

Inside Social Network Analysis

Kate Ehrlich¹ and Inga Carboni²

Introduction

A management consulting firm hopes to win a lucrative contract with a large international financial institution. After weeks of intense preparation, the team sends off a proposal. Shortly thereafter, they learn that contract was given to a competitor with whom the client had worked previously. Almost six months later, one of the team members finds out that another group at the management consulting firm had worked on a project with the prospective client and had gained an in-depth knowledge of its business operations. Why, asked the frustrated team member, wasn't this critical knowledge shared with the team?

This is just one example of the opportunities that large companies can miss if they fail to understand that success depends less on reporting structure and more on an informal web of contacts. In the past, companies that encountered a loss of business like the one described above, might conduct a survey and interview employees to discover what went wrong. A new approach called social network analysis or SNA, has been gaining currency among business consultants as an effective method for revealing the hidden connections that drive how work gets done.

In the networked organization, individual success and the success of a team depend less on reporting structure and more on who you know. This paper provides a brief introduction to the major concepts and measures in SNA and their application to business problems.

What is a Social Network Analysis?

A social network analysis examines the structure of social relationships in a group to uncover the informal connections between people. In a consulting setting, these relationships are often ones of communication, awareness, trust, and decision-making. As an approach to looking at these relationships, SNA has been around a long time.

¹ IBM T.J. Watson Research (katee@us.ibm.com)

² Boston College (inga.carboni@bc.edu)

Most network analysts cite Joseph Moreno's introduction of the tools and methods of sociometry, in 1934 as the year in which the formal analysis of social networks began. SNA became much more popular with researchers in the early 1970s when advances in computer technology made it possible to study large groups. Within the last ten years, SNA has risen to prominence in a number of fields, including organizational behavior, anthropology, sociology, and medicine.

Most recently, SNA has become an important tool for organizational consultants seeking to understand the connection between patterns of interactions and business outcomes such as job performance, job satisfaction, adoption of new ideas or technologies, likelihood of information getting shared, and creation of new ideas.

Together, these elements can be analyzed to reveal the social network and determine if it is meeting the business needs of the group. The outcome of an SNA lets us see where collaboration is breaking down, where talent and expertise could be better leveraged, where decisions are getting bogged down, and where opportunities for diffusion and innovation are being lost. The data give leaders the picture they need to create a set of remedial actions for individuals and leaders to improve productivity, efficiency and innovation. These actions include modifying roles and responsibilities to foster more effective patterns of communication, methods for improving trust, better use of technology to reach others, re-alignment of rewards and incentive programs.

How is a Social Network Analysis different?

SNA differs from conventional approaches to business problems in one very important way: SNA assumes that *people are all interdependent*. This assumption is radically different from traditional research approaches which assume that what people do, think, and feel is independent of who they know.

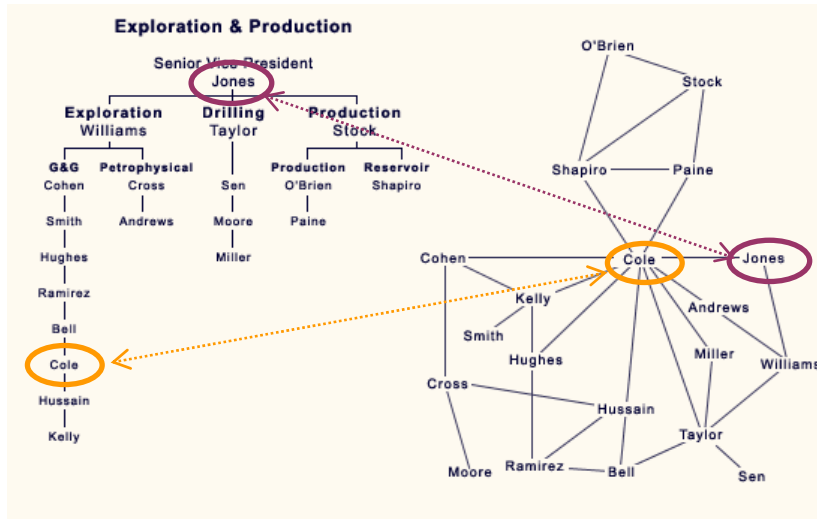
The focus on interdependence means that SNA can ask -- and answer -- questions such as:

- Is Sales effectively communicating with Marketing to share and coordinate information about the customer?
- When two companies or organizations merge, how can management use the informal network to spread important messages?
- Are decisions in a distributed software development team being made and carried out efficiently or are one or more people acting as bottlenecks?
- In R&D group, are there enough people bringing in ideas from outside and are those ideas being acted upon?

