Mamoun Darkazanli	7	1	0	0	0
12	Marwan Al-Shehhi	11	0.5455	0.0411	0	0.0037
13	Mohamed Atta	17	0.3676	0.1738	0	0.0102
14	Mohand Alshehri	2	1	0	0	0
15	Mounir El Motassadeq	8	0.9286	0.0011	0	0.0001
16	Mustafa al-Hisawi	7	0.381	0.0343	0	0.0049
17	Nabil al-Marabh	6	0.3333	0.028	0	0.0047
18	Nawaf Alhazmi	13	0.3846	0.098	0	0.0075
19	Osama Awadallah	2	1	0	0	0
20	Raed Hijazi	4	0.6667	0.0016	0	0.0004
21	Ramzi Bin al-Shibh	12	0.4394	0.0604	0	0.005
22	Rayed Abdullah	2	1	0	0	0
23	Saeed Alghamdi	11	0.4182	0.0575	0	0.0052
24	Said Bahaji	10	0.6667	0.0175	0	0.0017
25	Salem Alhazmi	10	0.5111	0.0324	0	0.0032
26	Satam Al Suqami	8	0.2143	0.0367	0	0.0046
27	Wail Alshehri	5	0.6	0.0034	0	0.0007
28	Waleed Alshehri	8	0.3214	0.0479	0	0.006
29	Zacarias Moussaoui	3	1	0	0	0
30	Zakariya Essabar	8	0.9286	0.0011	0	0.0001
31	Ziad Jarrah	11	0.5091	0.0592	0	0.0054
32	KSM	8	0.5357	0.0311	0	0.0039
33	Mohammed Haydar Zammar	7	1	0	0	0
	Mean	7.697	0.6102	0.0329		0.0032
	Variance	12.905	0.0667	0.0017		0.0000098
	Average Path Length	1.51				
	Edges	127				
	Density	0.241				

9.2 Evolution timeline

This section gives a succinct timeline of the relationships established between agents in the lead up to 9/11. The timeline is written using numbers for convenience. Each number corresponds to the terrorist's node number in the final stage of the network (which can be found in the statistics tables).

Family ties

- 9 and 18: Cousins in law
- 18 and 25: Brothers
- 27 and 28: Brothers
- Tribal/family affiliation: 2, 3, 7 and 23

Long histories (Pre-1998)

- 18 and 9 fight in Bosnia together against the Serbs in 1995
- 3 and 28 live together in Florida in 1997 and again in 1999-2000 in Virginia
- 17 and 20 attend terrorist training camp together in Afghanistan in 1992; they then live in Boston together for several years
- 3, 4 and 23 use the same address on licence applications in 1996-1998
- 5 and 8 lived together with 22 in 1997 (pre-existing friends)
- 17 and 20 meet in an Afghanistan training camp in 1980s

1998-1999

- Connection is made between 18, 9 and 25
- 12, 13 and 31 moved to Hamburg in 1998
- They become members of the Hamburg cell featuring 11, 15, 24, 33 and 30

- 13, 24 and 21 moved into an apartment together (13 and 21 have a relationship that predates this, however). 30, 12, 28 and 32 also stay during this time.
- 32 also meets 12 sometime while in Hamburg
- 12 and 33 make several phone calls to each other
- 32 repeatedly visits 13 in Hamburg during 1999
- In April 1999, 31 is married. A photo found from the wedding shows 13, 15, 21 and other unspecified members of the Hamburg cell
- In late 1999, 3, 7, 18 and 23 undertook training together in Turkey
- In December 1999, 32 gives further training to 18 in Afghanistan
- 16 finances 29's international travel within the Dublin al-Qaeda cell

By May 2000

- 18, 9, 21, 32 and attend the Malaysia al-Qaeda meeting in January
- In Spring 2000, 8 was sent to 32 for specialised training

By December 2000

- 8 moved in with 18 and 9 briefly in the Summer of 2000
- 9 and 18 moved into 2's house; 13 and 8 visit frequently.
- 26 and 10 have training in Afghanistan (sometime before November)
- 18 works with 19 at a gas station in Autumn 2000
- 19 also helps 9 adjust to US at same time (speculatively- no date but makes sense)
- 27, 28, 3 and 23 jointly swear an oath to jihad in Spring 200

- 16 transferred a total of $100,000 to the joint account of 13 and 12 in August to September 2000
- Jordanian government informs US that 17, 20 and 7 are connected through phone numbers
- 14 and 7 travel to Kuwait together in October 2000
- 32 sent 29 to Malaysia in Fall 2000 for flight training, then onto the US later

By May 2001

- 25 and 26 living together in early 2001; 31 lives there too for a while.
- 3, 8, 10 and 18 meet in Connecticut in March 2001
- By Spring 2001, customs agents further connect 17 and 20 to financial deals with 3 and 26 (but no indication of when the transactions actually took place and the relationship was established).
- 16 was coordinating with 13 to get the hijackers into the US (via phone and money transfers), at the request of 32 (connecting 16 and 32 at this date)
- 26 and 28 attempt to travel to Florida together
- 7, 14 and 3 arrive in Miami on May 5 2001
- 26, 27 and 28 attend a gym together over Summer 2001
- 18 and 13 have established a relationship by Summer 2001 (phone calls and airport trips- according to documents filed in the Virginia District Court).

By August 2001

- 16 transfers $4900 to 7 in June 2001
- 21 and 29 meet in Pakistan in June 2001; 21 pays for 29's flight training in July 2001 & transfers him some funds

- 32 transfers money to 21 before July 2001
- 2 arrives with 27 in Miami on June 5 2001
- 6 and 23 arrive in Orlando on June 21 2001
- 6 and 16 have established a tight financial relationship by July 2001 with a significant transaction list (they have joint accounts and 16 has POA over 6)
- 1 arrives with 25 in New York on 23 June 2001
- Connection between 20 and 3 established (but not known when connection actually made)
- 1 lives with 8 and 25 in Paterson, New Jersey from March to Sept 2001
- 18, 23 and 13 are seen coming and going from the apartment in New Jersey
- 10, 9 and 3 are also seen living there in the summer
- 16 sends $10,000 to 21 in July 2001
- Witness statements put 12 and 13 together with 17 between Jan and Sept 2001
- 31 and 2 share an apartment. No date is given as to when this occurred (speculation of date here).
- 2 and 3 are living together in August 2001 in Florida
- 1 moved to Florida in August 2001 to join his flight team (no specified connections)
- 6 and 12 live together from late June onwards
- 32 is given access to 16's bank accounts