Introduction

Luke M. Gerdes

Regardless of their party affiliations, contemporary policy makers agree that dark networks threaten contemporary global security. For example, the fight against terror networks, most especially al-Qaida and the Taliban, became the centerpiece of President George W. Bush's foreign policy in the wake of the September 11 attacks. In his final State of the Union address, the Republican president described this contest as "the defining ideological struggle of the 21st century" (Bush, 2008). Despite the increasing polarization of American politics, his Democratic successor, Barak Obama, attributed similar importance to terror networks. Speaking in 2014, he acknowledged that "today's principal threat no longer comes from a centralized al Qaeda leadership," and instead originates from a complex network-of-networks comprised of "decentralized al Qaeda affiliates and extremists, many with agendas focused in countries where they operate" (Obama, 2014).

Other international leaders share these American presidents' concern over violent dark networks. In 2014, following a spate of well-publicized beheadings by the Islamic State in Iraq and the Levant (ISIL), British Prime Minister David Cameron characterized Islamist extremism as "a poisonous political ideology" that threatens the United Kingdom (Cameron, 2014). The rise of ISIL also motivated Secretary-General of the United Nations Ban Ki-moon to state that it is "undeniable – and the subject of broad international consensus – that these extremist groups pose an immediate threat to international peace and security"

This work was supported by the Office of the Secretary of Defense, Minerva Initiative. The views expressed herein are those of the author and do not purport to represent the official policy or position of the United States Military Academy, the Department of the Army, the Department of Defense, or the U.S. government.

(Ki-moon, 2014). It remains important to recognize that violent extremists have not entirely replaced traditional concerns about state-on-state conflicts (Lake, 2011), but terror networks feature prominently in the contemporary security discourse.

However, because dark networks include all types of hidden and clandestine groups, about whom data are uncertain and ambiguous, the term encompasses far more than terror cells (Tilly, 2005; Everton, 2012b; Gerdes, 2014). Other types of dark networks threaten the economy and human security. For example, policy makers have long recognized the pernicious impact of narcotics traffickers, money launderers, counterfeiters, and other criminal actors (Kerry, 1997; Clinton & Espinosa, 2010). Indeed, the UN Office on Drugs and Crime estimates that criminal networks collectively earn $870 billion per year, much of it at the expense of the legitimate economy (UNODC, 2011). The U.S. Senate's unanimous passage of a 2012 resolution condemning Ugandan warlord Joseph Kony and the Lord's Resistance Army ("Senate condemns," 2012), whom the United Nations estimate to be "responsible for more than 100,000 deaths" and tens of thousands of child abductions (Ki-moon, 2013, p. 14), demonstrates that there is also broad agreement that violent non-state actors (VNSAs) pose substantial humanitarian risk. Thus, many of the most pressing contemporary threats to security and stability stem from dark networks.

Given the success of network science in studying overt, "bright" networks, it is unsurprising that analysts studying dark networks have attempted to bend this methodology to their field of inquiry. But the fit between method and topic is imperfect. Standard network tools were designed to study companies, friendship groups, and other transparent organizations. The resulting data is typically well-bounded (i.e., it is clear who is a member of the organization under study), and effective data collection can ensure reasonably complete data. Indeed, analysts often gather data by conducting surveys of the networks' members (Valente, 2010, pp. 43–50). By contrast, it is often unclear who holds membership in a given dark network; data omissions and errors are the norm rather than the exception, and the inaccessibility of clandestine actors dictates that standard approaches to data collection are problematic. Imagine, for example, the impossibility of surveying the members of an Islamist terror group or narcotics trafficking organization about their relationships with peers.

Consequently, data on dark networks are incomplete, inaccurate, and often difficult to find. Moreover, dark networks are often organized to undertake fundamentally different tasks than transparent networks, so resources and information may follow different paths through each. To understand dark networks, analysts require unique tools and methods designed to fit these structures' inherently clandestine nature.

