Proceedings of the 1999 Computational and Mathematical Organization Theory Workshop. Ju-Sung Lee and Kathleen M. Carley (Editors). May 1-2, 1999. Cincinatti, OH.

INFORMATION PROCESSING

Organizational Structure and Perpetual Innovation: A Computational Model of a Retail Chain

Myong-Hun Chang
Cleveland State University
And
Owen Graduate School of Management
Vanderbilt University
Nashville, TN 37203
615-343-8094 (Voice)
615-343-7177 (Fax)
myong.chang@owen.vanderbilt.edu
http://www.csuohio.edu/changm

Joseph E. Harrington, Jr.
Department of Economics
The Johns Hopkins University
Baltimore, MD 21218
410-516-7615 (Voice)
410-516-7600 (Fax)
joe.harrington@jhu.edu
http://www.econ.jhu.edu/people/harrington

Abstract

Historically, Sears, Roebuck has been a relatively decentralized organization (Worthy, 1984) which, it has been argued, is a source of their success (Emmet and Jeuck, 1950). Analogous comments have been made to explain the superior performance of Wal-Mart (Collins and Porras, 1994). While anecdotes are many, there has been no formal analysis as to how organizational structure influences performance. Is there one organizational structure that is best? Or does it depend on a chain's environment? If so, what are the pertinent features of the environment and how do they influence performance? The objective of this research is to provide some theoretical insight into these questions by constructing a computational model of a retail chain in which corporate headquarters (HQ) and stores are constantly adapting to their respective environments.

Central to our work is the perspective that innovation is an act of information creation and that the performance of a retail chain depends critically on the way organizational structure influences the effectiveness of communicating and utilizing this information. The dimension of organizational performance that we focus on is the dynamic one alluded to in the description of Wal-Mart: the rate of improvement in store practices as the result of adopting useful innovative ideas. Given this conceptual framework, we model a retail chain as a two-level hierarchy consisting of a corporate headquarters (HQ) and a number of local stores, each of which is in a distinct market. Generation of innovative ideas is distributed within the organization: they are generated at the HQ as well as at the local stores. The chain's organizational structure determines where authority lies for deciding whether these ideas are adopted or discarded. A centralized organization is one in which the authority to implement new practices (ideas) rests at high levels in the hierarchy. As such, all new ideas that are deemed useful must receive the HQ's approval for the ultimate adoption at the local store level. A decentralized organization allows the store managers to independently implement new practices without the HQ approval. Our task is to characterize how organizational form influences the profit performance of a retail chain in both the short-run and long-run attained through dynamic adoption of new discoveries.

Our model of a retail chain begins with the view that a store, at any given point in time, is characterized by its current operating practices or what Nelson and Winter (1982) would refer to as "routines." The practices of a retail store determine, for example, merchandising, marketing, pricing, and personnel policy. The merchandising dimension encompasses issues of product selection, inventory, and contractual arrangements with vendors while marketing includes displays and advertising. Pricing may concern the decision of offering everyday low prices or sales. The personnel dimension entails the allocation of employees (for example, how many workers are on the floor), training programs, and compensation schemes. The profit of a store is presumed to be the result of the interaction between its current routines along these dimensions and the practices desired by the consumers in its market.

Given the set of embedded routines, an innovation is viewed as a new way of running a store as represented by a change in a routine. The dynamic of innovation consists of two primary processes: 1) *idea generation* **-** the process by which new ideas are created; and 2) *idea adoption* **-** the process by which new ideas are sifted through and either

adopted or discarded. While we fully recognize the importance that organizational structure may have on the incentives to generate ideas, the focus in this paper is on how organizational structure impacts the adoption process. The idea generation process as well as the flow of information through the chain are exogenously specified and assumed to be independent of organizational form.

A store's profit depends on how its current set of operating practices matches up with what is desired by its consumers. New ideas represent a new point in store practice space and associated with that new point is a level of profit. Store profit, being defined over this store practice space, then forms a landscape over which the store manager can search for better practices through a standard hill-climbing rule. In that markets are allowed to differ, the landscapes faced by different store managers can differ. Analogously, corporate headquarters searches over a landscape based upon chain profit. In that we presume it does not have detailed information about stores, headquarters is presumed incapable of selectively instituting practices but must instead mandate a practice chain-wide. Thus, *centralization involves a uniformity of practices, while decentralization allows for divergent store practices*. We compare the path of chain profit under these organizational forms. This investigation is initially carried out for when the market environment is fixed so that stores and the chain converge to a local optimum. The environment is then enriched to accommodate a changing market environment so that innovation is perpetual.

The relevant set of parameters for the computational model includes the degree of market heterogeneity, the length of time horizon, and the sensitivity of the consumers to store practices, among many others. For each set of parameter values, the computational exercise consists of 500 replications with each replication involving a randomly drawn vector of initial store practices and a sequence of TM new practices, one for each of the M stores in each of the T periods. Two classes of output data are collected: i) the *ex ante* optimal organizational form based on average chain profit over the horizon (averaged over the 500 replications); and ii) the frequency with which the organizational form is the *ex post* optimum (out of the 500 replications).

When the market environment is fixed over time so that the distribution of consumers in each of M markets is assumed stable, the following properties are observed.

Property 1: Centralization is more likely to outperform when markets are sufficiently similar, while decentralization is more likely to outperform when markets are sufficiently different.

The basis for Property 1 is the following trade-off associated with a more decentralized organizational structure. Moving authority down the hierarchy allows store managers to modify their operational routines over time so that each store's operation is reasonably well adapted to its local market environment. However, as stores migrate to different parts of their landscape, a new practice uncovered and adopted at one store is likely to be incompatible with the current practices of other stores in the chain. In essence, stores gradually come to target distinct consumers and this limits the extent of inter-store learning. This divergence slows down the rate of innovation as stores end up searching independently. In contrast, a centralized structure enhances inter-store learning by keeping stores near each other in terms of their practices so that they are targeting similar consumers. With these two countervailing effects, a decentralized structure outperforms only when markets are sufficiently heterogeneous.

Property 2: Centralization is more likely to outperform over short horizons while decentralization is more likely to outperform over long horizons.

An important factor in this result is that the global optimum under decentralization has higher chain profit than the global optimum under centralization. A decentralized structure allows for the possibility that each store achieves its global optimum by exactly tailoring its practices to what is desired by its consumers. Centralization rules out that possibility by virtue of mandating common practices. In that achieving a global optimum is more likely with a longer horizon, decentralization tends to perform relatively better as the horizon increases.

Finally, the analysis is extended to accommodate the possibility of a changing market environment by allowing the landscape to shift continuously due to movements in consumer preferences. How does a continually changing environment alter the relative performance of these organizational structures? In this setting, it is no longer appropriate to imagine the chain or stores converging to a local optimum because the set of optima will itself be continually changing. Measuring the performance of the chain by the average profit along the steady-state profit path, the following property is observed.

Property 3: Centralization is more likely to outperform decentralization when market fluctuations are sufficiently large, while decentralization is more likely to outperform when market fluctuations are sufficiently small.

When the analysis is extended to accommodate the possibility of a changing market environment, the short-run superiority of a centralized retail chain is extended to the long run. With the market environment continually changing, the spillovers between stores that is promoted under centralization becomes a dominant force. Contrary to the usual claim that volatility in markets requires a more flexible decentralized organizational form, we find that it is the centralized organization with coordinated search that is more effective in responding to a new environment.

References

Collins, J. C. And J. I. Porras, *Built to Last: Successful Habits of Visionary Companies*, New York: HarperBusiness, 1994.

Emmet, B. And J. E. Jeuck, *Catalogues and Counters: A History of Sears, Roebuck and Company*, Chicago: The University of Chicago Press, 1950.

Nelson, R. R. And S. G. Winter, *An Evolutionary Theory of Economic Change*, Cambridge, MA.: The Belknap Press of Harvard University Press, 1982.

Worthy, J. C., *Shaping an American Institution: Robert E. Wood and Sears, Roebuck*, Urbana, IL.: University of Illinois Press, 1984.

Agent Modeling of Information Assurance - Results

Rob Nehmer
Berry College
Campbell School of Business

The American Institute of Certified Public Accountants (AICPA) is involved in the development of a major new classes of services called assurance services. These services can be understood as a generalization of the traditional audit service which attests as to the "fairness" of the financial statements. Assurance services will "assure" the information consumer of the quality of the information signal sent from any business information source. The problem has been modeled by the AICPA as being composed of the following components: business risk, decision relevance, and systems reliability.

Last year, I presented a conceptual overview of a simulation experiment involving CPAs and assurance services. In that model, a simplified situation was developed in which some agents would be producers of information and other agents would be consumers of information. This type of situation actually models capital markets well, where companies are producing annual reports (information production) and investors are evaluating the reports prior to making investment decisions (information consumption). The CPAgent in the model acts in one of two modes. In the first mode, it can increase the reliability of information production and thereby increase information efficiency. In the second mode, it can promote the relevance of the information to the information consumer, thereby increasing information effectiveness.

The model was given some additional theoretical interest by appealing to Habermas' Communicative Action Theory. In this theory, communication between individuals (agents) is seen as being influenced by four validities: performance, sincerity, legitimacy, and truth. Performance validity refers to the technical construction of an utterance, in this case, to the construction of an informational signal by a information producing agent. Sincerity refers to fit of the signal to the information environment into which it is communicated, that is, that the signal is complete and understandable. In the model, these two validities are affected by CPAgent's working with information producers. A signal is considered legitimate in Habermas' theory if the receiver of the information believes the sender has the social standing to send the signal. This is akin to investors recognizing valid investments. Truth validity refers to the perceived truth of the information signal by its consumer. Legitimacy and truth are two validities that the CPAgent can influence on the information consumer's side.

The paper presents results of simulations run using the above model. Factors that were varied include the information loss function from period to period and the costs of the CPAgents services. Service costs were differentiated between information production assurance costs and information consumption assurance costs. Reputation affects for all three types of agents were also varied. The results indicate a stable model under conditions which are both interesting and realistic. The final section of the paper presents an extension which addresses modeling agents as both information producers and consumers and CPAgents as having aptitude in both efficient information production and effective information consumption activities.

Risk Perception Propagation: A Validation of Virtual Design Team Methodology in Risk Management of Fast-Track Projects

Christopher Gordon
Construction Engineering & Management,
Department of Civil and Environmental Engineering,
Stanford University, Stanford, CA 94305-4020

Abstract

As project management complexity increases with decreased allowable project delivery time, the capacity of individuals to process information increasingly becomes a limiting factor in the effectiveness of project organization information flow. At the same time, this degree of activity interrelationships and time-to-market constraints creates the need for improved control techniques and risk management policy. Risk management efforts, however, are only as effective as the measure of risk they seek to mitigate. When the transmission of risk information is constrained, the ripple effect of incomplete mitigation can continue throughout the project as decisions continue to be based on incomplete information. A key element of project-related information flow, therefore, is the degree to which risk-based information is communicated. By creating a virtual organization prototype, one can simulate potential outcomes based on the organization's information-processing capability. The Virtual Design Team organizational model incorporates individual workload capacity with a project schedule, its activity durations, and requirements, and can therefore serve as a valuable instrument of risk perception propagation analysis.

Statement of Assumptions and Definitions

For purposes of this analysis, risk is defined as an undesirable event--not the probability of an undesirable event--which can be propagated within a activity sequence system (Niwa et al,1979). Risk-based performance is a product of sets of rules, knowledge, and skills applied to a perception of risk; risk perception is in turn a function of risk information and the time available to process this information. Actors are assumed to have finite rationality (March and Simon, 1993), and congruent goals.

Description of Contribution

Risk events and risk mitigation can be incorporated into a CPM project network to model risk event-related precedence relationships and activity duration. Risk events can be considered milestones of zero duration in this case. The Virtual Design Team model simulates the workload of individual actors relative to their capacity and project requirements over the course of the project. Activities such as risk and risk mitigation can support information exchange and failure dependency relationships, both affected by the relevant actor’s workload. Three simulation outputs significant to risk perception processing are a Gantt chart modified from the CPM baseline, a statement of actor backlog, and a heuristic assessment of communication risk—the measure of what percentage of communication is likely to be missed over the course of an activity. From a risk manager’s standpoint, simulating incidents of risk within a schedule network operationalizes the notion of risk as part of a sequence of events whose length, timeliness, and quality depend on the workload capacity of constituent actors. At the same time, this approach simulates and eases diagnosis of organizational risk readiness.

Conclusion

Organizations daily encounter risk events of varying magnitude such as design error, ball-in-court delays, and external or environmental incidents. The value of Virtual Design Team methodology in consideration of risk perception propagation is in real-time assessment of organizations' risk information processing capacity. By incorporating risk events into an existing CPM sequence, one can gauge the preparedness of the organization to respond to the events and minimize the impact on the overall network. This provides the opportunity to diagnose and address points of concern within the organization-project interface and minimize the impact of risk events.

References

March, J.G. and H.A. Simon (1993), *Organizations* (2nd edition). Cambridge Blackwell Publishers (1st edition 1958).
Niwa, K.(1989), *Knowledge-Based Risk Management in Engineering: A Case Study in Human-Computer Cooperative Systems.* New York, NY: John Wiley & Sons.
Thomsen, J., J.C.Kunz and R.E. Levitt (1998d), *A Proposed Trajectory of Validation Experiments for Computational Emulation Models of Organizations*. CIFE Working Paper #47, Stanford University.

SOCIAL NETWORKS

Modeling Relationships among Multiple Graphical Structures

Kari Chopra and William A. Wallace
Department of Decision Sciences and Engineering Systems
Rensselaer Polytechnic Institute
Troy, New York, 12180-3590

Correspondence should be addressed to:
William A. Wallace
Department of Decision Sciences and Engineering Systems
Center for Industrial Innovation 5015
Rensselaer Polytechnic Institute
Troy, New York 12180-3590
wallaw@rpi.edu

Abstract

Mathematical models of organizations frequently rely on graphical structures, such as social networks and cognitive maps, to model the relationships within a group and the mental models of its individuals. Questions frequently arise that require the analysis of multiple networks, graphs or maps. How can we analyze three-dimensional social network data; e.g., when we elicit from each member of a group their view of the social network? How can we compare and combine the mental models of the members of a group? How can we measure an individual's cognitive position within a group, i.e., the location of one's mental model in relation to the mental models of others in the group? Moreover, how can we test theories linking social position to cognitive position? The purpose of this paper is to define a mathematical model for capturing the relationships among multiple graphs and to demonstrate its application to several topics in computational and mathematical organization theory.

Three-way social network data is obtained when one performs multiple measurements of the social network on a single group of actors. This most frequently occurs when one obtains a series of measurements over time, or when one elicits a network from each actor of the group. In the latter case the data may be referred to as a cognitive social structure, and a primary question of interest is the identification of a consensus structure. In other words, given a group of actors, each with their own view of the social relationships among the group, what is the consensus view among the actors regarding the true social network for the group? Solution approaches frequently involve finding a central element, or identifying a social network that is central in some sense to the collection of individual networks.

Studies of organizational performance frequently cite a link to the degree to which its members exhibit similar or shared mental models. However, the measurement of differences between mental models and identification of a shared mental model of a group are still open questions. Mental models are frequently represented as graphical structures, such as cognitive maps, concept maps, influence diagrams, semantic networks, or belief networks, because they are flexible and rich enough to capture the elements of a mental model and the relationships between them, as well as precise enough to permit rigorous mathematical analyses. In such cases, differences between mental

models are frequently characterized by defining a distance function on their graphical representations, and shared mental models are constructed by determining a consensus representation, frequently focusing on finding a central element with respect to the graphical mental models of the group.

It has been hypothesized that cognition is related to social structure; i.e., that one's position in the social network affects what one knows. Pertinent questions involve not only pair-wise comparisons among actors, e.g., is the similarity between their mental models related to their distance in the social network; but also more holistic questions that attempt to link an actor's social position to their "cognitive position"; i.e., to how their mental model fits in with the mental models of the other actors. However, adequate measures of cognitive position and methods for relating it to social position need to be developed. This line of inquiry could be greatly enhanced by models that capture the relationships among mental models within the structure of a network or graph; i.e., a "cognitive network". Such a model would be amenable to conventional analyses of position within a network, as well as investigations of correlations between characteristics of the social network and characteristics of the cognitive network.

The present research approaches these issues by developing a mathematical modeling technique that consists of the following components. First, we define a set of graph metrics, functions that compute a distance measure for a pair of graphs, for modeling the differences between graphical structures, including social networks and graphical representations of mental models. Second, we define a new construct, the meta-graph, that captures the relationships among multiple graphical models within the structure of a graph or network. Third, we develop techniques for generating consensus structures from a collection of graphical models by identifying central elements of the meta-graph. We then proceed to illustrate how the meta-graph model and associated solution techniques may be applied to problems in organizational theory such as the construction of consensus social structures, the analysis of multiple mental models, the modeling of a cognitive network, and the analysis of correlations between the cognitive network and the social network

Characterizing Organizational and Social Networks: Some Issues and Findings

Ju-Sung Lee
Social and Decision Sciences
Carnegie Mellon University
Email: jusung@andrew.cmu.edu

Abstract

Social network analysis is a research methodology in which the focus or unit of analysis is the relation between two individuals rather than the individual alone. Often, behavior and dynamics of human systems cannot simply be gleaned from knowing how each individual operates. How each individual relates to others and how the interactions evolve are crucial to the understanding of macro-level social phenomena. In fact, the relationship between the individual and the structural focus is complementary. Often, we can infer individual traits by discovering the people to whom s/he relates. And conversely, we can make more robust statements about the macro-system when we understand how each individual fits into the scheme.

