Network-Enabled Collaborative Problem Solving

Steve Poltrock[1] and Paul R. Smart[2]

[1]*Boeing Phantom Works, Seattle, Washington, 98124, USA.*
[2]*School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.*

Introduction

People working together can accomplish tasks and solve problems that far exceed the capabilities of a single individual. In this paper we consider how people collaborate when solving problems and the role of networks in supporting this collaboration. We consider what aspects of problem solving are best performed alone, what is best done in collaboration with others, and the communication needed to support collaboration while solving complex problems. We also consider the cognitive processes and mental representations involved in problem solving by a single person. When people collaborate, these cognitive processes and representations extend across all participants, and we consider how these processes are replicated, shared, and/or allocated. Finally, we consider the role of both human networks and the networked technologies that support collaborative problem solving. Network properties such as bandwidth and latency may influence how problems are solved as well as both the quality and timing of solutions. Our objectives are to provide a theoretical framework for network-enabled collaborative problem solving that would underpin a research program for improved problem solving performance.

People solve problems (or contribute to solving problems) as individuals, in families, as part of organizations, and as part of communities and societies. As individuals, we frequently solve relatively simple problems such as deciding which clothes to wear or what food to eat, and even these simple problems may be solved with the help of other people. Families make decisions about when and where to take vacations, where to live, and how to manage family finances. Organizations formulate missions and determine how to accomplish these missions. Civilian organizations have missions that involve the design, development, production, and dissemination of products or services, whereas military organizations have missions that may involve finding and suppressing terrorists while providing a secure environment for civilians. The missions of both military and civilian organizations constitute challenging complex problems that require the coordinated activities of many people.

Consider simple problems and the role of collaboration in solving them. Does collaboration help solve simple problems, and if so, how? Liker and Bókony (2009) gave house sparrows a relatively simple problem: in order to obtain bird seed the sparrows had to lift a rubber stopper covering a well in a Plexiglas bird feeder. Sparrows were placed in an aviary in groups of two or six birds. The sparrows in larger groups were 11 times faster opening a well and they opened four times as many wells. All the sparrows in larger groups ate bird seed, but some of the sparrows in smaller groups got none.

We may wonder about the advantage of being in a larger group in this study. Did these birds solve the problem collaboratively? One possibility is that social birds are simply more active when in larger groups. Zajonc, Heingartner and Herman (1969) concluded that two cockroaches solved a runway task more quickly than one cockroach because of a 'social facilitation' effect (Zajonc, 1965). But observation of the sparrows in Liker and Bókony's (2009) study did not support this conclusion. Instead, the results suggest that larger groups were more likely to contain birds smart enough to solve the problem, and other birds quickly learned from observing the smart birds' success. Thus, one advantage of collaboration is that a larger more diverse group is more likely to contain individuals who know part or all of the solution or who are smart enough to find a solution; everyone collaborating benefits from the contributions of these individuals.

Collaborative problem solving among people can, of course, be much more complicated. Through language and images people can describe and discuss the problem, potential solution methods, and the solutions themselves. People can structure both problems and themselves to improve performance. People decompose problems into sub-problems and may allocate these sub-problems to different individuals or groups of people. Indeed, the organizational structure or collaboration network structure adopted by people strongly influences the solutions generated to problems. Furthermore, problems differ in the extent to which they can be decomposed, implying that different kinds of collaboration network structures, and different communication and collaboration patterns, are needed for different kinds of problems. Understanding how problem solving and networks are interrelated requires understanding how people solve problems and how people collaborate when solving problems.

In this paper we will briefly summarize theories of human problem solving beginning with a general theory developed by Newell and Simon (1972) that has been extraordinarily influential. More recent research is consistent with this general theory but places greater emphasis on how the physical and social environment may contribute to problem-solving success (e.g. Kirsch, 2009; Hutchins, 1995a). The general theory describes how an individual solves problems, but provides a foundation for thinking about how people collaborate when solving problems. We will explore the relationships between elements of the theory and human collaboration. We will then turn to properties of human collaborative networks and consider how those properties influence problem solving behavior and performance. We will see that properties of problems influence the kind of network that will be most effective. We will also consider how people construct collaborative problem solving networks. Finally, we will conclude with observations about ways in which network topology, bandwidth, and resources may influence collaborative problem solving performance.

Theory of Human Problem Solving

A General Theory

In 1972 Newell and Simon published *Human Problem Solving,* a landmark in understanding how individuals solve problems. This book reported research in which people talked about what they were thinking while solving intellectual problems such as cryptoarithmetic or playing chess. Viewing people as information processing systems, Newell and Simon proposed a general theory of problem solving consistent with these introspective descriptions.

This general theory of problem solving is summarized in Figure 1. The task environment, represented by the outer box in Figure 1, is the external environment in which the problem or task is instantiated. The inner box represents the information processing performed by the problem solver. A key element of this theory is the problem space, a mental construct that people create, invoke, and manipulate. The problem space includes the internal representation of the problem, methods that can be employed to solve it, and knowledge about the methods and problem.

The human information processing system creates and maintains this problem space, beginning with the initial state of knowledge about the problem. When creating the initial problem space human perception may play some role, but memory and experience play larger roles. Representations of a problem and the transformations that can be applied to these representations depend greatly on past experience with similar problems. Experience enables a solver to change the problem representation by constructing higher level representational units that diminish the memory capacity required to represent and reason about the problem.