This edited volume brings together three overlapping communities of methodologists to discuss the state of the art in the assessment of dark networks. Civilian academics working at top-tier research universities discuss insights obtained from the classroom and the laboratory. Military researchers, including uniformed officers and the civilian analysts who work alongside them, describe lessons learned through a blend of academic and practical experience. Finally, consultants, who provide assessments, advice, and training to the U.S. intelligence community, offer a practitioner's viewpoint. Thus, this volume provides a diversity of perspectives and consequently includes competing solutions to some of the most persistent problems involved in the analysis of dark networks.

This volume's first section deals with theoretical issues pertaining to dark networks. Bienenstock and Salwen describe the perils associated with attempting to "reverse engineer" social meaning from structural data. Using insights from exchange theory, the authors question whether the same network metrics can be used to study any social structure regardless of the nature of the relationships. Because this chapter also debates the applicability of standard network metrics, which were designed to assess small, intimate data sets, to the study of large and very large data sets, Bienenstock and Salwen ultimately call for a change in priorities among dark network analysts. Rather than developing additional graph-theoretic methods and algorithms to manipulate relational data, the authors advocate a return to fundamental research to test the underlying assumptions of such methods, so that end users can understand the scope and limitations of using dyadic relationships to infer social meaning.

Gerdes offers the second chapter in the theory section. He argues that analysts studying dark networks need to give special consideration to issues of weighting and dimensionality, which he defines as the determination of the types of relationship that exist in a network. He offers a detailed ontology to classify relationships among actors in dark networks and uses correlation tests to demonstrate that failing to consider the theoretical and technical implications of weighting and dimensionality results in inaccurate assessments at both the node level and the network level. His findings suggest that many studies of dark networks, which typically model relationships as flat, unweighted networks, require reevaluation.

The final chapter in the theory section is from Bright. He offers a broad literature review that documents contemporary efforts to employ network analysis and simulation to study dark networks, especially those organized to traffic narcotics and conduct terrorism. He argues that researchers must consider four variables to ensure an appropriate fit between subject matter and method: (1) the nature of the available data sources, (2) the type of interventions (e.g., organizational decapitation) to simulate, (3) appropriate outcome variables to measure the success of interventions, and (4) how to model adaptation within the dark network under study.

Kenney and Coulthart begin this volume's second section, which is on data collection. They use the al-Muhajiroun activist network as an extended example to describe the challenges associated with extracting and analyzing a dark network from text using automated software tools, specifically the AutoMap software created by the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. The authors argue that the accuracy of results depends critically on concept-extracting "thesauri" that reconcile collections of synonyms to a single term. If analysts neglect to invest significant time and labor in the creation of thesauri, automated approaches to text analysis will fail. Moreover, the thesauri-building process cannot occur in a vacuum; analysts must utilize substantial domain-specific expertise to ensure conceptual validity. Thus, these researchers provide a practical example of the limitations that Bienenstock and Salwen describe regarding the inference of social meaning from dyadic relationships, while also demonstrating the challenges of including the relational dimensionality that Gerdes outlines.

La Valley, Usher, and Halman continue the conversation regarding data discovery methods. They describe geotimehashing, a method to infer relationships in dark networks based on entities' repeated geographic co-location. They describe how to apply the method to geo-tagged social media data, such as that provided by Twitter and Facebook. If combined with the text-based approaches that Kenney and Coulthart advocate, geotimehashing could unveil communities linked by both physical proximity and the topics of their social media posts.

Davulcu and Woodward's contribution discusses a tool that provides another integrative approach to data collection. LookingGlass is a software package that performs multi-scale text mining to develop a visual intelligence platform for tracking the diffusion of online social movements. The tool's algorithms discover hotly debated topics and organize output according to organization, viewpoint, and strength of opinion. Ultimately, LookingGlass tracks the geographical footprint, shifting positions, and flows of individuals, topics, and perspectives between groups.