As with any measurement system or methodology, there are levels of granularity. Social network analysis is no exception. Measures have been devised to capture individual structural properties as well as system level properties. For example, the most common known property is degree centrality: with how many people do you have a particular kind of relationship, whether that relationship be one of friendship or advice or enmity? The measure applies to each individual; that is, each individual has his or her own degree centrality. This particular measure yields several interpretations: power, popularity, influence, flow of information, etc. An individual with many friendship ties is often described as popular or prestigious. If these ties are economic or political, he or she may be said to have a certain degree of power or influence. The aggregate version of this measure, called centralization, describes the distribution of the degree centrality measurements. It is this kind of measurement that will be the focus of this research.

The research on social network analyses has yielded a host of measures through which investigators can describe and characterize extant networks. Often, these measures provide an avenue for common interpretation. That is, given a particular context, the ramifications of these measures are deducible and often point to a small set of conclusions. However, there is little work in understanding the exact nature and importance of the measures and also patterns in the measurement of multiple networks. Furthermore, the issues surrounding hypothesis testing of these various measures remain to be fully explored. How networks and related measures are distributed and the

positions specific networks hold in their respective distributions are relevant issues for significance and hypothesis testing of social networks. The research proposed here also endeavors to address these issues.

Depending on the choice of a null-model, the issue of hypothesis testing can be easily or not so easily addressed. The first and obvious choice is the use of the random network for one kind of hypothesis testing. Specific networks can be compared to a distribution of random networks controlling for size and density; distributions of networks for different sizes and densities vary so these must be held constant. This type of testing can be used on most network-related measures such as reciprocity, asymmetry, triads, etc. Given enough data sets and enough variability across measurements, it might be possible to construct indices that adequately differentiate types of networks. For instance, friendship networks may have reciprocity levels that significantly depart from those obtained from random networks. Advice networks may exhibit strong asymmetry. And so forth. When network-level measures are meaningful, how some networks are more similar along these dimensions, and why are questions that still have yet to be systematically answered. This paper expands upon some ongoing research in the characteristics of network-level measures (Anderson et al 1998) by examining commonly-used network measures and their

significance in a host of empirical social networks, both organizational and non-organizational. Furthermore, the research attempts to uncover a typology of these networks using various multivariate techniques to explain the variance found across a specific set of measures including degree, closeness, betweenness, and eigenvector centralizations, transitivity, size, and density. I find that the analysis yields some insights into the boundaries between empirical networks and shows how often these boundaries are not quite intuitive.

I use a certain subset of network level measurements to characterize a set of twenty-five empirical networks. These networks vary considerably in size, origin, and type, from marriage unions of Florentine families in Renaissance Italy to African mining workers to University of Michigan students. I use the same set of networks in determining the significance of various network measures. These networks have been obtained from the UCINET network software and the INSNA web site; the latter correspond to networks found in the text, Social Network Analysis, by Wasserman and Faust (1993).

The literature and common-sense lead us to hypothesize networks of similar types will be similar statistically speaking. That is, friendship networks will have the same characteristics and have low distance ratings in the clustering exercise. Networks embedded in similar contexts should also be similar to one another. With this research, I have answered some questions and opened up a host of others. I find that unusual networks, or networks with significant measures, whether by context or measurement, are consistently and strongly unusual, so unusual that they drive some of the analyses. I have also found evidence for similarities in networks that are on the verge of undergoing a change process, in this case, one of conflict. Through all this, I discover that networks are characterized by the extent to which they have a popular elite (i.e. high centralization). Furthermore, while I can explain quantitatively the reasons why some fall in one group and others in the polar opposite group, the lack of common contexts across the networks give us little material with which we can base a strong theory. Certainly, more descriptive, preferably quantitative, information on each of these networks would have aided us in our deductive process. Also, by testing these networks against a simple null-model, I am able to not only characterize general classes of networks but also identify certain networks whose dynamics of formation may radically different from those typical of the entire set. Specifically, I find that cognitive networks, networks situated in special environments (e.g. prisons, mining facilities), and those constrained by cultural norms and boundaries (e.g. Florentine families) exhibit similar patterns in the significance levels of various measures. Understanding how these measures map onto the special dynamics is not obvious and requires further study.

This research represents a preliminary exploration into the usefulness of significance testing and multivariate categorizations. Clearly, we see that such methods have the potential for assisting the construction of a social network typology despite the use of a limited set of measures and only one, simple null-model as a basis of comparison. Other stochastic models such as P1 can and have been used to study significance of network dynamics and properties. The power of the explanation can be most enhanced by a confluence of such tests.

References

Anderson, Brigham, Carter T. Butts, and Kathleen M. Carley, 1999. "The Interaction of Size and Density with Graph-Level Measures". Working paper.

Wasserman, Stanley and Katharine Faust, 1994. *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press.

PROCESS ALGEBRAS AND NETWORK DYNAMICS

Corrado Pasquali
University of Genova
e-mail: pasquali@mclink.it

1. Overview

This presentation is based on the adoption of a "structuralist" perspective in explaining essential features of aggregate behaviors resulting from the interaction of dispersed, individual elements. This perspective emphasizes the organizational nature of individual interaction (who interacts with whom, which information flows are created...) its main claim being that the features of aggregate behaviors crucially depend on the interaction structures that constraint and permit individual interaction. My main aim is explaining and defending the thesis that process algebras can be effectively used a tool to represent and analyze the aforementioned problems in terms of network dynamics. In particular, I will focus my attention on pi-calculus (Milner et al. [1992]) and on the abstract notion of concurrency. Painted with a very broad brush, my claim is that process algebras can be conceived of as being an abstract, information theoretic approach to modeling organizational interaction and network dynamics.

2. Spatially Located Processes

A great deal of attention has been devoted to the problem of representing, analyzing and modelling coordination processes in decentralized environments in which global dynamics result from the interaction of a multitude of dispersed components in the absence of a central controller. An extremely interesting research hypothesis is that which emphasizes the role of interaction structures in inducing and determining the aggregate outcomes of individual actions. According to this hypothesis, the problem of the organizational nature of social interaction should be brought to the fore and the processes through which interaction channels, information flows and patterns of interaction are created, used and discarded should receive a detailed analysis. In the realm of economic theory, the need for a detailed explanation of how large scale, stable networks are induced (or “emerge”, as the saying goes) from initially random interaction, has been emphasized in some recent works by Alan Kirman. His claim that an economy is better analyzed as an “evolving network” (Kirman [1997]) is grounded on a strong argument against the idea, which lies at the very core of general equilibrium theory, that economic interaction takes place among anonymous agents that react in isolation to quantitative signals coming from the market through the price system. Kirman holds to the view that economic interaction has an essentially local nature that, in a sense, makes the consideration of its “spatial” character unavoidable. In this sense, the different ways in which agents interact, the capacity of agents to communicate, to observe the actions of others and to learn from these and the fact that interaction takes place through networks are problems that need to be considered in order to tackle the problem of micro-founding aggregate behaviors. Based on these considerations, a multitude of different perspectives can be considered: one can imagine, for instance, that the behaviors of single agents or classes of agents can be stochastically affected by those of others, or, in a second place, one can imagine that the specific underlying network structure is relevant in terms of the collective behavior resulting from individual actions. A first and very roughly cut distinction can be given with respect to the nature of interaction that can be global (i.e. every agent behavior is affected by that of every other agent) or local (i.e. every agent’s behavior is affected by a subclass of other agents) and with respect to the nature of interaction structures that can be fixed (and possibly given ex ante) or flexible (and possibly changing upon interaction). As stressed by Kirman, the most interesting problem is that of considering the case in which “interaction changes the nature of the relation between individual and aggregate behavior and the nature of aggregate behavior itself” (Kirman [1997], p.341). Not only: the other side of the same problem is considering the possibility that the very nature and architecture of network structures (i.e. communication channels, interaction structures) are affected and endogenously change upon interaction. In this sense, it is worth considering the possibility of approaching the problem of how “not only behavior changes as agents interact through a network, but also how networks themselves evolve” (Kirman [1997], p.346). This is exactly where existing approaches (Markov random fields, filtration models, stochastic graphs) show their main limit, that is: expressing how “links” (and that is: “structure”) change over time as a consequence of interaction.

3. Computation and Organization

This is the limit that, in my opinion, can be overcome by a computational approach to dynamical networks based on process algebras and concurrency theory. My thesis is to consider networks (and, broadly speaking: “organizations”) as grammars of interaction, that is as structures that rule and constraint their components interactions by allowing only specific subsets of interactions out of a multitude of possible ones. Now, the main appeal of this thesis with respect to the idea of endogenously evolving networks, is that not only a grammar is a set of rules, but also a generative mechanism. In this sense, a network or an organization can be viewed as a class of rules and constraints that acting upon a set of processes induce an endogenous motion in a dynamic space of networks. As a grammar (or, more precisely, a syntax) generates a language from a set of elementary building blocks and a set of rules for combining them into more complex expressions, an organization generates new configurations of processes and new organizational functionalities. More to the point, it is noteworthy that, just like a (formal) syntax, an organization might be thought to be closed under the operation of "construction". That is: functionalities, processes and configurations of processes constructed by the rules from old ones still belong to the organization. The grammar (i.e. the organization) can thus be thought to be the theory of the structure(s) of the language it induces and constructs. It is fairly evident how concurrent programming (and, broadly speaking, the theory of concurrent processes) being a theory of how component programs are executed in parallel giving rise to global information processing tasks might represent a rigorous computational tool to model, in a rich and constructive fashion, the relations between local interactions and global dynamics arising from them. With respect to this problem, it is possible to hold to the claim that computation theory, being the theory of possible constructions, can give powerful insights on how to overcome the problem of "observing and measuring the quantitative change of relevant state variables of a system once these are given together with their couplings and with the law that describes their dynamic evolution in time" and on how to approach the problem of "how those variables and couplings are constructed and originated from the dynamics of the system itself". In my opinion, this is the point that show how a computational approach to organizational phenomena can parallel and extend analytic approaches based on differential equations.

4. Pi-Calculus

Pi-calculus, in particular, is probably the best known “mobile process” calculus. Its firs attractive feature with respect to the problems at hand is its clear and simple syntax whose primitives are just processes that interact through channels connected by complementary ports. The second, and perhaps even more important feature of the calculus, is the fact the it can express how the topology of a network evolves dynamically upon the interaction of individual processes. This feature is obtained thanks to the elimination of any passive/active distinction at the level of the language and to the consequent possibility of directly expressing processes that create new channels/links and exchange channels (i.e. interaction possibilities) with one another. In this terms, the concurrent execution of a program (i.e. the accomplishment of an information processing task) is represented in terms of processes that perform a computational task according to a set of rules for communication . Computations in this sense, are viewed as a kind of topological objects that follow specific rules and laws.

References

Kirman, A. [1997] "The Economy as an Evolving Network", Journal of Evolutionary Economics, (1997) 7: 339-353.

Kirman, A., Oddou, C., Weber, S. [1986] "Stochastic Communication and Coalition Formation", Econometrica, Vol.54, No.1 (January 1986), pp.129-138.

Milner, R., Parrow, J., Walker, D. [1992] "A Calculus of Mobile Processes" (part I & II), Information and Computation, (1992) 100: 1-77.

ORGANIZATIONAL DESIGN

Exploring “Complex” Organizational Designs

Tim Carroll
Fuqua School of Business
Duke University
Durham, NC 27708-0120
tnc1@mail.duke.edu

Richard Burton
Fuqua School of Business
Duke University
Durham, NC 27708-0120
rmb2@mail.duke.edu

Currently visiting at: Odense University

Organizational adaptation to changing environments remains a central theme of Organization Theory. This paper seeks to explore how developments in Complexity Theory may shine new light on traditional approaches to this issue. Kauffman's (1993) NK model has been used to explore the dynamics of the adaptation of systems. However, the model has rather strict uniformity assumptions regarding the elements in the system (i.e., the actors and the links between them). The model is extended here by drawing on the communication network research of Bavelas (1948) and Leavitt (1952) to build a more realistic, organizational NK model. Using the VITE software program, we varied attributes of both the organizational units and the linkages between them to determine whether and how the extended model differs from the original.

(Key words: Complexity theory, NK models, organizational design)

Introduction

The issue of how organizations adapt to their environment remains a central concern of Organization Theory. Lawrence and Lorsch (1969) found that organizations operating under conditions of environmental uncertainty performed best when they increased both the level of differentiation (defined as differences in formal structure and managerial orientation) and integration (defined as the quality of collaboration between managers and departments). Recent work in Complexity theory also explores the dynamics of systemic adaptation to changing environments. One of the most intriguing findings from work in this area is that the most adaptive systems exist in a precarious balance known as “the edge of chaos”. In this state, systems are able to find not optimal, but very good solutions to complicated problems in a relatively short period of time. They are able to change behavioral and structural patterns in order to adapt to new stimuli.

Theoretical Motivation

Many studies in Complexity Theory find that systems undergo a phase transition between order and disorder, between stability and chaos. Thus, systems exhibit one of three types or ranges of behavior: ordered, complex, and chaotic. In the ordered regime, elements of the system are sparsely connected and have a tendency to ‘freeze’ in fixed states. Since each element is connected to few, if any, neighbors, then minor alterations are unlikely to have a major effect on the system, due to dampening feedback. This condition encourages order, but discourages adaptation. In contrast, there is little order to be found in the chaotic regime. Since each element is connected to many neighboring elements, then any small change will quickly diffuse throughout the entire system. In this case, understanding and prediction of macro effects from micro inputs is difficult to impossible, respectively (due to amplifying feedback).

The transition between chaos and stability is where the most ‘complex’ behavior occurs. It is in this boundary region where systems are able to optimize the conflicting demands of task performance and capacity to evolve. Importantly, Kauffman (1993) found that this phase transition between order and disorder may be the result of a control parameter, such as the number of inputs per variable or the structure of the organization of connections.

The motivating question here is what implications these new findings have for organizational design. Complexity Theorists posit that many, if not most, natural systems spontaneously self-organize to the edge of chaos. If so, is it possible to design an organization that will achieve this state?

The approach that we take in this paper starts with Stuart Kauffman's work on NK fitness landscapes. Fitness landscapes were first conceived by Sewell Wright (1931), a biologist working in the early part of this century. He developed the notion of a topological space where a given location on the landscape stands for a genetic element and the height of the peak at that location corresponds with the relative fitness contribution of that element. For example, certain species of birds with the attribute of long beaks would be more “fit” if this attribute improved their ability to forage for food. Landscapes are said to be “rugged” when attributes are correlated, that is, when the characteristics of attributes carry implications for other attributes, their neighbors on the landscape.

Wright's work was extended by Kauffman (1993) into his NK model. In this model N stands for the number of elements in the network of attributes. K means the number of inputs that each N receives. It is the number of connections for any given single element, N. Thus, the contribution of a given attribute of the organism (organization in our case) to the organism's overall fitness is assumed to be influenced by K other attributes.

Kauffman finds evidence of "edge of chaos" behavior in his NK fitness landscapes. When the level of connection between elements in the system is too high (as N → K) then good solutions are never founds since each individual change causes multiple changes throughout the other elements. Similarly, too little connection (closer to K=0) and poor compromises are never challenged.

Lawrence and Lorsch would say that as a problem becomes more complex, organizations ought to divide the work into smaller, self-contained units, and work at integrating these units together. Thus, work needs to be divided, but each unit needs to collaborate for the greater good of the organization as a whole. They never explicitly state how much integration needs to take place. Should every sub-unit coordinate with every other one? What is the minimal amount of coordination needed and what amount is too great, and thus counter-productive?

NK models give a surprising answer to these questions. The finding is that complex problems are best divided into a number of semi-autonomous sub-units (much the same as Lawrence and Lorsch's study found). However, as N increases (the number of self-contained units), the system performs best as K (the number of connections) is kept low. In essence, this implies that the greatest good occurs when each unit ignores a most of the other units and concentrates on its own welfare.

The finding from the work of Kauffman, among others, is that generally systems, and by implication organizations, best adapt using these intra-organizational units. The size of the units is key. If you have too few then the system will have a high degree of coordination or integration, but not enough differentiation to explore alternatives. Conversely, if there are too many units then there is too much differentiation and not enough integration. Any solution that an individual unit finds will not diffuse through the organization. Thus, from an organizational design standpoint, the task is to choose the appropriate level of both N and K.