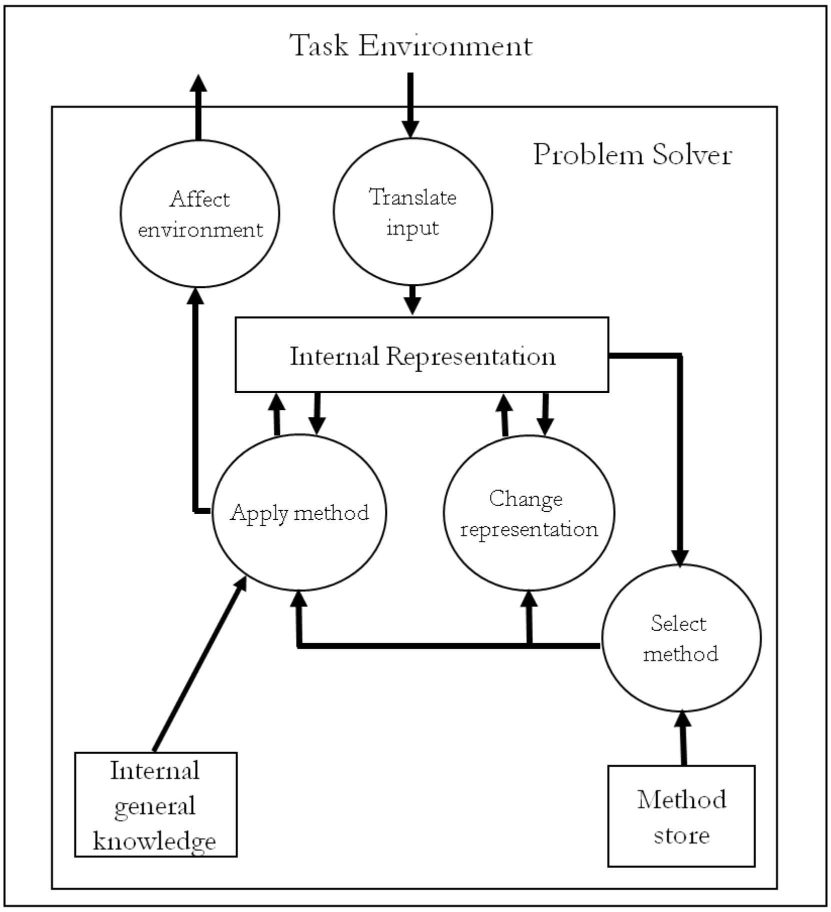

Figure 1: A schematic of Newell and Simon's (1972) general theory of problem solving.

The solution steps may involve reasoning processes, associations with related problems stored in memory, or heuristics. The solver may try various transformations and compare their results with the desired solution state (generate and test method) or employ other heuristics to choose transformations. Solving a problem can be viewed as a search task through a problem space of potential state transformations, and may include the generation of sub-problems that are solved using the same mental apparatus. Problems are solved, according to this theory, by an information processing system that manipulates a problem space.

Hard and Ill-Structured Problems

Why are some problems harder to solve than others and why are some people able to solve problems more often and more quickly? According to the general theory, the ability to solve a problem depends primarily on how the task or problem is initially presented, prior experience with this kind of problem, and cognitive abilities. The presentation of the problem can influence the internal representation in ways that determine problem difficulty. Consider, for example, the famous nine-dot problem shown in Figure 2. The problem is to connect all nine dots with only four lines (the solution is shown in the figure). When people first encounter this problem, they generally attempt to solve it with lines that do not extend outside the rectangle formed by the dots, which makes the problem insolvable. Solving the problem requires changing the internal problem representation to allow transformations with longer line segments.

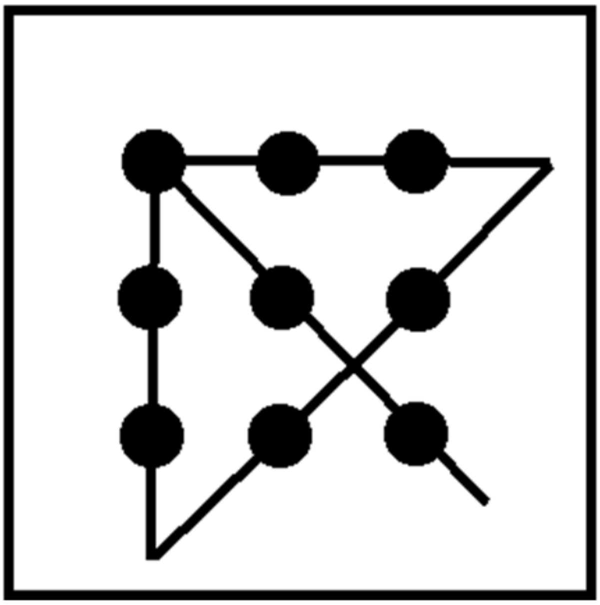

Figure 2: The nine-dot problem requires connecting all nine dots with four straight line segments (a solution is shown).

If problem solving is achieved via a process of search through a problem space, then memory load must contribute to problem difficulty. Remembering which paths have been explored is difficult when the search space is large. Novice chess players, for example, often have difficulty remembering which possible moves they have considered and their evaluations of those moves. Some problems are difficult because they require special knowledge or expertise. Astrophysicists, for example, would have little chance of solving problems in their domain without expertise in mathematics.