Many traditional statistical techniques are based on the assumption of independence. For this reason, traditional statistics, such as comparing the means of two groups, cannot be conducted on interdependent data. To deal with this problem, SNA has developed a set of SNA-specific statistics such as **centrality** and **density** that provide measures of interdependence

Why do a Social Network Analysis?

If we want to understand how a group functions we might go to an organization chart to find the senior people who are empowered to make decisions or to see how the work is divided up functionally. But in the evolving networked organization, this chart is no longer an adequate guide to how the group really works. Consider the following diagrams which represent the production division of a large petroleum corporation (Cross et al., 2001).



Upper management wanted to know how this group was preparing to share important drilling knowledge so they undertook an SNA. In the organizational chart on the left, we see that Jones is the senior person in the group while Cole is in a more junior role. An SNA revealed that, in contrast to the formal chart, mid-level managers, and Cole in particular, were playing a pivotal role in the group's communication network. Not only was Cole linked to many people, making him very central to the group, but he was also the only link between the cluster of people at the top who represented production and the rest of the group who were involved in other distinct but critical activities.

Through additional interviews with key people in the network, including Jones, it was apparent that Jones had become removed from many of the day-to-day workings of the group. His lack of responsiveness and participation in the group often resulted in delays which impacted the rest of the group. Without an SNA, Jones' lack of active participation in the group and Cole's importance may have gone unnoticed. After conducting the SNA, upper management was able to consider a range of options, including formalizing Cole's role.

Business applications of SNA

SNA applies to a wide range of business problems, including:

Knowledge Management and Collaboration. SNAs can help locate expertise, seed new communities of practice, develop cross-functional knowledge-sharing, and improve strategic decision-making across leadership teams.

Team-building. SNAs can contribute to the creation of innovative teams and facilitate post-merger integration. SNAs can reveal, for example, which individuals are most likely to be exposed to new ideas.

Human Resources. SNAs can identify and monitor the effects of workforce diversity, on-boarding and retention, and leadership development. For instance, an SNA can reveal whether or not mentors are creating relationships between mentees and other employees.

Sales and Marketing. SNAs can help track the adoption of new products, technologies, and ideas. They can also suggest communication strategies.

Strategy. SNAs can support industry ecosystem analysis as well as partnerships and alliances. They can pinpoint which firms are linked to critical industry players and which are not.

Conducting a Social Network Analysis

A social network analysis in a business setting has three important elements.

- **A group.** The first step in conducting a social networks analysis is to determine the group under study. In the management consulting example described at the beginning of this article, the group may be all of the consultants in the particular practice. A group can be composed of people who are dedicated to a particular task, such as a software development team or it can be a collection of people such as a community, upper management across different business units, or members of a temporary task force. Groups typically range in size from 25 to 200. The individual units in a group are commonly referred to as actors or nodes.
- **Interactions.** An SNA looks at the relationships between actors. The nature of the relationships we are interested in will vary according to our reason for conducting the study. In the case of the management consulting company, we might be interested in analyzing communication patterns by looking at which consultants interact with each other for the purpose of obtaining client-related information. Interactions are also referred to as the links or ties between people. The pattern of interactions in a group is called a social network.

- Attributes.** Attribute data can help determine whether there are systematic factors that influence interactions between people. For example, we often find that people in one business unit don't routinely share information with people in another unit. The factors that might influence these interactions range from incentive programs that motivate people to spend their time with people in their own business unit to "cultural" differences, such as language or work ethos, that make it difficult for people to communicate easily with one another. In a study on diversity, for example, attributes that influence interactions might include ethnicity, gender, and attitudes about affirmative action. In the management consulting case, relevant attributes might include where someone works (e.g. country, geographic region), which business unit they are in (e.g. sales, marketing, development), their level of seniority, and how long they have been with the company. They may also include personality measures, such as Myers-Briggs personality types or scores on aptitude tests. Only attributes that are believed to influence interactions are included in an SNA.

What can ties reveal?