Research Design

Kauffman's work is interesting, but not realistic in an organizational sense. Most importantly, there is a very strong uniformity assumption. There is no variety of either actors or linkages. We begin by relaxing this assumption and building an NK model from non-uniform organizational building blocks. Our simulation model is built on the VITE simulation platform. VITE is a robust modeling environment that allows the researcher to build a project organization and specify a variety of actor attributes (such as skill set and experience) and linkage types (such as information exchange and failure dependency).

Our initial task is to build a VITE model that replicates the experimental results obtained by Bavelas (1948) and Leavitt (1952) in their experiments on types of communication networks (circle, wheel, Y and chain). Next we "calibrate" the model to ensure that the behavioral results generated by the simulation are roughly equivalent to those found by Bavelas and Leavitt in their experimental studies. Once this done we vary the quality of the interaction between the units, the "K".

What this approach allows us to do is to address an interesting and topical question – how to most usefully divide and coordinate a complex task. We develop a theoretically grounded model of a network of organizational sub-units, based on classic organizational communication studies. We first validate the model in terms of these communication structures and then "scale up" the model to explore the dynamics of an explicitly organization NK model. Studying a simple, plausible organizational system like the one outlined above may give us insights or an understanding that is difficult to comprehend at the larger level.

Implications

We expect that our model will yield results generally consistent with previous studies of NK models. However, we expect that changing the attributes of N and K will modify the previous findings in important ways. It is not known, for instance, what effect changing the type of organizational linkage will have. For example, introducing failure dependencies between organizational units could either encourage more stability or instability. Likewise, some character attributes, such as a higher skill set or experience level, may make higher levels of connectedness (more linkages) possible (i.e., act as dampening feedback). We expect to find evidence that there are "optimal" general levels of interconnectedness, but these change as a result of the type or character of the connection. If they do not change, or change in surprising ways, then we will have an opportunity to suggest revisions to the standard organizational theories and to develop new questions for further research.

Process Integration in Enterprise Engineering

Katy Atefi, Michael Grüninger, Mark S. Fox
Department of Mechanical and Industrial Engineering
University of Toronto
Toronto, Canada M5S 3G9
gruninger@eil.utoronto.ca, msf@eil.utoronto.ca

Enterprise design knowledge is currently descriptive, ad hoc, or pre-scientific. It is a collection of heuristics which are not applicable in all circumstances. We want to define a theory of enterprise design by discovering its underlying principles. We want to understand why different approaches and techniques work for certain enterprises and why they fail for other enterprises. What is needed is a distillation of the principles for enterprise design implicit within the heuristics, and the formalisation of these principles as logical theories. We can then test, compare, and validate different enterprise design theories.

This paper will focus on our research into the development of a science of Enterprise Engineering, an emerging discipline that focuses on the modelling, design and analysis of enterprises. The process of designing an enterprise is an exploration of the space of design alternatives where configurations of activities, processes, resources, information, organisation, authority, behaviour, incentives, etc. are evaluated with respect to how they affect the efficiency and quality of the enterprise's performance.

Much research to date has focused on the tangible aspects of enterprise engineering; that is, the design of processes comprised of activities, information, and resources. But it has not adequately addressed the ``people'' side of the enterprise. This paper extends previous work by focusing on the incorporation of the *socio-technical* view of enterprise design. Sociotechnical design is the integration of organizations from the perspectives of people, organization structure, processes, technology, and culture. Issues of roles, authority, empowerment, and incentives therefore play a pivotal role in the optimal operation of business processes.

Enterprise Design

Our research explores two foundational aspects of enterprise engineering. The first component is the development of integrated models for enterprise design. Previous approaches to enterprise design lack an adequate specification of the semantics of the terminology of the enterprise models, which leads to inconsistent interpretations and uses of knowledge. Analysis is hindered because models tend to be unique to the enterprise, and are rarely reused. Within enterprise operations, obstacles to interoperability arise from the fact that the legacy systems that support the functions in many enterprises, were created independently, and do not share the same semantics for the terminology of their enterprise models.

A rigorous foundation for enterprise design, analysis, and operations therefore requires a formal specification of the semantics of enterprise models through the use of ontologies.
An ontology is a formal description of entities and their properties, relationships, constraints, behaviours [Fox and Gruninger 98]. It provides a common terminology that captures key distinctions and is generic across many domains. We have been developing a formal approach to enterprise modelling where the semantics of the terminology of the ontologies are defined using first order logic.

By representing the enterprise as a logical theory using the ontologies, operational and management questions about the enterprise can be considered as deductive queries using first-order logic [Gruninger 97]. All of the relationships among the different constraints within the enterprise are therefore made explicit. In addition, this framework provides a characterization of classes of enterprises by sets of assumptions over their processes, goals, and organization constraints.

In this paper we will discuss the extension of existing ontologies to include the organisational, behavioural and motivational aspects of the enterprise. These will address ``softer'' issues such as goals and objectives, policies and strategy, responsibilities and accountability, empowerment, incentives, culture, and change management. In particular, we will focus on the ontologies required to characterize a wide class of business processes arising from collaboration with IBM Canada.

Enterprise Analysis

The second aspect of enterprise engineering is the development of design theories that enable the exploration of alternative enterprise models spanning organization structure and behaviour.

The knowledge required to characterize the best enterprise along each design perspective, such as cost, quality, or time, is formalized using different sets of axioms in the microtheories of enterprise design expertise. By formalize, we mean the identification, formal representation and computer implementation of the concepts, methods and heuristics which comprise a particular perspective. This not only enables a precise formulation of the intuitions implicit in practice, but it is also a step towards automating the execution of certain tasks involved in enterprise engineering.

To formalize the intuition of design perspectives, we specify microtheories, which are sets of axioms in first-order logic which capture the knowledge required to characterize the best enterprise along each perspective (such as cost, quality, or time). In [Gruninger and Fox 94c] we introduced the notion of advisors as encapsulations of the microtheories required to reason about alternative enterprise designs. Each advisor is used to evaluate, guide, and complete possible enterprise designs (which are themselves expressed as first-order theories using the ontologies).

The advisors and their microtheories are the formalization of the enterprise design heuristics, and they constitute the theories which must be tested and validated for different classes of enterprises. Using the integrated ontologies allows the flexible configuration of
enterprise models and operating scenarios for problems. This supports a ``plug-and-play" approach to the incorporation and change of constraints in a problem specification. Hypotheses for enterprise design heuristics are expressed as queries that can be deduced from the axioms of the ontologies and microtheories. Testing and validation of design hypotheses is therefore done using the deductive queries from the design microtheories in conjunction with the possible enterprise models and operating scenarios of a problem.

Process Integration - A Case Study

Our previous work has focused on the development of generic micro-theories representing different design perspectives such as Activity-Based Costing, ISO9000 Quality, and Time-based Competition. In this paper we will explore the development of microtheories for socio-technical design. The paper will focus on the Agility and Process Integration Advisors, which were applied to selected business processes within IBM Canada [Atefi 97]. These advisors were the first attempt at the formalization off "soft" business process reengineering (BPR) heuristics and expertise. The categories of BPR expertise formalized within the Process Integration Advisor included

- integration of information flows within the enterprise
- allocation of personnel to case management
- modification of the assignments of personnel and processes
- violated policies

The Process Integration Advisor was applied to IBM's World wide Customer Opportunity Management process. It succssfuly identified and recommended solutions to a number of problems in the process design, many of which were not previously known to IBM.

References

[Atefi 97] Atefi, K., (1997) "Formal Models of Business Process Engineering for Design and Design Validation", Enterprise Integration Laboratory Technical Report, University of Toronto.

[Fox and Gruninger 98] Fox, M.S. and Gruninger, M. (1998), Enterprise Modelling, *AI Magazine*, volume 19, no.3, pp 109-121. AAAI Press

[Gruninger 97] Gruninger, M. (1997), Ontologies for Enterprise Engineering, *Enterprise Engineering and Integration: Building International Consensus*, Springer-Verlag.

[Gruninger and Fox 94] Gruninger, M., and Fox, M.S., (1994), An Advisor-based Architecture for Enterprise Engineering, *Workshop on Artificial Intelligence in Business Process Reengineering*, AAAI-94, Seattle.

VDT Computational Emulation Models of Organizations: State of the Art and the Practice

Raymond E. Levitt,
Professor of Civil Engineering
Associate Director, Center for Integrated Facility Engineering
Stanford University

The Virtual Design Team (VDT) research was initiated in the late 1980s with the goal of developing new micro-organization theory and embedding it in software tools that could be used to design organizations in the same way that engineers design bridges, semiconductors or airplanes—by modeling, analyzing and evaluating multiple virtual prototypes of the system to be designed in a computer.

We recognized from the outset that this was a significant challenge. Micro-theory and analysis tools that embed the theory for designing bridges and airplanes rest on well-understood principles of physics, and involve continuous variables with interval and ratio metrics, and materials whose properties are relatively easy to control and calibrate. These analysis problems yielded to mathematical modeling using systems of differential equations, although solving the large systems of equations for a complex building or airplane could quickly become intractable. Subsequently, applied mathematics techniques for approximating differential equations and plentiful, cheap computing power allowed the models to be computerized so that more complex systems could be analyzed and designed. The approach used to develop this engineering science and technology was to embed well-understood micro-theory into the models and attempt to capture the interactions between elemental parts of a model through constraints (such as constraints that maintain consistency between the deflected positions of adjacent element edges in a finite element model). The result was increasingly accurate predictions of both micro and macro-behavior of many kinds of engineered systems. For many kinds of bridges, deflections under a variety of loading conditions can now be predicted to finer tolerances than the bridge can be constructed to!

Organization theories to date, in contrast, are characterized by nominal and ordinal variables, and poor reproducability for measuring even these variables. Verbal theories incorporating nominal and ordinal variables create a significant degree of linguistic ambiguity, so that experimental results can not be reliably replicated and different theories are difficult to reconcile or disprove. In the late 1980s, our research group concluded that attempts to model organizations computationally could benefit greatly from the use of non-numerical or "symbolic" representation and reasoning techniques emerging from computer science research on artificial intelligence. Early experiments convinced us and other researchers (e.g,. Masuch and Lapotin 89) that this was a fruitful modeling approach.

In scoping out the kinds of organizations that we would model, we picked project teams performing routine design or product development work. For this class of organizations, all work is knowledge work so that an information processing abstraction (Galbraith 77) could potentially be used. Galbraith's "information demand, capacity and throughput" model can be viewed as an analog to Newton's Laws of Motion in Physics, a simple and useful first order approximation. Second, goals and means are both clear and relatively uncontested for these organizations, so that we could finesse many of the most difficult "organizational chemistry" modeling problems inherent in the kinds of organizations that sociologists have usually studied—mental health, educational and governmental organizations. By operationalizing and extending Galbraith's information processing abstraction in the Virtual Design Team (VDT) computational model, and focusing on the easy corner of the space of all organizations, we developed several versions of VDT (Cohen, Christiansen, Thomsen,) and validated their representation, reasoning and usefulness using the trajectory shown in Figure 1 (Levitt et al, 94; Kunz et al., 98; Thomsen et al, 98).

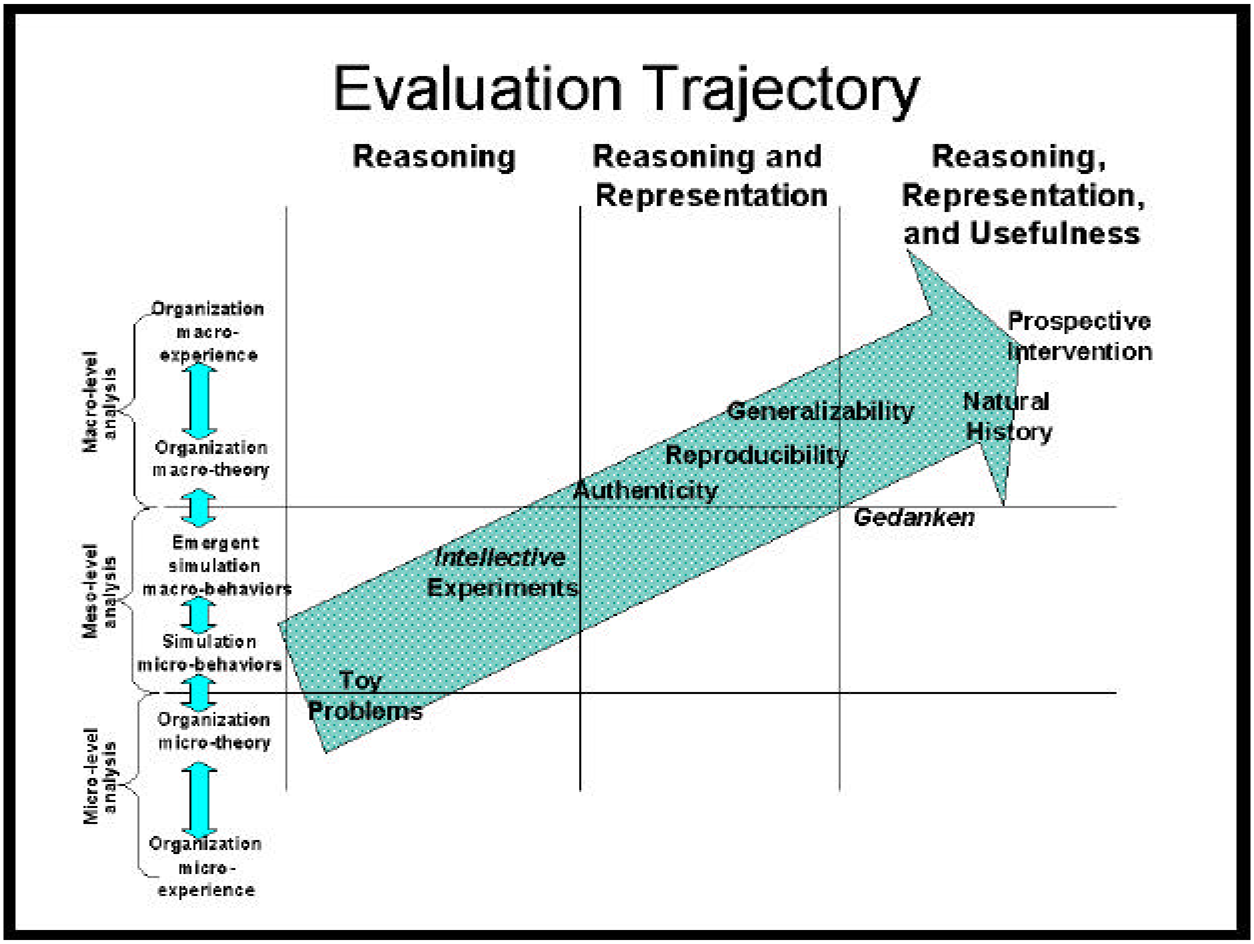

Figure 1. Evaluation Trajectory for Computational Emulation Models

Moving through the steps of validation, we developed sufficient confidence in the predictions of our theory and tools that managers in several companies are now redesigning their project work processes and organizations prospectively, based on the predictions of VitéProject™, a commercial implementation of the VDT model. Our VDT theory and analysis tools for project organizations have thus enabled true "organizational engineering" of project teams with congruent goals and routine—albeit complex and fast-paced—design or product development work.

Our intention was always to start with the "organizational information flow physics" and then progressively add elements of "organizational chemistry" to the modeling framework. This would allow us to move out of the "easy corner" of the organizational space and address a wider range of tasks and organizations. It is useful at this point to position our completed and ongoing versions of VDT in the space of organizations and modeling issues (See Figure 2).

The Cohen/Christiansen VDT-2 framework has been fully validated through all of the steps shown in Figure 1. It is applicable to routine work, in which: (1) All activities in the project can be predefined; (2) the organization is static, and all activities are pre-assigned to actors in the static organization; (3) exceptions to activities result in extra work volume for the predefined activities and are carried out by the pre-assigned actors; and (4) actors are assumed to have congruent goals. These conditions are satisfied for many kinds of design and product development work. VDT-2 has been commercialized as VitéProject™, and is in use by companies in a variety of industries for designing their project work processes and organizations.