The problems studied by Newell and Simon were well-formed; the starting state, goal state, and possible transformations were specified or already known by the participants. Well-formed problems are generally easier to study in a laboratory setting, but the problems that plague people in the real world are often ill-formed. Solving ill-formed problems may require determining the starting state, the goal state, and/or the permissible transformations. Ill-formed problems are not necessarily hard; examples (ranging from simple to extremely difficult) include what clothes to wear, how to capture an enemy position, and how to end world hunger. Newell and Simon's general theory describes how problems are solved once the problem attributes have been determined.

Situated Cognition

Situated cognition is a theoretical perspective that emphasizes the essential role of the environment or situation in which cognition occurs. It demands a different empirical approach to that traditionally encountered in cognitive psychology. Instead of studying people in laboratories solving problems posed by a researcher, this perspective requires that problem solving be investigated as people cope with the problems they encounter in the real world.

Kirsh (2009) summarizes some of the major criticisms of Newell and Simon's theory from the perspective of situated cognition. Advocates of situated cognition note that a problem's framing strongly influences how it is solved. A problem's framing is defined by elements of context that appear to be outside the scope of the general model. A chess master, for example, is likely to conceptualize the game and play very differently when playing in a tournament than when playing against her grandchild. A problem's framing may strongly influence motivation to achieve an optimal solution and even the definition of an optimal solution.

Situated cognition may also require registration between an internal representation of a problem and the physical environment. Consider, for example, the problem of finding a particular store in an unfamiliar shopping center. A map that identifies the location of the store offers a quick solution, but the shopper must also identify his current location and establish the relationship between objects on the map and objects within view. Finding a path on a map is easy, but registering the map with the physical environment may be harder. Similarly, the mapping between the internal representation or problem space and the task environment may be difficult.

Newell and Simon's general theory provides for the solver to take actions that affect the environment, as shown in Figure 1, but situated cognition emphasizes interactivity with the environment, which may include other people. People generally do not solve problems simply by manipulating an internal representation of the problem space; they interact with the physical elements of the problem to enhance understanding of potential transformations and to achieve partial solutions. A person preparing a meal, for example, may collect together ingredients and organize them in a way that will provide reminders of how and when each is to be used. The general theory permits this kind of interaction but does not emphasize it to the same degree as situated cognition perspectives. Indeed, Hutchins (1995a) suggests that the classical information processing view of human problem solving may best be seen as a model of human agents interacting with real-world, bio-external resources. He writes:

> *"The model of human intelligence as abstract symbol manipulation and the substitution of a mechanized formal symbol-manipulation system for the brain result in the widespread notion in contemporary cognitive science that symbols are inside the head. The alternative history I offer is not really an account of how symbols got inside the head; it is a historical account of how cognitive science put symbols inside the head. And while I believe that people do process symbols (even ones that have internal representations), I believe that it was a mistake to put symbols inside in this particular way. The mistake was to take a virtual machine enacted in the interactions of real persons with a material world and make that the architecture of cognition." (pg. 365)*

Situated cognition theorists observe that our physical environment is often designed to help us solve problems. The entire field of human factors is dedicated to identifying designs that facilitate interactions, and our environments are filled with signs and objects intended to help us achieve our objectives. Reflecting on the extent to which humanity has designed its physical environment to support its tasks, Clark (1997, p. 191) writes, "The coherence and the problem-solving power of much human activity, it seems, may be rooted in the simple yet often-ignored fact that we are the most prodigious creatures [sic. creators] and exploiters of external scaffolding on the planet." Consider, for example, the problem of driving to an unfamiliar location. Today this problem is easily solved with the aid of various resources, including maps, web-based services for printing maps and directions, and GPS devices that provide real-time instructions while driving. Lacking those resources, a driver may be able to navigate by reading street signs. But imagine trying to find a location without any such resources. An English-speaking person might experience this problem when trying to drive to a street address in Tokyo.

When people collaborate to solve complex problems, objects are often an important element of the problem solving process. Hutchins (1995a) proposed a theory of Distributed Cognition to describe how people and objects contribute to the solution of difficult tasks such as navigating ships into San Diego Harbor or flying a passenger plane (see Hutchins, 1995b). The instrument panels and the plane itself are essential elements of this latter task, and the approach of distributed cognition is to focus on how each person interacts with and interprets these objects.

Collaborative Problem Solving

Suppose that two different problem solvers are separately given exactly the same problem. Although the problem is the same, the problem solvers are different. Employing the model shown in Figure 1, they may have different internal representations of the problem, different general knowledge, different skills, and may select different methods. From the perspective of situated cognition, they may frame the problem differently because of differences in their values; one person may be highly motivated by the challenge of solving a problem, for example, and another person may consider the problem to be of little interest or importance. They may reach the same or different solutions, but the path to the solution is likely to be different for any reasonably complex problem. As the problem-solving sparrows exemplified, variation in the problem solving of different individuals is one reason for the success of collaborative problem solving.