Carley closes the section on data discovery with a contribution on open-source exploitation in the study of dark networks. She presents a semi-automated means of collecting and representing network data from news articles, blogs, and Web sites. Because such sources often provide information about dark networks that is incomplete, filled with errors, out of date, and spread across numerous hard-to-find documents, open-source data on dark networks present several major challenges. The dimensionality that Gerdes describes and the construct validity questions that Bienenstock and Salwen raise figure among these, as do Kenny and Coulthart's concerns regarding thesauri building. However, scalability, data integration, data layering, and data fusion are also problematic. Consequently, formal method is often difficult to apply in field settings. Carley outlines an integrative approach that seeks to overcome such challenges by using a tool-chain to parse open-source data into useable social scientific outputs.

This volume's third section focuses on measurement issues in dark networks. Arney, Bell, Coronges, and Merkl offer a chapter that explores new methods of assessing centrality. Because traditional network measures are often sensitive to changing, unknown, or inaccurate structural topology, analysts who rely on them often draw inaccurate conclusions. This problem appears to be particularly pronounced for dark networks because analysts often assume that these structures are inherently decentralized, even though emerging evidence, namely the preliminary results of an ongoing meta-analysis of data on dark networks conducted by the Mitchell Centre for Social Network Analysis at the University of Manchester, shows that clandestine networks display tremendous variance in degree centralization, density, and other network-level variables (Oliver et al., 2014). By taking into account observed node-level subdivisions, such as those that exist between the quasi-legitimate and overtly illegal portions of many violent organizations, Bell's subgroup techniques offer a generalization of centrality that may be more robust to the inevitable structural variations in dark networks. Merkl's sampling technique provides an additional method to estimate the confidence analysts should attribute to measurements in dark networks. When applied to test networks, this approach demonstrates that analysts are typically far too optimistic regarding estimates of centrality. Therefore, this chapter ultimately affirms the efficacy and validity of traditional centrality measures on the micro (subgroup) level, while simultaneously questioning their relevance at the macro level.

Shakarian and colleagues extend the conversation about measurement in dark networks to a new topical context: street gangs. This chapter introduces the GANG ("GANG Analyzes Networks and Geography") software, which is designed to aid intelligence analysis for law enforcement operations against violent street gangs. GANG uses new techniques in social network analysis to address several police analytical needs concerning street gangs. Specifically, the software can determine extent of membership for individuals who do not admit they are part of a street gang, can quickly identify sets of influential individuals (under the tipping model), can identify criminal ecosystems by decomposing gangs into subgroups, and can also consider geography in gang decompositions. After describing the design decisions underlying this measurement tool, the chapter provides results based on an analysis of anonymized real-world police data, and presents feedback from the Chicago Police Department, which currently uses the GANG software on an experimental basis.

Breiger and Pinson offer the final chapter on measurement. Their research views the network paradigm through a novel prism and measures similarity among 175 chemical, biological, radiological, and nuclear (CBRN) events, which are described by a comprehensive new open-source database. While standard investigations into CBRN proliferation emphasize the relations among variables and model predictor variables as having homogeneous effects on the outcome, Breiger and Pinson turn the usual regression models "inside out" to illustrate clusters of cases within which relations among key variables (e.g.,the effect of a perpetrator group's religious extremism on CBRN weapons pursuit) operate in opposing ways, thus aiding the identification of multiple threat scenarios.

The final section of this volume deals with tie formation in dark networks and emphasizes simulation-based approaches to the topic. Cunningham, Everton, and Murphy explore the factors that contribute to the evolution of dark networks. The authors propose a longitudinal, actor-based network model to capture the endogenous and exogenous factors associated with the evolution of a particular dark network: the terrorist network of Noordin Mohammed Top, who, until his death in September 2009, was Indonesia's most wanted terrorist and the mastermind behind a series of bombings. The chapter, which relies on data that Gerdes also critically reviews in this volume's theory section, suggests that actors' personal histories, including affiliations with specific schools and mosques, offer some of the strongest predictors of tie formation in dark networks, thereby providing evidence to question the long-accepted conclusion that dark networks are built on vague concepts of trust. The chapter also advances the application of stochastic-based models to the domain of dark networks, by analyzing undirected longitudinal data over several time periods, when the number of actors varies from one period to the next.