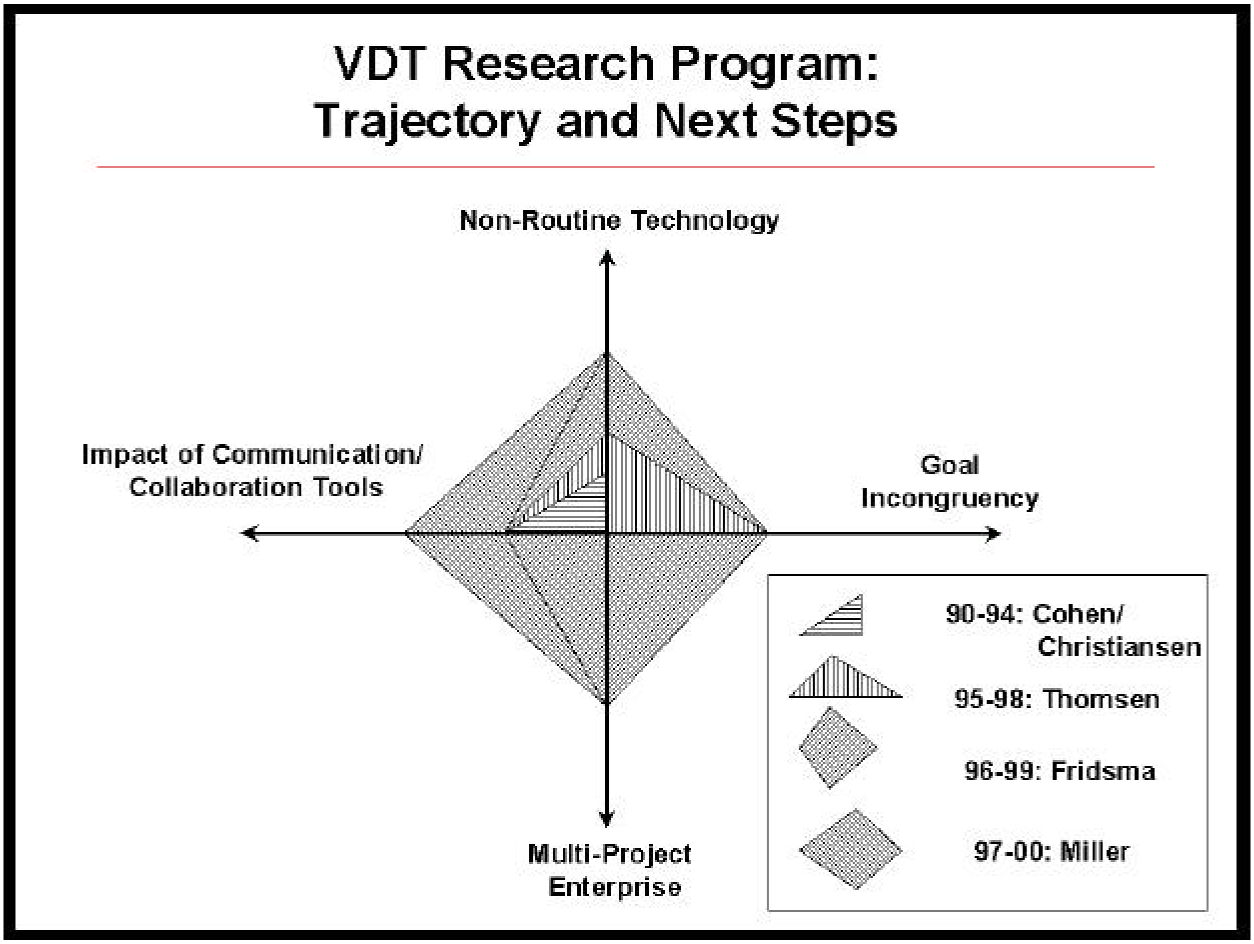

Figure 2. VDT: State of the Art and of the Practice

Thomsen's VDT-3 extends the range of work processes that can be modeled, to encompass less routine design or product development work, in which tasks are still predefined, but there can be flexibility in how they are executed. Actors can have incongruent goals causing them to disagree about how best to execute activities in the project plan. Goal incongruency levels between pairs of actors effect vertical and horizontal communication patterns. VDT-3 has been validated through "gedanken" experiments—thought experiments, in which the model's predictions are compared to managers' predictions of results. It has not yet been prospectively validated—in which its prospective predictions are tested against subsequent real project performance data.

Fridsma's VDT-4 is an attempt to move aggressively beyond the previous VDT limits on work process routineness and static organizational structure. He is attempting to model the highly non-routine work involved in health care delivery for bone marrow transplants and similar complex, multi-specialty medical protocols. For this purpose, we have to relax the constraint that all activities are known in advance. Diagnosis activities indicate needed repair activities, and any side effects must be diagnosed and treated contingently. This requires significant extensions to the VDT-3 framework. Fridsma (98) has extended the information processing micro-theory in VDT-3 to include a variety of more complex exceptions that can cause activities to be added, resequenced, deleted or reassigned, and actors to be dynamically added to the organization and assigned activities as needed. Referring to Figure 1, some parts of this extended framework have been implemented and internally validated on "toy problems". When the framework is complete, we will commence "intellective experiments," in which the model's predictions are tested against extant macro-theory, and so on, along the validation trajectory.

VDT-2,3 and 4 implement a simple model of communication tools that transmit information between actors' in-trays. Miller's VDT-5 attempts to extend the tool model to include the interactions between task, tool and organization for more complex collaboration technologies that effect not only communication of information between actors, but also the speed and accuracy of information processing by actors working together. Miller is developing the micro-theory to describe how attributes of tools, tasks, actors and organization structure interact, based on extant micro-theory in the CSCW and HCI literatures. When he has selected the "canonical" microbehaviors that he will implement as extensions to VDT-3 or VDT-4, he will begin testing their implementation on toy problems, and proceed with intellective simulations and other kinds of validation as the model develops fully and is refined.

A longer range goal of our work not shown on Figure 2 is to begin modeling even more flexible organizations that can be viewed as dynamically shifting "communities of practice," in which actors can communicate with anyone they choose inside or outside the organization their local "organization". Software development teams and some consulting organizations are already beginning to approximate this organizational form. Theories based on concepts such as Public Goods or reciprocity can be used to describe how these links form and persist or dissolve in cyberspace (Monge and Contractor, ??). This extension to our framework remains a gleam in our eye at this point.

This summary of the state of the art and of the practice of our VDT emulation models should give the reader a feel for how powerful and flexible the information processing view of organizations appears to be. We have found it provides an excellent first order theory to employ as the basis for building models that can predict the flow of knowledge work through organizations, and help managers proactively to diagnose and address information bottlenecks, delays and quality problems arising from failed communication attempts.

REFERENCES

Burton R and Obel B. Strategic Organizational Diagnosis and Design: Developing Theory for Application. Boston: Kluwer Academic Publishers, 1998.

Christiansen, T. R. (1993), *Modeling the Efficiency and Effectiveness of Coordination in Engineering Design Teams.* Ph.D. Dissertation, Department of Civil Engineering, Stanford University. Published as Det Norske Veritas Research Report No. 93-2063, Oslo, Norway.

Cohen G.P. The Virtual Design Team: An Object-Oriented Model of Information Sharing in Project Teams [Ph.D.]. Stanford: Stanford University, 1992.

Fridsma DB and Thomsen J. Modeling Medical Protocols for Organizational Simulation: An Information-Processing Approach. Computational and Mathematical Organization Theory 1998;4(1):71-95.

Galbraith, J. R. (1977), *Organization Design*. Reading, MA: Addison-Wesley.

Kunz, J. C., T. R. Christiansen, G. P. Cohen, Y. Jin, and R. E. Levitt (1998), "The Virtual Design Team: A Computational Simulation Model of Project Organizations," *Communications of the Association for Computing Machinery*, 41:11, pp. 84-92.

Levitt, R. E., G. P. Cohen, J. C. Kunz, C. I. Nass, T. Christiansen, and Y. Jin (1994), "The Virtual Design Team: Simulating How Organization Structures and Information Processing Tools Affect Team Performance," in K. Carley and M. Prietula (Eds.) *Computational Organizational Theory*, Hillsdale, NJ: Lawerence Erlbaum Associates.

Masuch M and LaPotin PJ. Beyond garbage cans: An AI model of organizational choice. Administrative Science Quarterly 1989;34:38-67.

Monge, P. R., & Contractor, N. S. (1999). Emergent communication networks. In L. Putnam & F. Jablin (Eds.) *Handbook of organizational communication (second edition).* Newbury Park, CA: Sage.

Thomsen, J., M. A. Fischer and R. E. Levitt (1998), *The Virtual Team Alliance (VTA): An Extended Theory of Coordination in Concurrent Product Development Projects*. CIFE Working Paper #44, Stanford University.

ORGANIZATIONAL LEARNING

Organizational Structure for Multi-agent Learning of Non-decomposable Problems

James D Thomas
Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
jthomas+@cs.cmu.edu

and

Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
katia+@ri.cmu.edu

This paper uses computational modelling to explore the ability of differing organizational structures to learn difficult tasks in an information processing context. Traditional models of organizational learning have focussed on learning relatively simple, fully decomposable tasks whose subtasks can be perfectly learned by individual agents in isolation. Here, we emphasize the role of tasks that are only partially decomposable, tasks that have sparse connections across recognizable subtasks. Most interesting tasks faced both by human organizations and artificial intelligence fall into this category.

Since we are interested in organizations as information processors, we model them as functions. Each agent sees a vector of inputs (from the outside world or other agents in the organization), computes a function over that vector, and passes on an output. The organization attempts to learn some function by examining sets of inputs and appropriate target outputs, and adjusting its internal structure until the outputs it produces match those of the target. If the target function is fully decomposable, an agent or set of agents can be assigned to specific subtasks and can learn them in isolation. Only when there are connections across these subtasks that makes this sort of learning difficult; no subtask can be learned without reference to potentially every other subtask.

In order to model our intuition of how such partially decomposable tasks can be learned, we take as our model of target functions second order polynomials. We hope that this class of functions is simple enough to be tractable and provide intuitive understanding, while retaining the essential character of partially decomposable tasks. The intuition goes as follows: second order polynomial functions can be represented as matrices. Fully decomposabe tasks can be thought of as block-diagonal. Each block represents a specific subtask. Provided that the blocks aren't too large for individual agents, each block can be handled by a single agent or a small group of agents. However, adding arbitrary off-diagonal elements creates connections across agent input domains. Given this model, we examine the abilities of organizational structures to compute such functions and conclude that standard conceptions of teams and hierarchies literally cannot learn such functions. The only organizational structure that allows for the computation of arbitrary is one with agents with input domains that overlap, handling the connections across subtasks. However, without knowing the full functional form in advance, coverage of every possible connection (requiring huge numbers of agents) is necessary to guarantee perfect learnability.

This poses a problem: we are reduced either cheap organizations unable to learn tasks completely, or hugely inefficient organizations that can. The solution we propose is to think about organizational learning in terms of the structure of the organization itelf in addition to learning at the agent level. In addition to agents simply learning the values of coefficients that lie within their input domains, a mechanism is needed which identifies which off-diagonal elements (connections across subtasks) are non-zero and assigns agents to expand their input domains accordingly.

This forces us to reconceptualize the function played by team structures vs hierarcical structures in organization learning. We introduce an algorithm based on the perceptron learning algorithm. It examines correlations between the outputs of other agents, and the organizational error on the target task. Correlation with large error usually indicates a connection across sub-tasks. This is where the difference between teams and hierarchies comes into play. In the team structure, agents directly look for correlations between the outputs of other agents with their own inputs. In the hierarchical structure, a supervisor agent looks for correlations between low-level agent outputs, and assigns low-level agents to explore promising connections. This sets up a contrast: the team approach has access to better information than the hierarchy, but at a higher computational cost.

We set up computational simulations in order to test these intuitions. We implemented our algorithm and ran both organizational structures over a wide variety of conditions. As a measurement of performance, we used the number of training epochs needed to reach a
certain level of error (set low enough that the organizations needed to identify each cross-domain dependency to reach it). Our core results bear out our basic intuition. Measured in training epochs, the team approach learns faster the hierarchy approach. However, the hierarchy approach needs fewer computations per training epoch. As a result when performance per training epoch is similar, the hierarchy approach is preferrable, but often the team approach is so much faster that it proves superior in total computation time. It is this dynamic that drives all of the results.

As a baseline, we ran simulations of an organization consisting of a large, single agent. Somewhat surprisingly, the single-agent approach did not learn the tasks significantly faster than either the hierarchy or team approaches (we speculate this is because the single agent spent so much time adjusting irrelevant coefficients down to zero), and the total computation required was enormously larger.

We examined a few spefcific issues. The first was a comparison of the two approaches' abilities to learn tasks of varying difficulties. As a proxy for task difficulty, we used the number of off-diagonal elements in the target function. With only one off-diagonal element, both the hierarchy and team approaches take about the same amount of training epochs to learn the function. As the number of off-diagonal elements increases, the training epochs required increases linearly for both organizational structures; however, the number of epochs required for the hierarchy approach rises significantly faster than the team approach. Measured in raw computational costs, this means that the hierarchy approach is cheaper on simple tasks, and the team approach wins out as the tasks grow more complex.

Next we examined the abilities of these organizational structures to handle dynamic tasks. To address this, we ran simulations allowing the organizations to learn simple tasks, and then introduced new off-diagonal elements at regular intervals to represent a changing target function. Here, the team approach was able to adapt quickly to changing tasks, keeping pace as quickly as new elements were introduced. For slow rates of change, the hierarchy approach performed nearly identically (in terms of training epochs); however, at a certain rate of change the hierarchy approach 'overloaded' and became significantly slower.

A further crucial issue is one of scaling; how well these algorithms performed as the task size (and hence size of the organization) increases. We find that as task size increases, the time required (in training epochs) increases very slowly, suggesting that scalability to large tasks is possible. Of course, given the increased number and size of agents, the total number of computations required does increase significantly.

Finally, we discuss the idea of cost. From a computational perspective, the concept of cost is simple: number of numerical operations. It is not clear that this is the correct way to measure costs for models of human agents. We discuss the idea of quadratic costs, and how changing assumptions about costs affects the relative performance of the two core organizational structures.

Taken together these results suggest the following. When tasks are simple, hierarchical structures are the most cost efficient way to learn; as task complexity increases, the team approach has the advantage. In dynamic situations, either method works well, so long as change is slow; however, the team approach is clearly superior when change is fast. Task size (and by extention, the number of agents and their sizes) appears to not be an important variable in these considerations. The exact tradeoffs involved depend on how one measures cost; for computational purposes a straightforward counting of operations is the clear metric, however, models of costs to human agents introduce added complexities.

Adaptive search revisited:
Replicating and testing an influential model of organizational learning

Anjali Sastry
University of Michigan Business School
701 Tappan Street
Ann Arbor, MI USA 48109-1234
phone: 313-763-1591
fax: 313-764-3146
email: masastry@umich.edu

Tom Fiddaman
Ventana Systems
8105 SE Nelson Road
Olalla, WA 98359
phone: (253) 851-0124
fax: (253) 851-0125
email: tom@vensim.com

In the study of organizational learning and adaptation, computer models of the processes involved in searching for new technologies and implementing them constitute an important and very influential contribution. Over several decades, March and his colleagues have drawn on computer simulations of organizational decision making to explain patterns of change in the evolution of organizations and to predict a number of counterintuitive phenomena, such as competency traps, slack search, and what March has labeled the "technology of foolishness." The findings from this program of research have shaped how we look at organizational change, technological innovation, and judgment and decision-making in organizations.

In this paper, we present our analysis of an early and influential computer simulation in the organizational learning research program: Levinthal and March's 1981 model of adaptive organizational search. We survey the evidence of its impact by looking at how the ideas presented in the paper are used by others (some 55 primary citations in refereed papers) as well as by March and his colleagues in subsequent work (such as his 1988 survey of Organizational Learning with Levitt, which in turn is cited 246 times). With this survey of key notions setting the context, we begin our analysis of the original model by replicating it using a system dynamics approach. The replication process generates a number of important questions about model formulation and interpretation. Simulation experiments allow us to link model behaviors to their structural underpinnings at the equation level. This exploration yields a number of intriguing insights and puzzles.

Although the formal model represents several features of behavioral decision making others have found to be important, such as exponential smoothing of expectations, we discover that in other respects the model differs significantly from the premises of behavioral decision theory. For instance, we find that the Levinthal-March model does not represent explicitly the lags required for collecting data, interpreting it, making decisions, and implementing them. Such lags may be missing because the modeling approach involves discrete-time simulation rather than the continuous-time approximation characteristic of system dynamics.

The discrete time representation also creates significant oscillatory tendencies in the model that do not have a clear behavioral basis. For instance, early results depict an oscillation in the early time periods, across a wide range of conditions, that is difficult to interpret. We also find an interaction between the frequency of environmental change and time lags in the organization. In addition, model formulations pose important questions. The representation of innovation draws, for instance, involves a heavy right tail that can generate severe problems in a number of simulations. Some of our findings suggest questions about both the original model and our theories of organizational change. Our observation that learning *requires* noise, for instance, contrasts to the arguments of March and his colleagues that learning happens *despite* noise.

In building the system dynamics simulation we translate the model to continuous time and reformulate several key components. We discus these innovations and show how they incorporate Levinthal and March's original assumptions as well as recent findings about organizational change, decision making, and technological change. Some of the behaviors found by March and his colleagues now appear in a narrower set of cases than the previous research shows, suggesting that several of the earlier predictions should be qualified.

We conclude with a discussion of several implications for those who both study and manage organizational change and reflect on the value of revisiting existing models and examining simulation results through the lenses of multiple modeling methodologies. The research also provides dramatic evidence of the value of a modeling approach capable of representing bounded rationality and behavioral decision-making.