Consider how problems are solved collaboratively. Because people differ in experience, knowledge, and cognitive abilities, they have different problem spaces even when working on the same problem. Indeed, if everyone had exactly the same problem space there would be little point in collaborating because they would all employ exactly the same methods on the same problem elements. Strategies for collaborating take advantage of individual differences in problem solving skills in different ways. In a strategy that Lazer and Friedman (2007) call parallel problem solving, all collaborators try to solve the same problem at the same time, sharing information, insights, and partial solutions as they proceed. Another strategy is to decompose the problem into sub-problems to be solved by different individuals or groups. A third strategy is to generate multiple potential solution paths that are explored by different individuals or groups. These latter two strategies permit aligning different parts of a problem with the expertise of different problem solvers. Small groups often decompose problems based on their mutual knowledge of group members' abilities and experience. Collaboration among many people on large complex problems often involves an organizational structure with leaders who maintain a high-level representation of the problem and allocate sub-problems to specialized groups.

When people collaborate on the same problem or sub-problem, they must construct a shared representation of the problem, tasks, and related knowledge, which constitutes a *joint problem* *space* (Roschelle & Teasley, 1994). The joint problem space can be viewed as the intersection of their individual problem spaces; it is the part of their problem spaces that is shared across all individuals. They construct the joint problem space by communicating with one another, often with the aid of external memory devices such as chalkboards, computer files, or paper notes. While collaborating, they diagnose and repair deficiencies in the joint problem space, particularly when someone acts in a way that suggests a different understanding of the problem. Large scale organizational problem solving requires construction and maintenance of many joint problem spaces corresponding to each of the sub-problems and to the integration of the sub-problem solutions into a complete problem solution.

Investigations of human problem solving have often studied how people play chess. This ancient game is appealing to researchers because it is a well-formed but potentially very hard problem with a very large number of alternative solutions. Consider how chess would be played collaboratively. Suppose teams of two people determined the moves of their chess pieces collaboratively. We can imagine that the two teams are in different rooms so they cannot see or hear the opposing team's discussions.

How might chess teams collaborate effectively? There are many possible strategies they could adopt. They could both consider each possible move of each chess piece, and discuss their evaluations of each move. This would, of course, be extraordinarily time consuming. They could both try to figure out a line of play, discuss the alternatives with one another until they had reached a consensus, and then continue discussing the pros and cons of each move as the play unfolded. This would also be time consuming, but much less so than an analysis of all possible moves. Alternatively, they could somehow attempt to decompose the game into sub-problems, but there is no obviously good way to decompose a chess game. Some options would include taking turns deciding which move to make, dividing responsibility for pieces on the left and right sides of the board, dividing responsibility between the pawns and the major pieces, dividing responsibility for strategy and tactics, or dividing responsibility for the opening, middle, and end game. None of these options for decomposing the play appear particularly attractive, which may explain why team chess is not widely practiced. Chess appears to be an example of a problem that does not greatly benefit from a collaborative approach that involves decomposition. Instead, chess is like the problem faced by the hungry sparrows; the benefit of having a team would be the increased likelihood of a more expert player. There appears to be little research, unfortunately, investigating the problem characteristics that determine whether collaborative problem solving is likely to be effective.

Collaborative Human Networks

Collaboration means working together toward a common goal, which people achieve by communicating with one another, coordinating their activities, and sharing information with one another (Grudin & Poltrock, 1999; Poltrock & Grudin, 1998). These three constituents of collaboration are tightly interconnected. People coordinate their communication with one another, they share information by communicating, and they may both communicate and share information in order to coordinate their activities. All these actions take place between two or more people, and these interactions between people define one or more networks. Networks can be defined by all acts of collaboration and communication. Social network theory describes these networks and provides methods and tools for analyzing them. The social network of a large problem solving organization may be of a hub-and-spoke or tree structure, but there may also be abundant interaction between the leaf nodes of such networks. Individuals often play essential informal roles in maintaining the flow of information within a social network.

We noted that collaborative problem solving involves the creation of joint problem spaces, and this activity is aided by external memory aids. We can extend social networks to include these memory aids and information repositories where people access and deposit information. The network that interconnects these artifacts and the people who use them may be represented as multiple layers similar to the knowledge network model proposed by Huang, Contractor, and Yao (2008).

Properties of these social networks influence problem solving (e.g., Hutchins, 1995a; Mason, Jones, & Goldstone, 2005). The network topology and the properties of each link influence the quantity, quality, and latency of the information that flows between any two individuals and/or information repositories. If people are geographically or temporally distributed, then technology underpins the social network, and the effectiveness of the technology network will determine the viability of the social network. Media richness theory (Daft & Lengel, 1986) claims that ambiguous and uncertain tasks such as those involved in complex problem solving require richer communication media and thus greater bandwidth and lower latencies. When rich media are not available, people adapt to the communication constraints by modifying their network and reducing the flow of information (Setlock, Fussell, & Neuwirth, 2004).

The Network Implications of Problems

We observed that network properties can influence problem solving, but conversely problem properties can influence network requirements, demanding the kind of adaptivity described by the adaptive coupling thesis proposed by Smart, Huynh, Braines, Sycara, and Shadbolt (2010). Some problems are decomposable into sub-problems that can be solved independently or perhaps with partial coupling, whereas other problems are not decomposable (such as the team chess example considered earlier). This section considers two approaches to analyzing the network implications of a complex problem: coordination theory and organizational analysis.