Scheinert continues the conversation on tie formation in dark networks by assessing the reciprocal impact of individuals on organizations and organizations on individuals. He argues that traditional approaches to the study of dark networks cannot accurately model embedded agent populations because such methods require analysts to identify either groups or individuals as the sole unit of analysis. Using data on Southeast Asian Islamist extremism in the period preceding the rise of Noordin Mohammed Top, Scheinert illustrates an approach to model groups, individuals, and their simultaneous impacts on each other, thereby enabling the examination of embedded units of analysis and their interaction. Scheinert's chapter offers proof-of-concept, rather than a fully calibrated model, but he ultimately posits rules regarding terrorists' operational and planning partnerships.

Pearson, Singer, and Airoldi describe an additional approach to simulation that attempts to ground a network-based approach to tie formation in dark networks against a more detailed model of the general population. The authors argue that the label "dark network" implies that covert group activity is driven by, and thus can be understood with, network metrics like degree distributions, paths, and centralities. Unfortunately, the central assumption of network analyses, which represented entities and interactions meaningfully reduced to a few categories, has limited empirical basis for covert groups. Analysts should, therefore, be skeptical of such approaches. To confidently apply a network-based analysis, it is necessary to ground that approach against more detailed models of underlying phenomena. These detailed models also have limited empirical support, but analysts can reason about them in a more principled way and consider alternative scenarios. Plausible guesses about underlying mechanisms make it possible for analysts to simulate activity and then project observed events onto a network for analysis. The authors demonstrate this approach by embedding synthetic covert groups into a real data set on the general population, before discussing the implications of their results for conventional social network analyses on covert groups.

Finally, Johnson, Johnson, and Arney close this book by detailing dynamic actor-oriented models (DAOMs), which can be used to replicate the general structure of dark networks. These models describe networks where the structure is determined by the decisions of a sample of agents, who represent the known actors in a dark network. DAOMs present a method based on nodal decision parameters (NDPs) that govern nodal decisions. In place of longitudinal data, the authors apply a bootstrap-like (Efron & Tibshirani, 1993) sampling algorithm to a test network depicting terrorism in Iraq. The process, which is time independent, requires only a partial list of actors, and allows for multiple connections at once, which lets it serve as a means to infer the overall topology of the "real" dark network from the subset of actors and relationships on which analysts have been able to gather information. Thus this piece emphasizes subdivisions between observed and unobserved actors, thereby echoing the same foundational distinctions as Bell's subgroup measures of centrality, while leveraging available data to explain broader patterns of tie formation.

1

Covert Network Analysis: An Exchange Network Theory Perspective

Elisa Jayne Bienenstock & Michael Salwen

The ultimate objective in using social networks to understand covert groups is to reliably and accurately leverage small amounts of revealed information about the social structures of these groups to gain insight into their operations. The rush to achieve this has put a priority on creating tools to apply graph-theoretic metrics more widely and to larger data sets with the underlying assumption that the metrics are valid regardless of the character of the connection, and that these metrics carry the same meaning regardless of context or scale. This chapter questions whether this headlong rush has "put the cart before the horse," and advocates for more basic and deductive research testing metrics under a broad swath of conditions to develop standards to measure the reliability and validity of results. To make the point explicit, and to recommend a path for new research, this chapter revisits the classic research on social exchange networks, unfamiliar to many who study covert networks, that casts doubt on some fundamental assumptions that underlie the current approach.

Two different types of risks are associated with the use of data on social networks in operational settings before the requisite science is completed. The first is that shortcomings attributed to the approach may discredit the method, discouraging its use when it is appropriate and can provide real value, proverbially throwing out the baby with the bath water. The second is that the naïve application of network metrics to data can yield results that are incorrect, which in an operational environment can have serious implications.

Consider, for instance, the case of Abu Zubaydah, also known as Zein al-Abideen Mohamed Hussein. In 2002, President George W. Bush described him as "al-Qaeda's chief of operations," thought to be "the number three name" in the organization (Finn & Tate, 2009). As a consequence Abu Zubaydah was captured, detained at Guantanamo Bay, and subjected to "enhanced" interrogation eighty-three times. Four years later, in 2006, journalist Ron Suskind revealed that in fact Abu Zubaydah was only "a minor logistics man, a travel agent" (2006). Having been wounded in an earlier jihad, he was considered mentally limited, and al-Qaida used him to book the travel of wives and children.