STRUCTURAL AND CULTURAL LEARNING IN AND AMONG ORGANIZATIONS

Kathleen M. Carley
Carnegie Mellon University
Pittsburgh PA 15213
kathleen.carley@cs.cmu.edu
http://sds.hss.cmu.edu/faculty/carley/carley.htm

Conventional wisdom often tells us that when it comes to "getting things done" it is "who" you know rather than "what" you know that matters. Contrarily, this same wisdom advises that "knowledge is power". The conflict between these perspectives is reflected in research that separates these two factors, focusing on how structure influences behavior devoid of personal attributes, or on attributes excluding the relationships that contextualize the individual and their attributes. This research has shaped our thinking about behaviors in organizations, often focusing on either the relationship or the attribute unable to reconcile the interrelationship between these two factors. Today, however, a new view of organizations as inherently complex, computational and adaptive systems is emerging. This new perspective urges a reconsideration of the roles of relations and attributes. These two mechanisms are symbiotic, simultaneously impacting the behavior in organizations. Organizations are composed of intelligent adaptive agents constrained and enabled by their positions in networks linking agents and knowledge. Consequently, organizations are themselves synthetic agents in which knowledge and learning reside in the minds of the participant agents and in the connections among them.

This paper presents a knowledge level approach to organizational learning. This approach defines learning, culture, structure and behavior at the individual and organizational level in terms of both "who one knows" and "what one knows". Relationships among individuals are important as they facilitate individual access to knowledge and serve as a form of organizational knowledge. Learning is conceptualized as the acquisition or loss of nodes and relationships, as with each node or relationship gained or lost knowledge is likewise acquired or forgotten. This is true whether the nodes are pieces of information or agents. This is true whether the relationships are among individuals, among information, or between individuals and information. Thus learning and memory exist at the individual and organizational level. When organizations, as synthetic agents learn important organizational behaviors emerge. Such behaviors reflect the emergent structuration of the organization's culture and structure through learning at the individual and structural level.

Herein, this view is explicated and given precise form by defining the primary constructs in terms of the meta-network linking people and knowledge and the processed for evoking change in this network. Essentially a knowledge level perspective is used to extend the traditional social network perspective on organizations to include both people and ideas. Social networks are seen as affecting a wide range of behaviors ranging from power to consensus to adaptability. According to the common formulation such networks are in terms of ties among personnel. However, networks are more ubiquitous than this conception implies. Networks of ties link not just people, but people, knowledge, resources, tasks etc. (Krackhardt and Carley, 1998). We explicate this idea, and describe a computational model that incorporates this perspective. This computational model is then used to examine the relation between structural stability and flexibility, performance, diffusion of new information, and consensus within organizations.

From a knowledge level perspective learning and cognition exists both within and among individuals. The former, being the familiar experiential learning and attributable to the individual gaining experience and receiving feedback and so over time becoming socialized to a certain way of doing business. In the aggregate this is cultural learning. The latter, the learning among individuals, is structural learning and is attributable to the structuration of ties and the stabilization of the networks within and among organizations. These ties act as a codification of process and so enable synthetic adaptation.

Using a computational model of organizational learning at both the cultural and structural level, CONSTRUCT-O, a series of virtual experiments are conducted to explore the impacts of complexity (in terms of population and cultural size and differentiation) and pro-activeness on organizational performance. Results indicate that while complexity in terms of size makes the organization take longer to reach its peak performance, the degree of pro-activeness in acquiring information actually has a curvilinear effect. Information seeking and information reception serve at times as opposing processes and at times as synergistic processes. This effects whether the organization reaches peak performance in diffusion, triadic stability, or consensus first. Which type of behavior peaks first may have dramatic consequences for the organization's culture and its perceived performance.

COMPETITION AND POPULATION LEVEL LEARNING: A SIMULATION APPROACH

Alan B. Eisner
Lubin School of Business
Pace University
Management Department
861 Bedford Road
Pleasantville, New York 10570
(914) 773-3517
Fax (914) 773-3920

Stephen J. Mezias
The Stern School of Business
New York University,
Management Education Center
Department of Management
44 W. 4th St., 7th Floor
New York, New York 10012
(212) 998-0229.
Fax (212) 998-4234

Abstract

This paper focuses on the effects of pooled risk and expense of firm experimentation on the propensity of firms to refine current technologies and develop new innovative technologies. Innovations or refinements to firm capabilities can enable firms to enhance and maintain their competitive advantages. To explore these effects, a simulation model based on the integration of the organizational learning perspective and the evolutionary approach to economic change is used.

We explore the effects of population level institutional practices that may affect competition and imitation among a group of firms producing substantially similar services or products. We explore the competitive behavior of firms under the conditions where population level institutional routines are present. We focus on competition because it may have ramifications for the strategies firms select to gain and maintain advantage (Scherer and Ross, 1990).

Population level institutional effects include the formation of industrial research consortia or foundations where member firms share both associated costs and benefits of these organizations. Some populations of firms develop these cooperative processes that become institutionalized within their industries. We explore the effects on competitiveness of those populations that adopt these practices.

Following Mezias and Eisner (1997) our approach to modeling populations of firms and the ability of these firms to adopt innovations, imitations, and refinements involves several assumptions. The competitive environments of firms are viewed from a perspective that is articulated well in Nelson and Winter's (1982:278) description of Schumpeter's propositions about the nature of competition in industries.

The recent literature of organizations has included discussions of industries with rapidly developing technologies, decreasing product and process cycles, declining profits, and fierce price competition. Following Mezias and Eisner (1997), our conceptualization of competition is sufficiently broad to capture the impacts on firm performance that may result from intense price competition, overproduction, and declining profits. Mezias and Eisner (1997) define competition as:

> "… a direct negative impact on the performance of the focal firm resulting from the presence of other firms that compete for the markets that the focal firm serves. More specifically, competition occurs whenever multiple firms offer substantially similar goods or services. It is assumed that the level of competition increases with each additional firm that enters the market. It is further assumed that competition persists because collusion or cooperation is not feasible."

Following Mezias and Eisner (1997), we take competition and imitation not as mutually exclusive, but as multifaceted dimensions. Nelson and Winter argued that "... market structure influences the rate at which transient quasi-rents are eroded away by imitators"(1982: 280). This suggests that the effects of competition captured by market structure are systematically distinct from the effects of imitation. Whether it is possible for one firm to imitate the processes of other firms is theoretically distinct from examinations of the relation between market structure and competition.

In attempting to understand the impact of population level cooperative routines on competition and imitation, we adopt a Schumpeterian perspective. Thus, developers of new technologies or processes play a key role in

moving an economy towards the production of more valuable goods (Nelson and Winter, 1982; Chou and Shy, 1993; Segerstrom, Anant, and Dinopoulos, 1990; Cheng and Dinopoulos, 1992). This is an appropriate choice because the analyses from this perspective have been concerned with the role of innovation in an economy. Competition and imitation are essential ingredients in the central role of innovation of an economy suggested by a Schumpeterian perspective.

This paper uses the organizational learning perspective as the basis for modeling firms (March and Olsen, 1976; Levinthal and March, 1981; Levitt and March, 1988; Lant and Mezias, 1990; 1992). We believe this is an befitting choice because the ideas of organizational learning and innovation have been entwined previously, both in conceptual work (e.g., Angle and Van de Ven, 1989; Brewer, 1980; Stata, 1989; Tushman and Nadler, 1986; Tushman and Nelson, 1990; Brown, 1991; Mezias and Glynn, 1993; Mezias and Eisner, 1997) and in empirical research (e.g., Henderson and Clark, 1990; Cohen and Levinthal 1990; 338; Sahal, 1981). A few papers specifically delve into modeling competitive conditions explicitly by having organizations pay penalties for expending resources on innovating in the absence of significant increase in returns (Mezias and Glynn, 1993; Mezias and Eisner, 1997). They address the question of imitability and model returns to innovations that are directly to organizations' abilities to enjoy economic benefit from those innovations.

The model of learning used in this study is an adaptive organizational learning model that view populations of organizations that are goal-oriented systems responding to experience (Cyert & March, 1992; Glynn, Lant, and Milliken, 1994; March & Simon 1993; March & Olsen, 1976; Milliken & Lant, 1991; Levinthal & March, 1981; Levitt & March, 1988; Lant & Mezias, 1990, 1992). Organizations in the population continue successful actions and search for alternatives in light of unsuccessful actions.

In short, an adaptive organizational learning model implies three major sets of routines that organizations follow. First, organizations have performance aspirations to which they compare their current organizational performance (Cyert & March, 1992; Herriot, Levinthal, & March, 1985; Lant, 1992). Organizations analyze their performance feedback and scan their environment to assess their level of goal realization. Second, organizations search for alternative strategies under conditions of failure, where the process of gathering and processing information about alternative behaviors is a relatively costly (Cyert & March, 1992; Nelson & Winter, 1982). Third, organizations change their strategies based on alternatives selected by their search processes (Mezias & Lant, 1994).

Using an organizational learning model, this paper demonstrates how firms' routine-based behavior patterns may impact the level of innovation in a population of firms. This research focuses on how firms that are experiential learning systems might respond to varying levels of population level cooperation and the effects on the levels of competition and imitation. Mezias and Eisner (1997) found that intense competition was necessary to spur organizations into making the innovations necessary to distinguish themselves from competitors among firms modeled as experiential learning systems. They found that competition was not an obstacle to firms' decisions to refine their current technology. They also found that when firms can easily imitate the technologies of competitors, the level of both innovation and refinement that organizations undertake was reduced.

At the population level, Nelson and Winter (1982: 279) argue that imitators gnaw away at the returns innovators gain from investments in new development. However, in a population where institutional cooperative processes may be at work, there may be a public good generated by firms imitating the shared resources and information of the industrial cooperative effects. A raising of the competitive technological floor or starting point may, in effect, raise the overall level of innovation in the population. In this scenario, many firms a population contribute to an industrial consortium and share the cost of risky research and development. While free-riders may still exist in the population, they are counterbalanced by the relatively low cost of undertaking seeking innovations. Thus, population level institutional cooperative processes may increase the amount of innovation in a population.

PROPOSITION 1: Population level institutional processes will decrease innovative search costs.
PROPOSITION 2: Population level institutional processes will increase the levels of innovation.
PROPOSITION 3: Population level institutional processes will increase the levels of imitation.

The formation of information sharing mechanisms, such as industry trade groups, trade consortia, and bench marking groups, can enhance organizations abilities to assimilate new organizational processes. Organizational members of these groups can access specific information about best practices throughout their industry and often have the opportunity to discuss alternative organizational processes within their groups. These information sharing mechanisms may increase organizational flexibility, ease the cost of change, ease the cost of locating refinements of current technology, or ease the cost of locating the best technologies (i.e., imitation).

PROPOSITION 4: Population level institutional processes will decrease refinement search costs.

PROPOSITION 5: Population level institutional processes will decrease imitation search costs.
PROPOSITION 6: Population level institutional processes will decrease the cost of organizational change.

It is also possible that industrial cooperative efforts may create a "red queen" effect (Barnett and Hansen, 1996; March, 1996). In this type of situation, organizations are forced to innovate as a necessary condition of remaining viable in a rapidly changing environment. Thus, even though the technologies are available from the population level mechanism at relatively low research and development cost, the larger burden for a population may be the cost of implementation of new organizational routines. In these situations where implementation is the greater expense in advancing the technological capabilities of the population, we expect to find that population level institutional cooperative processes will actually decrease the level of innovation in a population.

PROPOSITION 7: Population level institutional processes will decrease the levels of imitation.
PROPOSITION 8: Population level institutional processes will decrease the levels of innovation.

EVOLUTION AND INNOVATION

Organizational Changes and Lamarckian/Darwinian Biological Evolution

Benoit MOREL[*] , Rangaraj RAMANUJAM[#]
Carnegie Mellon University

Abstract:
The idea that organizations behave as evolving systems has been very influential in the literature on organizational change[1]. Tushman and Romanelli in a well-known paper, used the expression punctuated equilibria to characterize the dynamics of organizational changes[2].
Punctuated equilibria are precisely the words used by S.J. Gould[3] to characterize biological evolution: period of steady growth, punctuated by extinctions of more or less large amplitude. In the same way, according to Tushman and Romanelli, the life of organization is made of long periods of "equilibrium" where changes are small and tend to improve incrementally the ability of the organization to achieve its goals, punctuated by short periods of revolutionary changes or strategic reorientations for the organizations.
The choice of words by Tushman and Romanelli was apparently fortuitous. But one of its consequences has been that the Darwinian theory of evolution has become a major paradigm in the design of models of organizational changes.
In their classical work entitled "An Evolutionary Theory of Economic Change", Nelson and Winter[4] had already pointed out that Lamarckian evolution is more appropriate than the Darwinian theory for that kind of study. The Lamarckian theory of evolution preceded the Darwinian theory and is considered as invalid as a theory of biological evolution. It assumes that evolution is progressive: a mix of adaptation to the environment and its changes together with a drive toward incessant improvement where new structures are modifications of existing structures. It involves two kinds of events[5]:

---

[*] Department of Engineering and Public Policy
[#] Graduate School of Industrial Administration

[1] MT Hannan, J. Freeman: *Organizational Ecology*, (1989) Harvard University Press.

[2] M Tushman, E. Romanelli, *Organizational evolution. A metamorphosis model of convergence and reorientation*, in LL Cummings and B. Shaw (eds), Research in Organizational Behavior **7** (1985): 171-222

[3] Gould, S.J. and N. Elridge (1977), "Punctuated equilibrium: the tempo and mode of evolution reconsidered", *Paleobiology* **3**, 114. See also: Gould, S.J. (1989) *Wonderful Life*, Norton.

[4] R.R. Nelson, S.G. Winter: An evolutionary theory of economic change, Cambridge MA (1982), Harvard Univ. Press.

[5] S. J. Gould, *Ontogeny and Phylogeny*, Harvard university Press, 1977, p.86

• Progressive changes by "the force that tends to complicate organization", and
• Adaptations to specific environments, which materialize in the neck of giraffes for example. Imitation can be a driver of evolutionary change.

This is different from the Darwinian evolution changes are in principle random, and the improvements are the result of the selection of the fittest. Genetic algorithms is supposed to be the mathematical translation of that dynamics into an optimization algorithm. In fact there is a fundamental difference between the two. Genetic algorithm is an algorithm to maximize a pre-existing cost or fitness function, whereas Darwinian evolution is a random process. There is no explicit fitness function.

Of relevance here, is the fact that the dynamic of change associated with Darwinian evolution is based on mutation and cross-overs, and its effect is seen in the following generations. This does not corresponds to the process of organizational changes. Organizational are clearly more progressive and affect the organization itself. The paradigm of Lamarckian evolution is more consistent with what we know about organizations. First, organizational actors are boundedly rational in their decision making and this cannot be modeled as an optimization algorithm.

Second, decisions in organizations are directed towards discovering a course of action that satisfies a whole set of constraints rather than towards achieving a goal. It is this set that it most accurately described as the goal of action[6]. This suggests that the basic problem is to identify the set of constraints that drive the decisions of organizations and use it as rules of evolution. This is very different from optimizing a fitness function. Some of the factors that are known to trigger organizational changes include decline in performance, change of leadership, change of environment, technological changes, etc. Taken together these suggest that the dynamics of organizational changes do not, as implied by a Darwinian approach, take the form of a random search in a space of possibilities that lets ecological interactions select the best solution. Rather, the dynamics of change are best represented by a model of evolution based on adaptation, imitation, fitting models as suggested by a Lamarckian view of the world.
There are several reasons why the Darwinian view has been very influential, among which there are the apparent relevance of the concept of punctuated equilibria[7] and the stochastic nature of the outcome of organizational changes.

But we show, through examples, that those features can also come out of models of organizational changes based on a Lamarckian perspective. More precisely, we present a computer simulation of a simple model of evolution that incorporates a Lamarckian view and demonstrate that it exhibits dynamics such as punctuated equilibria that are consistent with empirical findings reported by Tushman and Romanelli[8].

In this model, an organization is represented as a collection of routines. Evolution is measured in terms of changes that affect the performances of routines or the architecture of the organization by adding new connections between existing routines or by adding new or subtracting some routines. The algorithm is organized around the worst-performing. It consists to identify and change it. The performance of the routines directly connected to that routine is affected by the change. The effect is stochastic. This algorithm is an evolutionary algorithm inspired for the model of Bak and Sneppen[9], and is not an optimization algorithm.

---

[6] H. Simon, On the Concept of Organizational Goal, *Admin. Sci. Quaterl*. **9** (1964), p.1-20.

[7] Gersick, C.J.: Revolutionary Change theories: a Multilevel Exploration of the Punctuated equilibrium Paradigm, *Acad. Man. Journ*. **16**, (1991), pp.10-36.

[8] E. Romanelli, M Tushman, Organizational Transformation as Punctuated Equilibrium: an Empirical test, *Acad. of Managt. Journ*., **37** (1994), 1141-1166)

[9] Bak, P. and Sneppen, K. (1993), "Punctuated equilibrium and Criticality in a simple model of evolution", *Physical Review Letters*, **24** , 4083.

Although this is a very simple model, it displays rich dynamic behavior. We show that organizational change in this model exhibits punctuated equilibria. We use the findings from this study to reinterpret the data in punctuated equilibria reported by Tushman and Romanelli (1994). A Lamarckian model provides a conceptual basis for reconciling the seemingly irreconcilable positions of theories of strategic choice and theories of organizational evolution. It suggests that managerial determinism can exist without contradiction within a stochastic dynamics of change.