Coordination Theory Links Coordination to Dependencies

According to coordination theory, collaboration occurs in order to manage the dependencies among activities (Malone, Crowston, & Herman, 2003). Collaboration would be easy if everyone could work independently toward a shared goal without any need for interaction. In reality, collaboration requires managing the dependencies among people, processes, and objects, and coordination theory describes these dependencies and the mechanisms for managing them. The theory identifies three basic types of dependencies: flow, sharing, and fit. Flow dependencies arise when one activity produces a resource needed by another activity. For example, an engineering design activity necessarily precedes a stress analysis of the designed part. A sharing dependency arises when multiple activities require the same resource, such as a person, a machine, or a data object. A fit dependency arises when multiple activities produce some resource and all the parts of that resource must fit together.

Coordination theory offers an approach for modeling the process of solving complex problems collaboratively. Poltrock, Klein, and Handel (2007) modeled how changes are planned in a large aerospace program, and Poltrock and Klein (2007) modeled the military decision making process as described in the *US Army Field Manual 101-5* (Department of the Army, 1997). They identified the key dependencies, mechanisms for managing these dependencies, exception conditions that may occur, and procedures for handling these exceptions. Figure 3 shows a small part of the model of military decision making, with fit, flow, and share dependencies indicated by labeled arrows pointing rightward. The bold model elements are coordination mechanisms and methods for handling the exceptions encountered by these coordination mechanisms.

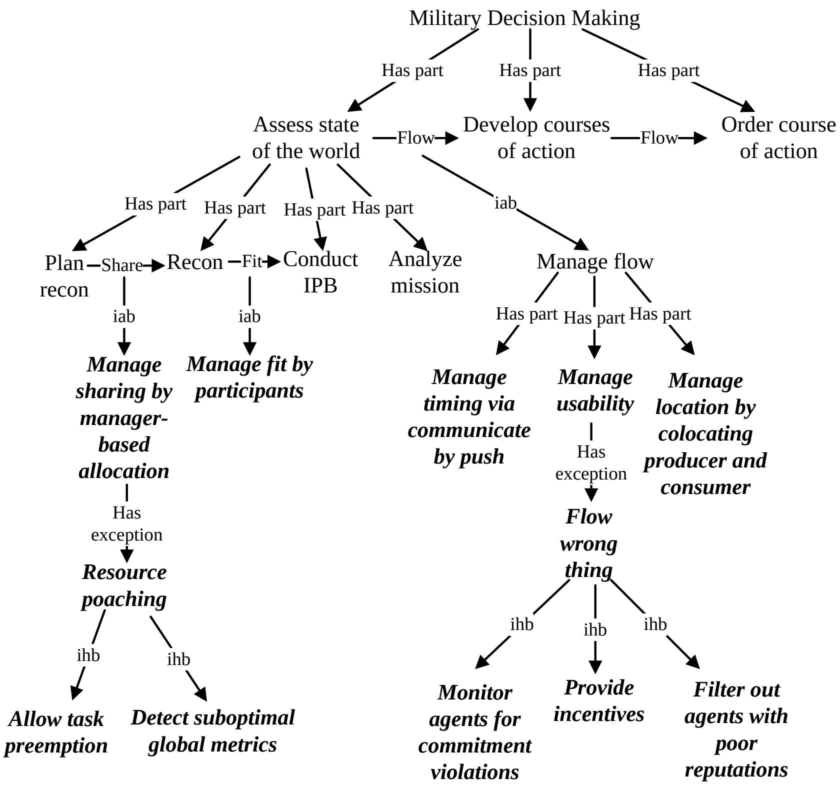

Figure 3: Part of a coordination theory model of military decision making.

These models capture how organizations solve complex problems such as deciding upon a military strategy, but they do not indicate who or what performs any of the elementary processes. The dependencies indicate the type of connection that must be provided between collaborators. A flow dependency requires a channel or medium for transporting the substance that flows from one collaborator to another; the model describes how the timing of the flow will be managed, how the quality of the delivered product (e.g., surveillance information) will be ensured, and where the flow will take place. These may be viewed as properties required of the network that connects the producer and supplier within this flow dependency. These models could be employed to anticipate when communication will be required between collaborators in organizational problem solving and to estimate required bandwidths. Software agents could use these models to recognize dependencies and either initiate coordination mechanisms or simply alert other human or software agents.

Problem Complexity Impacts Organizational Topology

The purpose of organizations is to divide and manage labor in a manner that efficiently coordinates the results of that labor. An organization is an ensemble of partially interrelated tasks and processes that produce goods and services. An optimal organizational structure is one that minimizes the dependencies between organizational elements. If the tasks and processes are minimally interrelated, then each element of the organization can work independently with little communication or interaction with other elements. In practice, of course, frequent communication is required within organizations because tasks are interrelated to some degree. If tasks and processes are highly interrelated, the demands of coordination are overwhelming and chaos ensues.

An organization intending to solve a large, complex, intractable problem (e.g., defeat a foreign insurgency or construct a new airplane) must decompose the problem into more tractable sub-problems, as Simon (1969) observed long ago. Ideally, these sub-problems should be independent so that the solution of one does not depend on the others. But as Simon observed, boundedly rational agents (such as people) cannot decompose complex problems into independent elements. Finding such a decomposition of the problem is equivalent to solving the problem, meaning that it is no longer complex or intractable. And solving sub-problems independently when dependencies exist cannot achieve an overall optimal solution.