The result of collecting SNA data is a matrix showing, in numerical form, the existence, type, and/or quality of interactions between pairs of people. Information about these kinds of ties is commonly collected through interviews or surveys, often administered online. A typical survey might list all of the people in a work team and then ask each individual who they go to when they need client-related information. The result of all the responses might look something like the tables below which shows, for each pair of people, who goes to whom for information.

	BILL	JIM	CAROL	PAM	PAT
BILL	0	0	0	1	1
JIM	0	0	0	1	0
CAROL	0	1	0	1	0
PAM	0	0	0	0	1
PAT	0	0	1	0	0

In this example, you can see that Bill, Jim and Carol all go to Pam for client-related information. Pam, however, goes only to Pat for information and Pat goes only to Carol. Thus, Carol, who only one person turns to directly for information, may actually be the source of much of the client-related information shared in the group (via Pat and Pam).


Collecting information about ties is not limited to surveys. This information can also be inferred from a number of existing data sources, such as email exchanges, direct observations of group interaction, billable hours (i.e., who works on projects with whom?), professional citations (i.e., who publishes with whom?), corporate board interlocks (i.e., who serves on a board with whom?), consulting contracts, package transmission (i.e., who is delivering packages to whom?), or charitable donations (i.e., who is giving money to whom?). Any of these methods can generate relationships between two or more people which together can represent the social network for the whole group.

Beyond simply looking at who is connected to whom and who is central in the group, ties can indicate the strength and direction of a relationship, as well as if the relationship is direct or indirect. To draw these inferences it helps to understand some of the concepts and implications of ties.

Are ties strong or weak? Strong ties, indicated by a higher number, are characterized by frequent interaction, feelings of closeness, and multiple types of relationships. For example, a strong tie may provide you with emotional support, job-related information, and a person to go see your favorite sci-fi movies with. On the other hand, it also requires a good deal of energy to maintain. Weak ties may not provide as much social support but, since they are easier to maintain, you can have many more of them. Weak ties can be critical for innovation. For example, a researcher is much more likely to learn about a line of relevant research in an otherwise unrelated field from a casual acquaintance than from a good friend. This is because good friends tend to have access to the same information whereas casual acquaintances tend to offer new information.

In business settings, it is important to have a good balance between strong and weak ties. Too many strong ties – always going to the same set of people for information or advice, may limit access to important new information. In our first example of the management consulting company, the client team had been interacting mainly with existing trusted connections who didn't know about the other team's relationship with the same client because the other team came from a different part of their company.

Are ties reciprocal? Some of our relationships are naturally two-way. If I meet with you, then you are probably also meeting with me. However, many relationships can be one-way and the directionality can provide additional information. For example, I may respect you but that doesn't mean that you necessarily respect me. In a network diagram, the arrows indicate directionality.



The tie between A and B is reciprocal. The tie between A and C is not reciprocal; A goes to C but C does not go to A


When a tie goes in both directions we think of it as being reciprocal. It can be important to know when a relationship is reciprocated. For instance, a group will generally function better when a key decision-maker is not only sought after for information but he or she also seeks information from the group. In general, reciprocated ties tend to be stronger than non-reciprocated ties.

Are ties direct or indirect? A direct connection is the connection between two people. As we saw in the petroleum company example, Cole had many direct connections in the network of people in exploration and production. He was *central* in this network. Central people have more influence in their network, tend to receive better performance reviews, and tend to be more satisfied with their jobs than people who are less central.

A person can also be indirectly connected to another person through knowing a “friend of a friend.” Tracking the number of connections it takes for one person to reach another is one way to monitor the flow of information and opportunities in network. A person who is indirectly linked to many top level people, is more likely to get promoted than someone who is not.

The related, proverbial “small world problem”, refers to the likelihood that two people who are not directly connected to each other will be able to communicate via indirect connections. Due to random connections between groups of people, people who appear unconnected may, in fact, be just a few relationships apart.

Ties can represent a range of relationship types. For example, a tie can indicate if one person *likes, trusts, respects, reports to, communicates with, or gets information from* another. Similarly, ties can represent *sells to, buys from, delivers to, contracts with or collaborates with* relationships. The presence of a network tie indicates that a relationship exists. The absence of a network tie indicates that the relationship does not exist.