THE DIVISION OF INNOVATIVE LABOR EVOLVING

Paper submitted by Fabio Pammolli[1] and Massimo Riccaboni[2]

1. THEORETICAL MOTIVATION

The paper aims to move a step forward in the direction of establishing a connection between the structure and evolution of knowledge bases and the structure and evolution of organizational forms in innovative activities.

First, some stylized structural and dynamic properties of search spaces in scientific and technological systems are characterized. The emphasis is drawn on the nature of relevant uncertainties and the evolutionary processes that characterize the growth of knowledge and promote the division of innovative labor. Then, the paper analyzes the evolution of the structural properties of the network of collaborative agreements that enables and implements such division of labor under conditions of strong uncertainty.

Thanks to the focus on licensing agreements in R&D, the nature, the functioning, and the economic implications of a market for intermediate technological inputs are explored.

In this perspective, the network is considered as a type of organization of economic activities which, through *market mediated knowledge flows*, allows individual agents and firms to balance the need for a *variety* of a priori plausible trajectories to be explored and the *irreversibility* associated with the cost and degree of *specialization* required in search activities.

2. METHODS

The paper looks at the micro-dynamics of network formation and evolution in the pharmaceutical industry after the “molecular biology revolution”. It builds upon a comprehensive data set that covers more than 8.000 licensing agreements in pre-clinical R&D in the period 1978-1998. Firms are classified according dates of foundation and dates of entry within the network. Agreements are classified according underlying research techniques.

The paper is characterized by its explicit focus on the dynamics of the network.

First, the fundamental features of the growth of knowledge in molecular biology are mapped into the evolution of the industry and the aggregate dynamics of the structure of the network. Algebraic network analysis techniques of the degree of *structural stability* and *structural equivalence* are used in such a way to produce information on the dynamics of the network over time (see also Orsenigo, Pammolli, Riccaboni, et al., 1998).

Second, the *mathematical analysis of the bipartite graph* generated by R&D licensing agreements shows the coexistence of three fundamental roles within the network, that is a downstream sector composed by a group of established firms and an upstream sector composed by two groups of firms:

---

[1] Associate Professor of Economics and Management, Faculty of Economics “*Richard M. Goodwin*”, University of Siena, Piazza San Francesco, 7, Siena (Italy), 53100. E-mail: pammolli@unisi.it.

[2] PhD Student, *S. Anna School of Advanced Studies*, via Carducci 40, 56100 Pisa (Italy). E-mail: riccaboni@sssup1.sssup.it.

firms that seem to be specialized in technologies that work for specific purposes, and firms that seem to control general purpose technologies, applicable to N different applications.

The analysis of the variations of *the connectedness properties of the graph* over time unravels the fundamental mechanisms that structure the division of innovative labor and source-user relationships within the industry.

The mentioned analysis is carried out building upon a fundamental theorem on *canonical decompositions in bipartite graphs* by Dulmage and Mendelsohn.

3. RESULTS

Concerning network structure, it is found that, while the size of the network increases over time due to net flows of entry, its topological properties remain relatively stable, with strong first-mover advantages.

As far as age-dependent propensity to collaborate is concerned, the paper finds that the extent of inter-generation collaboration is much more significant than intra-generation collaboration.

The propensity of firms of a given generation to enter into collaboration with firms of a different generation increases with the distance between the two, while the total number of intra-generation collaborations decreases over time and, moreover, tends to decrease for most recent generations.

As far as the extent of role differentiation within the network is concerned, over time startups seem to come to establish areas of competitive advantage in a variety of generic technologies. In particular, one could conjecture that the large differentiation of the downstream market characteristic of pharmaceutical industry tends to promote the emergence of a new role within the R&D chain, that is companies that invest in general-purpose technologies/generic knowledge bases and interact with both established and co-specialized firms.

4. CONCLUSIONS

In the paper a unitary and coherent explanation for the above body of evidence is developed, coming to reveal the existence of a striking isomorphism between structural properties of the dynamics of knowledge and of the evolution of network structure.

The combination of variety, irreversibility of investment and co-specialization between search technologies and objects that characterize pharmaceutical R&D after the molecular biology revolution has strong economic and organizational implications. The huge variety of search hypotheses and investment trajectories generated by the progress of science produces two fundamental consequences. On one hand, learning new search technologies involves some fixed costs, which makes extremely costly to explore the whole tree of relevant hypotheses, or even large portions of it. The high level of co-specialization between objects and procedures of search, the nature of the dominant scientific paradigm and, finally, the presence of strong complementarities and feedbacks between theories and experimental practices in the process of search jointly induce high path dependency and conditions of strong irreversibility in the exploration activities carried out by firms in the industry. On the other hand, the required degree of structural differentiation is much higher than the one that hierarchy can handle, if it has to maintain internal consistency and identity. By its very nature, hierarchy tends to impose tight constraints to the process of generation of new competing hypotheses, and therefore tends to naturally limit the degree of cognitive diversity. To explore a search space which is characterized by large variety of plausible a-priori, it is necessary to use a mix of alternative coordination principles, including large integrated multi-technology corporations, different types of market institutions, and network of ties between firms and between firms and institutions. In sum, the very nature of search and learning in pharmaceutical research after the molecular biology revolution seems to require some kind of organizational form which realize a better compromise between flexibility and irreversibility than integration of R&D opportunities by large established firms.

In this perspective, the dynamics of the network of agreements is seen as an adaptive response to the structural features of the dynamics of search activities. To start with, one can observe that the hierarchical nature of the growth of knowledge helps to explain the relative stability of the network that is clearly demonstrated by our data. At the same time, the hierarchical nature of the growth of knowledge, combined with changing degrees of co-specialization between search technologies and research objects, can explain the expanding dynamics of the network of cooperative relationships. This can also explain one of the most relevant results of our analysis, that is that inter-generation collaborations are more frequent than intra-generation collaborations.

In this context, the network appears to be the emergent product of an adaptive process by which agents try to cope with the underlying tight exploitation-exploration trade-off. In fact, it can be considered as an organizational form that permits profitable opportunities to be exploited through the integration of downstream stages in the innovation chain, but also extremely uncertain prospects to be explored with adequate investment.

References

1) Anderson S. S. 1970, *Graph theory and finite combinatorics*, Markham Publishing Company, Chicago.
2) Atkin R. H., 1977, *Combinatorial connectivities in social systems*, Birkhauser, Basel.
3) Diersel R. 1997, *Graph theory*, Springer-Verlag, New York.
4) Dosi, G. 1982, "Technological paradigms and technological trajectories. A suggested interpretation of the determinants and directions of technical change", *Research Policy*, n. 11, pp. 147-162.
5) Graver J. E., Walkins M. E. 1977, *Combinatorics with emphasis on the theory of graphs*, Springer-Verlag, Berlin.
6) Harary, F. 1969, *Graph theory*, Addison Wesley, Reading, Ma.
7) Levinthal, D. 1996, Organizations and capabilities: The role of decompositions and units of selection, Wharton School, University of Pennsylvania, mimeo.
8) Liu C. L. 1968, *Introduction to combinatorial mathematics*, McGraw-Hill Book Compalny, New York.
9) Lovász L., Plummer M. D. 1986, *Matching theory*, Annals of Discrete Mathematics, 29, North Holland, Amsterdam.
10) March, J.G. 1991, "Exploration and exploitation in organizational learning", *Organization Science*, vol. 2, n. 1, pp. 71-87.
11) McCullagh, P., (1978), "A Logistic model for paired comparisons with ordered categorical data", *Biometrika* 64, pp. 449 - 453.
12) Mirsky L. 1971, *Transversal theory. An account of same aspects of combinatorial mathematics*, Academic Press, New York.
13) Nelson, R.R. 1961, "Uncertainty, learning and economics of parallel research and development efforts", *Review of Economics and Statistics*, vol. 43, pp. 351-364.
14) Orsenigo, L., F. Pammolli, M. Riccaboni, et al. "The dynamics of knowledge and the dynamics of an industry network", *Journal of Management and Governance*, vol. 1, n. 2, pp. 147-175.
15) Pattison P. 1993, *Algebraic models for social networks*, Cambridge University Press, Cambridge, Mass.
16) Powell, W.W., K. W. Doput, and L. Smith-Doerr. 1996, "Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology", *Administrative Science Quarterly*, vol. 41, pp. 116-145.
17) Simon, H.A. 1969, "The architecture of complexity", in *The sciences of the artificial*, MIT Press, Cambridge, Ma.
18) Slepian P. 1968, *Matematical foundations of network analysis*, Springer-Verlag, Berlin.
19) Wasserman, S., K. Faust. 1994, *Social network analysis: Methods and applications*, Cambridge University Press, Cambridge, UK.

20) Winter, S.G. 1984, “Schumpeterian competition in alternative technological regimes”, *Journal of Economic Behavior and Organization*, pp. 287-320.

The Evolution of Cooperation in Scientific Communities. A Socionic Approach

Peter Mutschke
Social Science Information Center
Lennéstr. 30
D-53113 Bonn
Germany
me@bonn.iz-soz.de

Abstract:

The question, how social agents develop the willingness to cooperate, is a key problem both for the distributed artificial intelligence (DAI) and for the sociology. In multi-agent systems (MAS), a special branch of DAI, up to yet only systems with few agents and simple social models were examined. To develop powerful algorithms in DAI, it is necessary to model more complex models in artificial systems, e.g. socially coordinated problem solving behavior as we can obtain from social reality. This is the aim of a new interdisciplinary research approach 'Socionic'.

In looking for specific problem areas, scientific communities are in almost ideal manner suitable to be an application domain for such research perspectives. Scientific communities indicate a set of characteristics, which play also a central role in multi-agent systems: They consist of a multiplicity of heterogeneous and constantly fluctuating agents. They are characterized by a mixture of team-like cooperation and hierarchically organized structures. Decision-making processes are organized decentralized and are characterized by a high degree of parallelism (aspect of the time as factor for the dynamics of the social system). Despite these characteristics it is nevertheless possible to form a society in shape of partly very large social networks which is able, to reach, for instance, a consensus concerning the specific contents of a research field or to establish a new innovative topic as a mainstream theme of the research field.

Those problems cannot be solved without cooperatively distributed problem solving. Cooperation is necessary, because no individual agent of the (science) system has sufficient knowledge, information and resources to solve such problems. Thus, there is obviously such a thing like an all-embracing problem solution ability (on the macro level), which is even not traced back to the exclusive problem solution authority of the individual actors, but arises as an emergent result of complex interaction processes.

Moreover, the pesented project can refer to scientometric investigations, in which socionic relevant hints for the modeling of social interaction in multi-agent systems were already given. One of the methods used in these studies is the system AKCESS which aims at the analysis of social interaction and science-structural relevance of persons in a scientific community on the basis of large bibliographic databases. The method examines two central features of cooperation between actors in a scientific community (co-authors etc.): one is their semantic proximity which is calculated by means of a modified equivalence index; furthermore, the roles of the actors in cooperation relationships and their meaning (e.g. project manager, publisher etc.) are considered as "hierarchical" relations between cooperation partners. On this basis transitive relations between all persons in a network and their semantics are derived, so that statements about the science-structural integration of scientists in a scientific community can be finally made regarding their degree of centrality and "hierarchical" relevance in researchers networks.

An another method is the cognitive mapping procedure LEXIMAPPE which is based on a co-word-analysis of terms of a research field. The method divides networks of keyword cooccurring in documents with the help of a clustering procedure into meaningful subunits called clusters or themes. On the basis of the categories centrality and density the thematic significance of a cluster can by expressed by its position in a two-dimensional space called cognitive map.
The surprising - and socionic interesting - is that over time cyclic process forms could be observed concerning the career of topics: After increasing centrality (which indicates increasing interest in the topic) a rising density of the topic follows. Finally the theme promotes to a "mainstream" subject of the research field (characterized by high density and high centrality). Leaving interest corresponds then with decreasing centrality but with still high density indicating that the topic moves (in tendency) into the area of "ivory tower" research. Thus, the clusters move within the map in predictable way, i.e. against clockwise direction.

The interesting now is that the interaction processes, which cause this phenomenon, are obviously dependent on sociostructural factors, i.e. the position of the actors within the communication context. In empirical tests using the tools LEXIMAPPE (topic development) and AKCESS (development of networking) a strong statistical connection could be actually proven between the centrality of persons and the centrality /density of themes. The inclination to select attractive subjects seems to rise with the membership and a central positioning in large networks.

Such processes seem to be - on the macro level - in tendency dependent on two central science-structural factors: sociostructural coherency of research fields (size and density of researcher networks) on the one hand side and the "closeness" of the communication context (research-field-internal consensus finding) on the other hand. By the example of the career of new innovative topics in sociostructural different research fields it could be proven that innovativeness (indicated by the consensus concerning the relevance of a new idea) arises rather not in the center, but in the periphery of a research field.

As a reason for this phenomenon the sociostructural caused topic orientation of the actors is seen: Highly integrated actors are more involved in research-field-internal consensus finding processes concerning the specific contents of the main topics of the research field. Therefore those scientists tend rather towards thematic conformity: We have an internal orientation of central actors. In contrary, scientists, who have peripheral positions in researchers networks, are more involved in external discussions, i.e. other research fields. This situation predestines them to the early transfer of new ideas. They function as bridges for the information flow. So, we can speak of an external orientation of marginal actors. Therefore, in contrary to the opinin leader model, a coherent system of mutual influence globally seems to hinder the acceptance of innovative ideas.

It is rather clear that such structure effects might affect also the co-operation behavior of the actors: The choice of subjects and cooperation partners seem to be strongly influenced by the social coherence of a research field and the degree of integration of the involved actors in cooperation networks. There is obviously such a thing like patterns of social interaction in scientific communities, which have even (at least on the macro level) a certain degree of "predictableness"and thus operationalizability for MAS. From a socionic perspective in particular the question arises whether the phenomenon of the cyclic topic careers corresponds with increasing social networking.

In order to be able to understand the evolution of co-operation in scientific communities network analysis on the basis of large bibliographic databases seems to be a suitable methodological basis. Network theory gives information about structure attributes and therefore the context and the borderlines in which agents could act. Simulation of network based social systems by artificial multi-agent systems, or multi-agent systems ruled by laws based on network theory will trigger both social science and DAI research.

COORDINATION AND COOPERATION

Application of Formal Theory of General Systems Theory to Organization Theory

Yasuhiko Takahara
Wataru Kimura
Dep. Industrial and Systems Eng.
Chiba Inst. Of Technology
Tsudanuma,Narashino,Chiba,Japan
takahara@cc.it-chiba.ac.jp

We have been working for development of a formal theory of an organization as an application of the mathematical general systems theory(MGST). This project tries to delve into deeper knowledges about an organization which can lay a foundation for the computational and simulational study of it. The starting point of our research is the systems approach to an organization and a cybernetics approach to it.

According to the systems approach to an organization is conventionally characterized by the following terms and concepts:Open socio-technical system;man-machine system;contrived system with a preassigned goal;hierarchical system;steady state or dynamical equilibrium;managerial system;in particular,feedback mechanism;coordination function;adaptive and maintenance mechanism;growth through internal elaboration. It is important to notice that in addition to the above characterization the systems approach asserts that a basic hierarchical structure of an organization consists of three levels,technical core(operational level),organizational level and institutional level. The technical core imports management resources and transforms them into an organizational product. This activity is a main function of an organization. The institutional level faces the environment of an organization and modifies the organization structure adaptively. The organizational level is responsible for the function of the middle management,coordination of the activities of the operational level. It is also important to notice that the institutional level is assumed an open system but the organizational level and the technical core are closed systems,shielded from the environment by the institutional level. The mangement resource input constitutes a part of the controllable variable of the operational level.

VSM(viable system model),which is a typical example of the cybernetic approach to an organization,details the management function of the systems approach. The task of the organizational level is departmentalized in VSM. The function of the organizational level is particularized as being responsible for management of “now and inside” of an organization status. On the other hand,the institutional level is for management of “future and outside” of an organization.

The hierarchy theory of MGST provides a layer model for a complex system, regulation/optimization layer, coordination layer and adaptive/self-organization layer. It is obvious that the functions of the three level model of the systems approach directly correspond to those of the multi-layer model. This fact suggests that if the multi-layer model ,which is a general model,is particularized in accordance to the specific characters of an organization,a formal theory of an organization can be derived from MGST.

Fig.1 shows the basic model for the mathematical organization theory we are developing. The model is obtained from the multi-layer model by specializing it in accordance to the characterization of an organization.

We can develop formal theories about many aspeects of an organization using the basic model. Typical theories are related with a formal theory of an organizational chart concerning with the hierarchy property or adaptation by variety matching of the intitutional level.

In the workshop we will discuss the coordination function of the organizational level to illustrate how we can develop a formal theory of an organization in our way.