Considering these challenges, what network topology or organizational structure is most appropriate for solving complex intractable problems? The traditional organizational structure is a hierarchy, similar to the military command and control hierarchy, but this management approach has been criticized for failing to benefit from the knowledge and expertise of lower levels in the hierarchy. Centralized networks provide good levels of coordination in the case of simple problems, while decentralized networks are more suitable for complicated problems (Leavitt, 1962). One of the features that may affect the relative suitability of network structures is therefore the complexity of the problem to be solved (see also Lazer and Friedman, 2007).

Many have predicted that organizations will move toward a decentralized structure in which activities are coordinated by competitive markets. In this model, an organization describes its requirements and others bid for the opportunity to perform the work. Marengo and Dosi (2005) concluded that this open-market structure cannot produce an optimal solution for complex problems. An organization cannot decompose the problem into independent sub-problems, and the open-market structure does not provide for coordination between suppliers. The higher the degree of decentralization, the smaller the portion of the search space that is explored and the lower the probability that optimal solutions are included in the space. They conclude that hierarchical organizational structures and communication are essential when solving such problems. Lazer and Friedman (2007) found that flatter organizational structures could employ parallel problem solving to solve complex problems, but only when everyone is interconnected via a network with a fast, efficient structure that allows for the rapid dissemination of information.

Consider the problem of developing a new passenger airplane. This requires the coordinated design of many systems and their components while optimizing flying range, efficiency, comfort, safety, and passenger capacity. A hierarchical structure is essential during the design and production of the first aircraft because the hierarchy provides a means of coordination across the dependent sub-problems. Once this problem has been solved and all of the interfaces between systems and components have been standardized, the ongoing production of the aircraft may benefit from a more decentralized structure. The standardization of interfaces removes the dependencies that would limit further optimization of the individual elements. Thus, the optimal network structure depends on the scale and complexity of the problem to be solved.

Network Construction

Some problems are solved by teams that are pre-defined, trained, and ready to perform. A team of fire fighters, for example, does not need to recruit new members before responding to an alarm. But new networks are formed to solve many problems, and here we consider two approaches to constructing such networks.

Enlisting Problem Solvers

One approach is to seek people with the expertise needed to solve sub-problems efficiently. In recognition of the difficulty and importance of finding such people, researchers have investigated expertise location systems. These systems build models of people's expertise and match this model to the requirements of the person seeking assistance (see McDonald & Ackerman, 2000, for a review and system architecture and Huang, Contractor, & Yao, 2008, for a recent example). In some cases, the models are based on explicit information about experience and expertise such as a resume, but more commonly the expertise system infers expertise from a person's activities, such as what web pages the person has authored or read.

Expertise location systems have proven highly successful in some contexts, such as a consulting company that builds a team based on the experience people have had on previous consulting activities. In this case, the expertise location system has good information about prior expertise and potential team members are likely to welcome being selected.

But expertise location systems are not always warmly welcomed. When an expertise location system was first introduced to an organization where one of us worked, several people in the organization complained that they did not want unknown people to call them with questions about their area of expertise. They did not want to be found by the expertise location system because they feared wasting time explaining their area of expertise to people who lacked even basic knowledge about the area. As Erlich (2003) and McDonald and Ackerman (1998) observed, in the workplace people balance a combination of social and organizational factors when deciding who is the most suitable and responsive person to approach, and these systems lack the information needed to make those judgments.

Of course, once experts have been found and located, successful problem solving requires knitting the experts into an effective network. They must decompose the problem into appropriate sub-problems, assign sub-problems to different experts, identify the dependencies between sub-problems, and manage these dependencies through appropriate and timely communication and information delivery. People must learn about the other members of the team, what they know, how they work, their communication styles, and their reliability. All these attributes play roles in the construction of a networked team.

Crowdsourcing

Another approach for constructing a problem-solving network is called crowdsourcing. In this approach, an open call for solutions to a problem is broadcast to everyone wanting to participate (the crowd), and people submit solutions. In some cases, the crowd also evaluates the solutions. This establishes a hub-and-spoke network for problem solving in the sense that everyone obtains the problem from a central source and reports their solutions to that central source. This approach has been successfully employed to solve a wide range of problems including geometric reasoning

(Jagadeesan et al, 2009) and improvements to online search (Ganjisaffar, Javanmardi, & Lopes, 2009). The problem of opening a well to obtain bird seed that Liker and Bókony (2009) posed to house sparrows can also be viewed as a form of crowdsourcing. The shortcoming of crowdsourcing, of course, is that each problem must be sufficiently small and tractable to be solved by at least one entity in the crowd. Crowdsourcing is an example of a decentralized structure coordinated by a competitive market, and consequently cannot produce optimal solutions for complex problems (Marengo & Dosi, 2005). This approach is not likely to be successful for large complex problems with many dependencies.

Conclusions

Many problems of modern life are too big and complex to be solved by a single individual; collaborative problem solving is essential. Indeed, we could argue that people engage in collaborative problem solving even in situations that may not, at first sight, appear to be collaborative. For example, obtaining food has been a problem throughout the history of the human race and remains a problem today for people in many parts of the world. In the industrialized world this problem is simply solved by visiting a local supermarket, selecting food items, and paying for them. We may feel that we are solving this problem without any assistance because all the components of the problem associated with planting, tending, harvesting, storing, processing, and transporting the food are hidden. However, from the standpoint of situated cognition, a vast network of people collaborates to enable us to have food in our homes when we need it.