This diagram, drawn from a study of participants in a workshop illustrates *who likes whom* (Borgatti et al., 1992). It also illustrates a common pattern; namely, that people who share a particular characteristic or attribute, in this case gender, frequently interact more with each other than they do with people with whom they do not share a particular characteristic or attribute. In the diagram the males are all interacting with each other as are the females with only two points of contact between them.

Actions following a network analysis

An SNA used in a consulting setting is primarily diagnostic of the group being studied. To improve the flow of information, knowledge, or decisions in these groups the SNA should be accompanied by actions. These actions may be the responsibility of individuals, the group, management or the larger enterprise.

An increasingly popular use of SNA is to feed back results to the entire network in a structured setting in which discussion can be productively facilitated. Viewing the SNA diagrams seems to be particularly powerful. For example, in the case of the management consulting company, the results surfaced hidden barriers between senior and junior consultants. Revealing these previously hidden barriers constructively was a powerful tool for creating new opportunities for cross-group relationship development. An SNA can also suggest ways of restructuring organizational charts or implementing new processes.

Similarly, an SNA conducted on a set of organizations can let leaders know where their firm stands in the social network. Leaders may then make strategic decisions about where and with whom to develop closer business relations.

Application of SNA: The case of the management consulting company

The management consulting company mentioned at the beginning of this article needed help improving its consultants' awareness of related projects. Before embarking on potentially costly organizational changes and other initiatives, the executive leadership elected to conduct an SNA to find how information was (or was not) getting passed on. The SNA was conducted with a group of consultants from different geographies, practices, levels of seniority, and tenure in the company.

Questions included: "How aware are you of the projects done by this person in the past 12 months?" and "How often does this person provide you with information you need to develop client proposals?" By looking at who was central in the awareness and information-sharing networks, the company could see which people the team tended to go to for information. By looking at how often people from each group interacted with each other, the SNA revealed that consultants who had been with the company the longest tended to get information from each other; newer employees only turned to each for more information if they were located in close physical proximity. Follow-up interviews revealed that time pressures left members of the team with few opportunities to develop relationships with newer or more distant employees.

As a result of the SNA, upper management took several actions. They instigated a mentor system to help new employees, sponsored events several times a year that brought people together from different parts of the company, and developed an award program for any new engagements that were won as a result of existing relationships from outside the immediate team.

SNA: A final word...

SNA is a new approach to solving traditional problems in business and management. It assumes that people are interconnected, that connections have real consequences for performance and satisfaction, and that connections can be structured to optimize individual, group, and organizational outcomes. SNA responds to the growing awareness that something intangible is in danger of being lost as the marketplace (and the workplace) becomes increasingly dispersed.


The interpersonal connections that we once took for granted are no longer an assured part of everyday life. SNA places a premium on these relationships and their role in the way we do business.

Appendix

SNA concepts

Social network analysis is an established subdiscipline of many academic fields, including management, social psychology, and sociology. As a result, research has already revealed relationships that may be helpful to practitioners. Below are three examples of key SNA concepts. To learn more about SNA, refer to the references listed at the end of this Appendix.


Centrality. Centrality is the extent to which a person is in the center of a network. Central people have more influence in their network, tend to receive better performance reviews, and tend to be more satisfied with their jobs than people who are less central. Measures of centrality include [degree](#), [betweenness](#) and [closeness centrality](#).



Person A is the most central in the network

Social capital. Social capital refers to the total sum of potential or actual resources that a person accrues as a result of interpersonal relationships. This translates into being able to get favors granted, getting rapid response to requests, or having access to influential people. Social capital may arise from having a dense network of strong support or being in a privileged role with respect to other people, referred to as structural hole (see below).

Structural holes. A structural hole occurs whenever a person (a) has a relationship with someone who is connected to a separate cluster of people and (b) has no other direct or indirect connection with the people in that cluster. In effect, the person is connected to a distinct group via one person.



A has the most structural holes in the network

A person who has a large number of structural holes in his or her network is likely to be exposed to more diverse information and opportunities than a person who has relatively few structural holes in his or her network. For this reason, a person with a large number of structural holes is more likely to be promoted, develop innovative products, and to gain power and influence by acting as a broker between otherwise unconnected groups.