We first formalize the basic model. An organization is assumed to have a goal,organizational goal. It is a function of the outcome of the operational level. It evaluates the organization performance. The operational level consists of goal seeking systems. They are interacted. The organizational level is given a goal which is called overall goal. The overall goal is an operationalized one of the organizational goal. At the same time the organizational level has a coordination variable and its own goal which eveluates the coordination activities. The organizational level is assumed to try to optimize the coordination evaluation goal.The coordination variable and the coordination evaluation goal are specified by the coordination principle.

In genreral it is not true that a successful coordination implies realization of the overall goal. If an organization has the property that a success coordination implies realization of the overall goal,the organization is called sound. In the workshop we will discuss a problem when a sound organization can be designed.

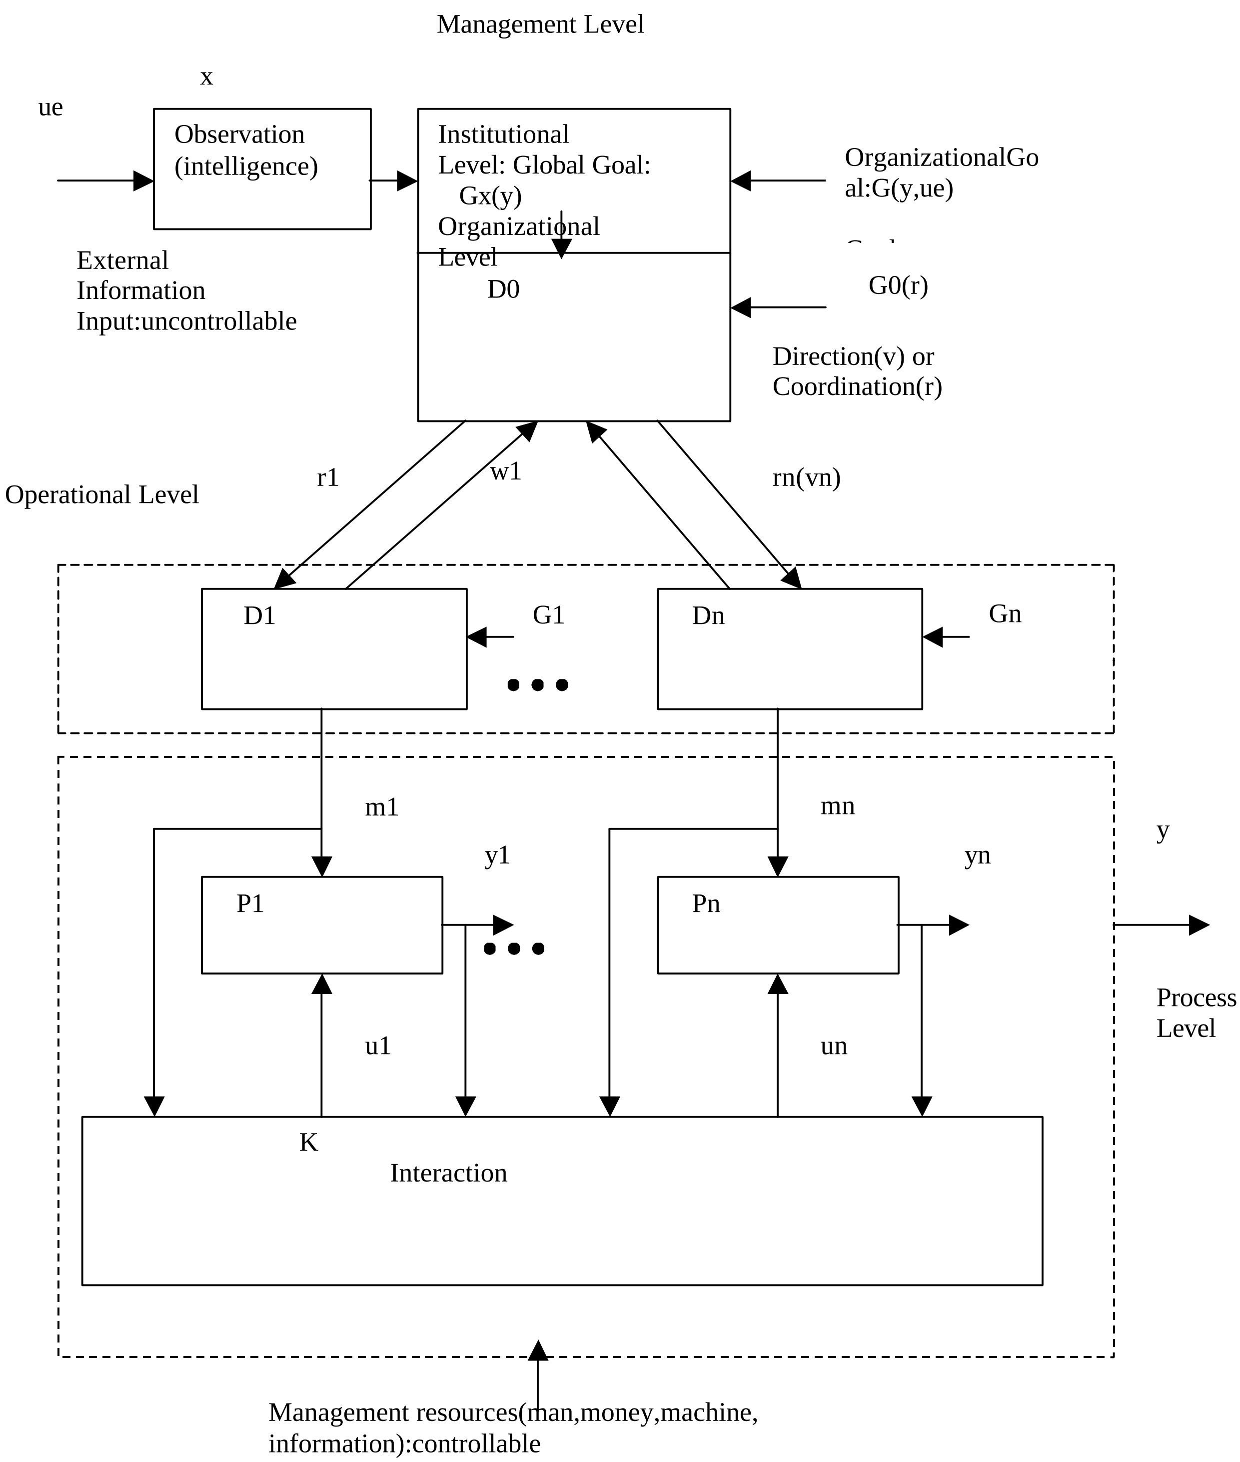

Fig.1 Basic Model of Organization

A Study of the Impact of Organizational Design on Organizational Learning and Performance

Ronald Vyhmeister and Aris Ouksel
University of Illinois at Chicago

Introduction

Since the beginning of the century, researchers have debated the best organizational structure [Taylor, 1911, Fayol, 1918; Weber, 1947; Mintzberg, 1985]. In recent years, downsizing (now often called "rightsizing") and reengineering have been posited as the "best" solutions to organizational design [Hammer and Champy, 1991].

Some researchers who believe that the elimination of middle management may be counterproductive [Cascio, 1995; Halley, 1995; Schenk, 1997] attribute the varying results of reorganization and downsizing to organizational design. Studies have found that reengineering (and its accompanying organizational redesign) often produces a less than desirable result [Davenport, 1996], motivating researchers to search for a better understanding of what constitutes good organizational design. Unfortunately, the results achieved through reengineering and downsizing are not the best [Davenport, 1996]. This leaves us with the question of what is the "best" way to structure an organization.

The Research Issues

The ability of organizations to adapt and learn is today considered one of the most important organizational characteristics. This ability is dependent on the existence of (a) organizational goals, (b) information which enables the organization to make decisions, and (c) feedback that enables individuals involved in the decision-making process to assess whether a decision is correct or not.

Many have attempted to model the organizational decision-making and learning process, but only one has clearly captured the various aspects of decision-making and learning. This model emerges from a long stream of research, beginning in 1990 [Carley, 1990; Carley 1991; Carley 1992; Mihavics and Ouksel, 1996; Ouksel, Mihavics and Carley, 1997]. This model (hereafter referred to as the OMC model) is a generalization and a formalization of a number of past studies [Carley, 1990, 1992, 1996; Lin and Carley, 1993; Carley and Lin, 1995; Ye and Carley, 1995; Mihavics and Ouksel, 1996]. Results from simulations of specific cases of the OMC model show that (a) organizations facing new tasks learn at different speeds and that (b) some structures learn better in the short term while others do poorly at first, but learn better in the long run [Mihavics and Ouksel, 1996; Lin and Carley, 1993; Carley and Lin, 1995].

While the results achieved to date are useful, the question arises as to their robustness. This question is especially important since past results are limited by (1) the limited organizational size used, and (2) the limited subset of organizational design parameters used in any one study, which ignores many of the possible interactions among the various design characteristics, and (3) every past study looks at results at a different time frame for decision-making.

Methodology

In order to address these three shortcomings and to get a more complete picture of the impact of organizational design on organizational learning and performance, we have used all the variables used in the various past studies. They are:

1. Number of Agents. The number of individuals at the bottom layer in the organizational structure.
2. Bits per agent. The number of elements of evidence that each agent processes for any given organizational decision.
3. Decision-making structure. The organizational structure used to evaluate learning ability (Expert teams, democratic teams, and hierarchies).
4. Evidence weighting. Typically weights have been assigned in one of three different ways: *uniform*, *non-uniform dispersed*, or *non-uniform clustered*. (See Mihavics and Ouksel, 1996).
5. Task Decomposition. Evidence is defined as seen by only one agent (*non-overlapping*), or by more than one (*overlapping*).

6. Incorrect information. The percentage of the evidence seen by an agent which is incorrect.
7. Missing information. The percentage of the evidence required by an agent which is not available.
8. Incorrect feedback. The percentage of decisions for which the feedback received by an agent is inaccurate.
9. Missing feedback. The percentage of decisions for which the agent receives no feedback.

For our study we model decisions that are two-valued (go/no-go). It is important to note that no study has used more than five parameters. In this study we look at all of these parameters simultaneously, looking to test the validity of past results, especially verifying whether their results remain valid in the presence of other parameters, or if any of the parameters change. In addition to using all of these past parameters, we introduce the concept of localized versus generalized feedback, where we distinguish between those organizations that provide the same feedback information to all agents from those that provide individualized feedback.

In order to test the robustness of past results, we insured that numerical parameters had multiple values, including high, medium, and low values [Carley, 1996]. In addition, we avoided using parameters which could lead to inconclusive information [see Mihavics, 1995, for a complete discussion of the problem]. Through the use of the three categorical variables: organizational structure (majority team, expert team, and hierarchy), weighting scheme (uniform, dispersed, and clustered), and feedback mechanism (generalized and localized) we divided all organizations into eighteen distinct groups.

Because the number of simulations needed to exhaustively test the model is excessively high, we used a sample in order to select representative cases for our study which satisfies two conditions: the sample size must (a) have a sufficient number of cases so we can test the robustness of past results of the model, and (b) it must be large enough so the results of the regression analysis of the asymptotic behavior are statistically powerful. We also insure that the number of decisions we simulate each organization are enough to ensure that the organizational performance approximates a steady state, while at the same time we should not simulate an inordinately large number of decisions, since organizations only look at similar decisions a finite number of times.

Results

Our results demonstrate a number of characteristics about the learning curve. First, we note that there are three important points during the learning process: (1) when learning begins, (2) when learning stabilizes, and (3) the maximum performance of the organization. Using statistical analysis we were able to determine which factors are statistically significant in determining each of these three points, as well as the magnitude and direction of their influence.

Our results verify the results of past studies using the same parameters. However, we find that many of the past results are only true for specific cases, or when certain combinations of design parameters are present. Others, such as the fact that hierarchies initially learn faster than other organizations, yet underperform them in the long run, are substantiated.

We find that giving accurate localized feedback as opposed to generalized feedback significantly impacts how quickly learning begins, but that feedback distortion is not a significant factor in determining the maximum potential performance of any organization. We also find that

One of the more interesting results is that the belief that hierarchies would be more resilient in the face of information distortion, are contradicted by our study. Is it possible that the increased speed of learning is dependent not on their resiliency, but merely on the fact that hierarchies effectively divide and conquer the problem?

More study is necessary, especially in the areas of data mining and complex decision functions, as studied in [Ouksel, Mihavics, and Carley, 1996] as well as field studies to test the results using actual organizations and people.

Bibliography

Carley, Kathleen. *CCOR Workshop, Carnegie-Mellon University*. Pittsburgh, PA: 1996.

Carley, Kathleen and Zhiang Lin. "A Theoretical Study of Organizational Performance under Information Distortion." *Management Science*. October, 1997, .

Carley, Kathleen M. "Computational and Mathematical Organization Theory: Perspective and Directions." *Computational and Mathematical Organization Theory*. 1, 1, 1995, 39-56.

Carley, Kathleen. "Coordinating for Success: Trading Information Redundancy for Task Simplicity." *Proceedings of the 23rd Annual Hawaii International Conference on System Sciences*. 1990.

Carley, Kathleen. "Organizational Learning and Personnel Turnover." *Organization Science.* 3, 1, February, 1992, 20-46.
Carley, Kathleen. "Designing Organizational Structures to Cope with Communication Breakdowns: A Simulation Model," *Industrial Crisis Quarterly*, 5, 19-57.
Carley, Kathleen M. and Zhiang Lin. "Organizational Designs Suited to High Performance Under Stress." *IEEE Transactions on Systems, Man, and Cybernetics.* 25, 2, February, 1995, 221-230.
Cascio, Wayne F. *Guide to Responsible Restructuring*. Washington, D.C.: U.S. Department of Labor, 1995.
Cascio, Wayne F., Clifford E. Young and James R. Morris. "Financial Consequences of Employment-Change Decisions in Major U.S. Corporations." *Academy of Management Journal.* 40, 5, October, 1997, 1175-1189.
Davenport, Thomas H. "The Fad that Forgot People." *Fast Company.* 1, 1, February, 1996, 70-74.
Fayol, Henri. *Administration Industrielle et Generale: prevoyance, organizstion, commandement, coordination, controle..* Paris: Dunod. 1918.
Halley, Alexis A. "Downsizing." *The Meridian International Institute on Governance, Leadership, Learning, and the Future.* Washington, D. C.: 1995.
Hammer, Michael and James Champy. *Reengineering the corporation: a manifesto for business revolution.* New York: HarperCollins, 1993 .
Lin, Zhiang and Kathleen M. Carley. "Proactive or Reactive: An Analysis of the Effect of Agent Style on Organizational Decision-making Performance," *International Journal of Intelligent systems in Accounting, Finance and Management.* 2:4, 1993, 271-287.
Lin, Zhiang and Kathleen M. Carley. "DYCORP: A Computational Framework for Examining Organizational Performance Under Dynamic Conditions." *Journal of Mathematical Sociology.* 20, 1995, 193-217.
Mayo, Elton. *The Human Problems of an Industrial Civilization.* New York: The Macmillan Company. 1933.
Mihavics, Ken. *A Model for the Study of the Effects of Organization Structure on Organizational Learning.* Univ. of Illinois at Chicago: 1995.
Mintzberg, Henry. *Structure in Fives: Designing Effective Organizations.* Englewood Cliffs, NJ: Prentice Hall. 1983.
Mihavics, Ken W. and Aris M. Ouksel. "Learning to Align Organizational Design and Data," *Journal of Computational and Mathematical Organization Theory*, 1:2, 1996.
Ouksel, Aris M., Ken W. Mihavics and Kathleen M. Carley. "A Mathematical Model of Organizational Performance." Technical report, University of Illinois at Chic: 1996.
Shenk, David. *Data Smog: Surviving the Information Glut.* London: Abacus. 1997.
Taylor, Fred Manville. *Principles of Economics.* Ann Arbor, MI: University of Michigan. 1911.
Weber, Max. *Max Weber, the Theory of Social and Economic Organization.* New York: Oxford University Press. 1947.

Virtual Teams: Its Dimensions and Coordination Implications

Sze Sze Wong
Fuqua School of Business
Duke University
Durham, NC 27708-0120
ssw2@mail.duke.edu

Richard M. Burton
Fuqua School of Business
Duke University
Durham, NC 27708-0120
rmb2@mail.duke.edu

(Key words: virtual team, coordination, team performance)

ABSTRACT

One evolving new form of organization is frequently described as the virtual organization. This emerging new form is viewed as the solution to managing global operations and external alliances, by linking distributed, and culturally and organizationally diverse participants using sophisticated networking technologies (Hedberg, Dahlgren, Hansson and Olve, 1997). Within the firm, it is conceptualized as a collection of virtual teams, which also bridge the boundaries of time, space, culture and organizational boundaries. Even though the concept of the virtual

organization has received wide currency in the popular business press (Davidow and Malone, 1992; King, 1994; McDonald, 1995; Barner, 1996), it has not been subjected to systematic analysis. This is unfortunate, as we need to develop a deeper understanding of the virtual organization, in order to propose strategies for managing this virtual structure.