One might argue that food production and distribution is not a problem because we have already worked out how to accomplish this in a general sense. Problems do arise, of course, when there is a food shortage, a transportation strike, or some failure of the system. Then we become aware of the network that underlies our solution and may have to find new solutions.

The example of food production and distribution illustrates that collaborative problem solving is not restricted to entirely new or novel problems. Indeed, many organizations thrive because they can consistently and systematically solve some class of problems. Software companies solve the problem of constructing and supporting new software products, aerospace companies develop new aircraft, and military organizations plan and execute missions. Each of these is an example of a complex, intractable problem demanding collaborative problem solving. Although organizations do not know the solutions to these problems, they know systematic methods, approaches, or strategies for solving the problems. These organizations contain experts in disciplines relevant to the problems they solve, and the organizations are structured into groups in ways that correspond to decompositions of the problems. Military organizations, for example, have separate branches for naval, air, ground, and space warfare. They also have separate divisions for logistics, reconnaissance, and operations. And they have established processes for solving problems. Software engineering is the discipline associated with the processes involved in constructing new software products, and military organizations follow the military decision making processes documented in the *US Army Field Manual 101-5* (Department of the Army, 1997).

Organizations have processes and structures for solving many frequently occurring problems. There are, of course, problems for which existing processes and organizational structures are suboptimal because the dependencies between problem elements do not align themselves with organizational structures, and there are problems that are too big to be solved by a single organization such as ending global warming. But organizational processes and structures are appropriate for the vast majority of societal problems; otherwise these organizations would be unsuccessful and replaced by others. The existence of processes and organizational structures for collaborative problem solving offers a way to investigate the network properties needed to support collaborative problem solving. Here we consider three network properties: topology, bandwidth, and resources.

Network Topology

The organizational structure defines, to some extent, the network topology. In a hierarchical organization, executives and managers coordinate the activities of the people who report to them, and managers report their progress and problems to the managers and executives above them to support this coordination. The purpose of this organizational structure is to manage the dependencies between the groups within the organization, and the organizational structure defines the most critical elements of the network topology, the flow of instructions and information up and down the hierarchy. As Marengo and Dosi (2005) observed, this management structure is the most appropriate when solving complex intractable problems.

An organization solving a large complex problem will decompose the problem into pieces and assign those pieces to different managers. They, in turn, decompose the problem further. This process continues until some individual or group of people is given a well-formed task with explicit goals and known alternative methods for achieving those goals. These groups of people communicate intensely as they construct a joint problem space and develop and evaluate problem solutions. The most frequent and highest bandwidth communication occurs within these groups.

Although the principal channels of coordination are through the management chain, the problem decomposition is imperfect and people must communicate between organizations. These channels of communication are less frequently used and generally require less bandwidth than those within a group.

Organizational structure and the network topology are emphasized here because they strongly influence the solutions that emerge. As Conway (1968) noted long ago, "organizations which design systems…are constrained to produce designs which are copies of the communication structures of these organizations." An organization with four engineering groups is likely to build systems with four major subsystems, and the interfaces between subsystems and their components will depend on the interpersonal communication that occurs between the members of these groups.

Network Bandwidth

According to Media Richness Theory (Daft & Lengel, 1986), richer media requiring greater bandwidth and shorter latencies are needed for ambiguous and uncertain tasks. Collaborative problem solving is generally most ambiguous and uncertain while exploring and seeking to understand an ill-formed problem. During this early phase of problem solving, participants characterize the problem and its potential solutions. They begin construction of a joint problem space, identifying those attributes of the problem that everyone must know and understand. They also consider ways of decomposing the problem and allocating sub-problems. They may begin building a core team with the expertise needed to solve or lead others in solving the problem.

An ill-formed problem is, by definition, ambiguous and uncertain, and the activities required to reduce the ambiguity and uncertainty require close coordination and communication among people with different expertise. Face-to-face meetings are often the preferred mode of collaboration in this phase, and if people cannot meet they generally seek high-bandwidth alternatives.

When the problem has been decomposed into sub-problems, the groups responsible for solving sub-problems are likely to engage in similar high bandwidth activities until they fully understand the problem and have determined a strategy for solving it. Bandwidth demands are likely to decrease when a strategy is identified, understood, and implemented. As people work toward solving and integrating the solutions of the sub-problems, there will be a greater emphasis on coordination and information sharing and substantially less ambiguity.

Network Resources

Constructing a joint problem space generally requires external memory services. A single distributed group or team may be satisfied with the electronic equivalent of chalkboards and flipcharts for capturing ideas as they explore a problem space. Larger problems that involve coordinated activity by multiple groups increase the demand for network resources. There may be multiple joint problem spaces, with each group working on a different sub-problem and leaders maintaining a joint problem space that integrates these sub-problems. Information repositories will be needed to capture these interrelated joint problem spaces and progress on sub-problems. Repositories that allow all participants to contribute and access information, such as wikis, are vital. These information repositories may play a key role in coordinating the work, tracking progress on each piece of the problem, delivering partial results to those who will benefit from them, and alerting people about both progress and any obstacles that are encountered.

Research Agenda

This paper has described theories of human problem solving and how these theories apply to collaborative problem solving. There are many different ways to collaborate when solving a problem, and characteristics of the problem are the most important determinant of which ways are most likely to be successful. Organizations solve very large problems using established methods and organizational structures. The way that people work together to solve a problem, including their organizational structure and processes, defines a network among those people and the resources they use. Communication flows intermittently over this network, with fluctuating demands for bandwidth and resources.