Brokerage. A person who connects two otherwise unconnected people is in a position to manage or “broker” information flow. How he or she chooses to do that can be revealed through an SNA. For example, Bill may broker the relationship between sales reps and product developers by receiving information from sales reps and providing it to product developers but not vice versa. Conversely, Bill may receive and provide information to both sales reps and product developers. Moreover, Bill may be a sales rep himself or he may be a product developer or he could be in an unrelated group, such as distribution. The division to which he belongs has implications for the role he plays (e.g., insider, outsider, etc.). Brokerage differs from structural holes in that (a) it does not assume that groups are unconnected and (b) it more clearly describes the flow of information.

Common SNA statistics and measures

Because SNA rejects the assumption of independent observations underlying traditional attribute-based research, several SNA-specific statistics have been developed. Some of the most common are listed below. A comprehensive list of SNA-specific statistics can be found in the UCInet reference guide (see references, below).

Centrality. The most common measure of centrality is degree centrality and is a simple tally of the number of people attached to each person. A person with three reciprocal relationships, for example, has a “degree” of 3. Other measures of centrality include betweenness and closeness.

Betweenness centrality indicates the extent to which a node lies on the shortest path between every other pair of nodes. For example, if Adam is connected to Bill and Bill is connected to Cindy but Adam is not connected to Cindy, then Adam must go “through” Bill to “reach” Cindy because Bill is between Adam and Cindy. The person with the highest betweenness centrality is the person that others in the network most commonly must go through in order to reach each other.



A is high in betweenness centrality.

B is high in degree centrality.

Closeness centrality is the number of links that a person must go through in order to reach everyone else in the network. If Adam is connected to Bill and Bill is connected to Cindy but Adam is not connected to Cindy, then Adam must go through two links (i.e., Adam-Bill and Bill-Cindy) in order to reach Cindy. The person with the highest closeness centrality score is the person who goes through the fewest number of ties to reach everyone else in the network.

Density. Density captures how closely a group or subgroup is knit. It is a proportion that indicates the number of actual ties present in the group relative to the number of possible ties in the group (i.e., if everyone had a relationship with everyone else in the group). Density can be calculated within a group or between two groups. When calculating the density of an individual's network, SNA looks at how closely connected a person's friends are to each other.

Cohesion. There are several measures of cohesion, including density. However, one common measure is the average number of ties it takes for a person in the group to "reach" another person in the group. If Adam is connected to Bill who is connected to Cindy, then Adam is at a distance of 2 from Cindy. The average distance for the group gives an indication of the group's cohesion.

Subgroup Identification. SNAs can identify the number of closely knit subgroups or "cliques" in a network. Within a clique, every unit is connected to every other unit. These subgroups can then be analyzed to see if they share overlapping members. A network that contains highly segregated subgroups is not as well integrated as a network in which individuals belong to several overlapping subgroups.

References

- Borgatti, S.P., Bernard, H.R., and Pelto, P. 1992. NSF Summer Institute on Ethnographic Research Methods. Available from Analytic Technologies www.analytictech.com
- Brass, D. (2000). *Achieving Success through Social Capital: Tapping Hidden Resources in your Personal and Business Networks*. Jossey-Bass.
- Cross, R. and Parker, A (2004). *The Hidden Power of Social Networks*. Harvard University Press
- Cross, R., Parker, A., Prusak, L. & Borgatti, S.P. 2001. Knowing What We Know: Supporting Knowledge Creation and Sharing in Social Networks. *Organizational Dynamics* 30(2): 100-120. [[pdf](#)]
- Cross, R. and Prusak, L. (2002). The People Who Make Organizations Go—or Stop. *Harvard Business Review* 80(6), pp 104-112
- Cross, R., Liedtka, J. and Weiss, L. (2005). *A Practical Guide To Social Networks*. *Harvard Business Review*.
- Hanneman, R. and Riddle, M. (2005). *Introduction to social network methods*. Riverside, CA: University of California, Riverside (published in digital form at <http://faculty.ucr.edu/~hanneman/>)

Krackhardt, D and Hanson, J. (1993). Informal networks: The company behind the chart. *Harvard Business Review*, 71(4), pp104-111

Borgatti, S. and Foster, P. (2003). The network paradigm in organizational research: A review and typology. *Journal of Management* 29(6), 991-1013.

Wasserman, S and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press

For access to UCInet software (the most commonly used software for SNA) and associated articles: www.analytictech.com

For information about InFlow analysis tools, articles and diagrams:
<http://www.orgnet.com/inflow3.html>