It is not known whether social network data were used to identify Abu Zubaydah as critical to al-Qaida's operations, but it is easy to imagine how a naïve application of ideas informed by social network theory could result in an analyst drawing a false conclusion from phone or e-mail traffic data. The most popular, easy to compute, and well understood social network metric, degree centrality, has been shown time and again to be correlated with power in social groups. This association of high power to structural degree centrality might easily lead an analyst to infer power or importance from degree centrality; however, this is a fallacy. In fact, organizational studies abound with examples of individuals with low power and high degree centrality, for instance, mailroom personnel, pizza delivery people, coffee shop clerks, and now travel agents. Moreover, even if there were some contexts or types of relationships for which degree centrality served as a reliable indicator of power, there is no justification for blindly generalizing this measure's use to other relationship types. If data on social structure is going to reveal the link between structurally discernable attributes and social meaning, a meticulous articulation of scope conditions is needed.

The next section provides background into the history of network research that indicates how network science got ahead of itself. Following that is a discussion of findings and insights from exchange network theory relevant to the analysis of covert networks. The argument advanced there is that the proper analysis of exchange relations requires metrics different from those used in the analysis of other types of relations, and recommends the creation and adoption of a priori standards to determine which metrics are useful under different sets of conditions.

I. Social Network Analysis (SNA) and Network Science

Until the turn of this century, SNA was an obscure collection of data analytic methods a small group of sociologists and anthropologists used to formally represent relations among members of the social groups they studied. These researchers, in collaboration with scientists from a host of disciplines, discovered that mathematical operations performed on these formal representations produced useful metrics that correlated well with constructs important to understanding social dynamics.

The primary focus of research on social networks was the development and empirical testing of structural metrics of social groups as one component of rich, multifaceted studies that included many types of data on the quality and contexts of relationships. In trying to organize and formalize these rich data, researchers recognized that by representing relations between pairs of individuals – dyadic ties – as edges in a graph or as values in an adjacency matrix, they could rigorously conceptualize and quantify concepts that had traditionally been described in purely qualitative terms (Freeman, 2004).

Researchers soon discovered that specific structural patterns were associated with social concepts that were useful for providing deeper insights into the functioning of the groups. Associations between certain structural patterns in the graphs and important social concepts were found consistently across data sets. For instance, powerful and influential people were also often "central" in the group, by having many direct connections (degree centrality), short paths to many others in the group (closeness centrality), or the exclusive ability to act as a liaison connecting people otherwise not connected (betweenness centrality) (Freeman, 1979). Moreover, robust evidence suggested that people surrounded by others who shared common connections (i.e., those with dense local social networks) were more likely to follow the norms of the group (Bott, 1955; Coleman, 1990).

Initial findings inspired more empirical work, both in the form of additional field studies conducted across a wide variety of topical contexts and in laboratory experiments designed to test and falsify conclusions drawn from observations relating network structural characteristics to social roles or social consequences. What emerged was a community of scholars who shared a theoretical perspective and a methodological approach to characterize data on social relationships as graphs or adjacency matrices (Leinhardt, 1977; Wasserman & Faust, 1994). The goal was to advance the techniques so that they could be more broadly used to augment traditional methods to characterize social groups; however it was also recognized that structure alone was not sufficient to infer meaning and that extrapolating too much from the representation of dyadic ties could lead to a faulty understanding of the group. Samuel Leinhardt summarized this sentiment:

> "The ease with which social networks can be represented mathematically has the potential of improving the precision and specificity of theories of social structure. Obviously, this does not mean theories of social structure based on the network paradigm are necessarily true. What it means is that when networks are employed, the precision of mathematical ideas and the power of mathematical operations can often be used to help us think more clearly and more exactly about structure in social relations." (1977, p. xiv)