The relationships between problem types, solution strategies, organizational structure and processes, and networks could be used to predict when network communications will be needed, by whom, and their bandwidth requirements. We noted, for example, that intractable problems are best solved by hierarchical organizations, but other problems may best be solved by markets using a method such as crowdsourcing. These two approaches have very different network topologies and network requirements. Network science would benefit from a more systematic study of problem types, organizational structures, and the corresponding network topologies.

References

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, Massachusetts: MIT Press.

Conway, M. E. (1968) How do committees invent? *Datamation*, 14(5), 28-31

Daft, R. L. & Lengel, R. H. (1986). Organizational information requirements, media richness and structural design. *Management Science*, 32, 554-571.

Ehrlich, K. (2003). Locating expertise: Design issues for an expertise locator system. In M. S. Ackerman, P. Volkmar & V. Wulf (eds.), *Beyond Knowledge Management: Sharing Expertise*. Cambridge, MA: MIT Press.

Department of the Army (1997). *Field Manual 101-5: Staff Organization and Operations*. Department of the Army, Washington DC, USA.

Ganjisaffar, Y., Javanmardi, S., & Lopes, C. (2009). Leveraging crowdsourcing heuristics to improve search in Wikipedia. *5th International Symposium on Wikis and Open Collaboration (WikiSym'09)*, Orlando, Florida, USA.

Grudin, J., & Poltrock, S. (1999). Groupware. In J. G. Webster (ed.), *Encyclopedia of Electrical and Electronic Engineering*. New York, USA: Wiley.

Huang, Y., Contractor, N., & Yao, Y. (2008). CI-KNOW: Recommendation based on social networks. *Proceedings of the 2008 International Conference on Digital Government Research*, Montreal, Canada.

Hutchins, E. (1995a). *Cognition in the Wild*, Cambridge, MA: MIT Press.

Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.

Jagadeesan, A. P., Lynn, A., Corney, J. R., Yan, X. T., Wenzel, J., Sherlock, A., & Regli, W. (2009). Geometric reasoning via internet crowdsourcing. *SIAM/ACM Joint Conference on Geometric and Physical Modeling*, San Francisco, California, USA.

Kirsh, D. (2009). Problem solving and situated cognition. In P. Robbins and M. Aydele (eds.), *Cambridge Handbook of Situated Cognition* (pp. 264-306). New York, NY: Cambridge University Press.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52, 667-694.

Leavitt, H. J. (1962). Unhuman organizations. *Harvard Business Review*, 40(4), 90-98.

Liker, A., & Bókony, V. (2009). Larger groups are more successful in innovative problem solving in house sparrows. *Proceedings of the National Academy of Sciences*, 106(19), 7893-7898.

Malone, T. K., Crowston, K., & Herman, G. (2003). *Organizing Business Knowledge: The MIT Process Handbook*. Boston, MA: MIT Press.

Marengo, L., & Dosi, G. (2005). Division of labor, organizational coordination and market mechanisms in collective problem solving. *Journal of Economic Behavior & Organization*, 58, 303-326.

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, 137, 422-433.

McDonald, D. W., & Ackerman, M. S. (1998). Just talk to me: a field study of expertise location. *ACM Conference on Computer Supported Cooperative Work*, Seattle, Washington, USA.

McDonald, D. W., & Ackerman, M. S. (2000). Expertise recommender: a flexible recommendation system and architecture. *ACM Conference on Computer Supported Cooperative Work*, Philadelphia, Pennsylvania, USA.

Newell, A., & Simon, H. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.

Poltrock, S., & Grudin, J. (1998). Computer support for cooperative work and groupware. Tutorial presented at the *ACM Conference on Human Computer Interaction (CHI'98)*, Los Angeles, California, USA.

Poltrock, S. E., & Klein, M. (2007). A coordination-theoretic model of the military decision-making process. *1st Annual Conference of the International Technology Alliance*, Maryland, USA.

Poltrock, S. E., Klein, M., & Handel, M. (2007). Understanding process differences: Agreeing upon a single way to skin a cat. *4th International Conference on Knowledge Systems for Coalition Operations*, Massachusetts, USA.

Roschelle, J., & Teasley, S. D. (1994). The construction of shared knowledge in collaborative problem solving. In C. E. O'Malley (ed.), *Computer Supported Collaborative Learning*. Heidelberg: Springer Verlag.

Setlock, L. D., Fussell, S. R. & Neuwirth, C. (2004). Taking it out of context: Collaborating within and across cultures in face-to-face settings and via instant messaging. *ACM Conference on Computer Supported Cooperative Work*, Chicago, Illinois, USA.

Simon, H. A. (1969). *The Sciences of the Artificial*. Cambridge, MA: MIT Press.

Smart, P. R., Huynh, T. D., Braines, D., Sycara, K., & Shadbolt, N. R. (2010) Collective Cognition: Exploring the Dynamics of Belief Propagation and Collective Problem Solving in Multi-Agent Systems. 1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition, Maryland, USA.

Zajonc, R. B. (1965). Social facilitation. *Science*, 149, 269-274.

Zajonc, R. B., Heingartner, A., & Herman, E. M. (1969). Social enhancement and impairment of performance in the cockroach. *Journal of Personality and Social Psychology*, 13(2), 83-92.