The relative infancy of this topic demands that we take a narrower and less ambitious approach to studying virtual organizations. Instead of the virtual organization, we decided to focus our research attention on the virtual team. This is a reasonable initial anchor, as a virtual organization is frequently viewed as a collection of virtual teams (McDonald, 1995; Simons, 1995; Lipnack and Stamps, 1997). Like the virtual organization, the virtual team is characterized by members who are physically dispersed, and are both culturally and organizationally differentiated. Therefore, the problems of coordination and communication, which plague the larger virtual organization, will apply to the virtual team as well. Notwithstanding, we acknowledge that management of the virtual organization will entail problems of inter-team coordination and integration, which will be unfortunately missing if we merely studied a virtual team. However, we believed that greater understanding of the internal dynamics of a virtual team is prudent before advancing to the study of coordination processes at a broader level.

Consistent with existing conceptualizations of virtual teams (Clancy, 1994; McDonald, 1995; Lipnack and Stamps, 1997), we argue that a virtual team manifests the following characteristics: a set of (1) culturally and (2) organizationally differentiated members, who are grouped together (3) temporarily, are (4) physically dispersed, connected by (5) weak lateral ties, and engaged in performing (6) non-routine tasks. However, despite its so-called prevalence, few pure virtual teams exist today. Some teams consist of members who are geographically dispersed, but are culturally and organizationally homogeneous. Other teams may contain members who transcend cultural and organizational boundaries, but are physically co-located. Hence, whether a team is virtual maybe more in degree than in kind, an approach taken by DeSanctis, Staudenmayer and Wong (in press) in their conceptualization of virtual organizations. In this regard, we created three dimensions of team virtuality: (1) context; (2) team members' characteristics; and (3) relationships between team members.

Context. One oft-cited advantage of the virtual team is that its members can be assembled quickly to exploit emerging opportunities, and then dissolved when the opportunity expires (McDonald, 1995; Lipnack and Stamps, 1997). As such, the team members typically do not have a history of collaboration. In addition, the tasks undertaken by these virtual teams are usually non-routine, and completed under time-pressured environments. The members of the team are also often not physically located together (Clancy, 1994; Barner, 1996; Lipnack and Stamps, 1997). Rather, they are distributed throughout the world, connected only through an array of interactive technologies. Therefore, a team that is embedded in a virtual context is typically characterized by low team history, novel tasks, and physically distributed members.

Team Members' Characteristics. Virtual team members are also culturally and organizationally diverse. The knowledge and talents of these members, rich in their unique cultural and organizational perspectives, are pooled together to maximize the potential of the team to exploit market opportunities (Lipnack and Stamps, 1997). Hence, another important dimension of virtuality is the heterogeneity in the team members' cultural and organizational background.

Relationships between Members. The relationships between these diverse members are more often than not, weak and formal, rather than strong and informal. The lack of face-to-face interaction, the span across cultural and organizational boundaries, and the lack of a history of cooperation suggest that the relationships amongst virtual team members are weak. However, due to their physical dispersion and the nature of their work, they are often connected by lateral ties. In sum, members are likely to be linked to one another laterally, but tend to treat one another formally, due to both a lack of a prior relationship, and the cultural and organizational barriers.

Depending on where a team falls on the virtuality continuum, the coordination dynamics within the team would differ. For instance, a team that is only virtual on the context dimension (e.g. the formation of a team comprising of plant managers from the firm's operations in the US, Europe and Asia, for the purpose of transferring best practices) will have different coordination problems and bottlenecks from a team that is only virtual on the team members' characteristics dimension (e.g. when the firm brings in external consultants to work with a racially diverse internal group on improving cycle time in its production plant). The former is likely to experience coordination problems arising from delays and misunderstandings due to the lack of face-to-face interaction, while the latter's coordination mishaps are likely to be attributed to conflicts caused by racial and organizational differences. By comparing the performance results, and the trade-offs, of different virtual aspects of a team, we hope to draw inferences of how those virtual dimensions impact coordination and communication within a team. Equipped with this understanding, we aim to propose strategies to managing virtual teams.

Using the VITÉ software program, which describes teams as information-processing structures (Levitt et al., 1994; Jin and Levitt, R.E., 1996), we modeled a virtual team engaged in a banking software customization and installation project. In this model, the team members are further divided into sub-teams, responsible for separate modules of the software design. These separate modules are later combined to form the integrated product. Depending on their assigned responsibilities, each member is differentially linked or interdependent with different parties. This software design team was modeled closely after the structure and workflow of real software design teams. Its members, tasks, and corresponding interdependencies were modeled after an extensive interview with an experienced systems analyst. This team structure served as our baseline model, whose simulated results are later compared with that of our models of different virtual teams.

To create different models of virtual teams, we manipulated the scenario properties of this baseline model in VITÉ. By varying the task nature, project error rates, communication processes, and team structure, we are able to model the context and behavior of team members to fit each virtual setting. Specifically, we created the (1) virtual context, (2) virtual team members' characteristics and (3) virtual relationships by manipulating (1) functional error rate, team experience, noise, meeting frequency and information exchange frequency; (2) degree of formality in information exchange and project error rate; and (3) tendency to reciprocate information requests and degree of centralization respectively. Since the VITÉ simulation model has been demonstrated to approximate information processing in real teams (Levitt et al., 1994), and we have been able to model the virtual characteristics in the simulation model to our satisfaction, we believed that our models provide a close correspondence to the different virtual team settings we see today.

Preliminary findings indicated that the team that manifests all three dimensions of virtuality experienced the greatest coordination problems. This was not surprising, given that such a team would not only encounter coordination bottlenecks caused by distance, but also communication problems arising from cultural differences and weak ties. Interestingly, the escalation of coordination difficulty did not inhere in all aspects of virtuality. Specifically, the team that was only virtual in relationships between members had lower coordination volume and costs, compared to the baseline model. With the exception of the team that is virtual only in the relationships amongst its members, there is no clear performance advantage between a team that is only virtual in terms of context and a team that is only virtual in terms of members' characteristics. The presence of tradeoffs suggests that the optimality of these two teams will be contingent on the situational demands.

These findings suggest that despite the potential flexibility and productivity gains from a virtual team, the coordination within the team to reap these gains is fraught with difficulties. In particular, the results indicate that different dimensions of virtuality have different effects on different aspects of team performance. For instance, a team that is virtual in context incurs the highest coordination costs, while a team that is virtual in team members' characteristics incurs the highest rework costs. Different patterns of interaction and communication seem to inhere in different dimensions of virtuality, suggesting that different types of strategies may be required in different virtual teams. Based on these findings, we proposed various ways to improve the performance of virtual teams.

It is clear that successful virtual teams do not only need to have a sophisticated networking infrastructure linking interdependent parties, they also require socially collaborative behaviors between team members. Rational coordination behaviors are embedded in a social network of relationships. It is not that we are ignorant of the interplay between task-oriented coordination behaviors and socially laden relationships, but is that we are perplexed by how they enable and/or inhibit the other. This complex relationship is played out in virtual team settings, where socially distant ties between team members overlap with physically distributed and tightly coupled task interdependencies. But different aspects of virtuality may generate different patterns of task interdependencies and social behaviors, thus resulting in different coordination and communication problems. By distinguishing between these different aspects of virtuality and relating them to different dimensions of team performance, we aim to simplify this complex interplay, and gain a deeper appreciation of how virtual teams work. In so doing, this study takes an initial step to developing a systematic understanding of the virtual phenomena.

REFERENCES

Barner, R. (1996). The new millennium workplace: Seven changes that will challenge managers and workers. *The Futurist*, 30, 14-18.

Clancy, T. (1994). The latest word from thoughtful executives – the virtual corporation, telecommuting and the concept of team. *Academy of Management Executive*, 8(2), 8-10.

Davidow, W. H., & Malone, M. S. (1992). *The Virtual Corporation*. New York: Edward Burlingame Books/Harper Business.

DeSanctis, G., N. Staudenmayer & S.S. Wong (in press). Interdependence in Virtual Organizations. In C. Cooper and D. Rousseau (Eds.), *Trends in Organizational Behavior*. NY: Wiley & Sons.

Hedberg, B., Dahlgren, G., Hansson, J. & Olve, N-G. (1997). *Virtual Organizations and Beyond: Discover Imaginary Systems*. NY: Wiley.

Jin, Y. and Levitt, R.E. (1996). The Virtual Design Team: A computational model of project organizations. *Computational and Mathematical Organization Theory*, 2(3), 171-196.

King, J. (1994). Network tools of the virtual corporation. *Network World*, 11, 28-30.

Levitt, R.E., Cohen, G.P., Kunz, J.C., Nass, C.I., Christiansen, T. and Jin, Y. (1994). The "Virtual Design Team": Simulating how organization structure and information processing tools affect team performance. In Kathleen M. Carley and Michael J. Prietula (Eds.) *Computational Organization Theory* (pp. 1-18). Hillsdale, NJ: Lawrence Erlbaum Associates.

Lipnack, J., & Stamps, J. (1997). *Virtual Teams: Reaching Across Space, Time and Organizations with Technology*. NY: John Wiley.

McDonald, T. (1995). Brain Trust: A new corporate climate favors loose, flexible, idea-based teams. *Successful Meetings*, 44, 16.

Simons, T. (1995). Virtual reality. *Inc*., 17(14), 23-24.

STRATEGY AND DECISION-MAKING

Improving and Extending Social Science Theories by Formal Logic: A Case Study of Resource Partitioning Theory

Ivar Vermeulen (Universiteit van Amsterdam, ivar@ccsom.uva.nl) and
Jeroen Bruggeman (Rijksuniversiteit Groningen, j.p.bruggeman@bdk.rug.nl).

The value of a field of science is in the theories it produces. For a scientific theory to be valuable, it must have a consistent, sound, and transparent argument, and unambiguous concepts. Many social science texts that present theory are not clear, though, leaving the burden of disambiguating concepts and phrases, and of checking consistency and soundness on the readers of these texts. To relieve this burden in some extent, we propose a five-step approach to systematically think through declarative texts presenting scientific theory, using formal logic and computational support. We show that by this approach new and possibly valuable conclusions can be "squeezed"' out of such texts in a rigorous manner. To show the resourcefulness of our approach, we think through resource partitioning theory (Hannan and Carroll, 1995) as a case in point. Resource partitioning has become an important part of organizational ecology (Hannan and Freeman, 1989), and deals with the population dynamics of competing generalist and specialist organizations. Its main claim is that ``Increasing concentration enhances the life chances of specialist organizations," as has been observed in the car industry, banking, newspapers, beer brewery, wine making, and hi-tech industries. This view conflicts with industrial organization economics that sees high concentration as a barrier to entry, especially for small (specialist) organizations. Our analysis of resource partitioning casts new light on both views. We therefore believe that the full-length-paper's results are of interest to organizational theorists, organizational strategists and organizational ecologists alike (Vermeulen and Bruggeman, 1998).

To find theory in social science texts, and to distinguish theory from auxiliary parts, we focus on the main claims in a text---addressing the research question(s)---and their supportive arguments. These claims and arguments taken together, we see as the core theory. The first step in our five-step approach, then, is to mark the core sentences in question.

Analyzing and sharpening key concepts in the core theory is our second step. Moreover, we look for relations between key concepts, which can be tacit but important for the argument.

In the third step, we focus on the structure of argument. We distinguish premises from conclusions, use the sharpened concepts from step 2, and (informally) axiomatize the core theory. To visualize the structure of argument, we draw a conceptual model that represents the key events and their relations as described by the premises. The interesting point is to see if the premises do support the conclusions, i.e., the claims. Loopholes in the diagram may point to tacit background assumptions or other flaws in the core theory. We subsequently make hidden background assumptions and boundary conditions explicit.

The first three steps of our approach are, for short, a rational reconstruction of a text. To be sure that our rationally reconstructed theory of resource partitioning is sound and consistent, we formalize it in step 4, in first order logic, and formally test it in step 5. Computer tools that we use in step 5---a theorem prover and a model generator---are freely available on the World Wide Web, so readers can check for themselves the logical rigor of our treatment of resource partitioning theory, or apply these tools to other theories.

In our case study, we passed resource partitioning theory ``through the purgatory of proofs and refutations," as Lakatos phrased it. As a result, we were able to prove two theorems that answer the question Under what conditions specialists proliferate and why? We showed that under certain general conditions made explicit, ``resource partitioning" and the proliferation of specialists are endogenous population processes, i.e., independent of changing consumer tastes. Our showing this independence generalizes the theory at hand. Several meanings of the concepts of specialist and generalist in organizational ecology were illuminated in the formalization process, and our analysis furthermore shows that not concentration enhances the life chances of specialists, but economies of scale instead. Under the conditions explicated, we logically infer that if scale economies come to dominate, the number of organizations in the population increases. This new theorem clarifies previously unexplained empirical findings.

Short Bibliography
Hannan, M.; Carroll, G. 1995. Organizations in Industry. Oxford.
Hannan, M.; Freeman, J. 1989. Organizational Ecology. Cambridge.
Lakatos, I. 1976. Proofs and Refutations. Cambridge.
Vermeulen, I.; Bruggeman, J. 1998. The Logic of Organizational Markets: Thinking Through Resource Partitioning Theory. Draft paper, Universiteit van Amsterdam.

The automated theorem prover, OTTER, and model generator, MACE, can be found at http://www-unix.mcs.anl.gov/AR/otter/

Acknowledgement. The heuristics used in our five-step formalization approach were developed in joint projects by all members of the CCSOM institute at the Universiteit van Amsterdam.

Operational Control in Interorganizational Policy Making

Johannes S. Timmermans and Giampiero E.G. Beroggi
jost@sepa.tudelft.nl; beroggi@sepa.tudelft.nl

School of Technology, Policy, and Management
Delft University of Technology
P.O. Box 5015, 2600 GA Delft, Netherlands

Abstract

Policy making in complex societal problems typically involves many organizations with conflicting interests and diverging control over crucial issues. The organizations, also referred to as stakeholders or actors, interact and communicate with each other in a network like structure, where interactions and negotiations can be bilateral and multilateral, and possibly involve elements of organizational, administrative or financial hierarchies. Organizational decision support has primarily focused on strategic and planning aspects and less on operational approaches. In this paper we discuss an approach to operational control in interorganizational policy making. The approach relies on

preference structuring theories from the area of decision analysis and network theories with respect to social systems and negotiation and conflict resolution. The concept has been implemented in a multimedia group decision support system at Delft University of Technology and it has been used with people in the field for real problems, including airport planning, facility location, and water resource and coastal zone management. Analysis of empirical findings, complemented by theoretical studies, confirm the viability of the approach. The paper ends with an agenda for forthcoming research.

The actors involved in complex policy making have different interests and control over a set of issues. Actors need to enter interactions and negotiations to identify profitable exchanges of control in order to improve their decision making position. The quantitative linear exchange model proposed by Coleman (1990) has been successfully applied in the area of political influence systems, where the actors gain power by exchanging the influence they have over the different issues (Pappi and Knoke, 1991). The scope of Coleman's theory is to determine the equilibrium control and to use it as an indicator for the development of policy making at the system level. Equilibrium control is defined as the optimal control resulting from appropriate exchanges, given the interests and influence structure of the actors. To compute the equilibrium control, the actors must express their interests and control for the different issues. Based on Coleman's linear system of action, the equilibrium control can be computed by simple matrix computations.

The purpose of our research is to operationalize the exchange model such that it can be used within a policy gaming setting or along an evolving real-world problem. The laboratory sessions that we conducted so far were aimed to find answers to three research hypothesis: (1) the linear system of action, describing complex problem solving as an exchange process, is a valid model for the policy process simulated in policy gaming; (2) interaction and negotiation contribute to exchanges of control (3) the operationalized system supports problem solving.

Four experimental sessions in a laboratory setting, using a computer-based multimedia policy game, have been conducted over a ten months period. They each lasted 8 hours. After a 20 minute introduction to the policy game and the problem situation and a 40 minute introduction aimed to facilitate interface control and role identification, the subjects were asked to make a first assessment of interest and control. Each of the 6 subjects had to assess its preferences for the 9 issues, and each subject had to assess for each issue the control intensity of each actor. These assessments provided a 6x9 interest matrix, and a 9x6 control matrix. These matrices were then used to compute the subjects' interdependence, the equilibrium control, the demand for control and the dependence on exchange. After this first round of assessments, the actual policy game started. Subjects could interact both through the networked computer system and in a face-to-face fashion. The networked computer system facilitates the actual exchange of control, for example, by applying for, or issuing, permits, implementing measures, or engaging in multiple actor investment projects. This interaction round lasted one hour and was repeated during the second half of the session. During the interaction rounds subjects could decide to engage in face-to-face negotiations in special break-out rooms. Any communication between the subjects, both through the computer network as well as in the break-out rooms, were registered. Computer interactions including motivations were registered in a database log file that can be queried for time, actions, and actors through an SQL report tool. Verbal face-to-face interactions were registered using a computer based observation system applying interaction categories congruent to the issues and actors in the linear system of action. Three laboratory sessions entailed a second measurement round after the first interaction round. The fourth session involved three measurement rounds. The purpose to increase the number of measurements was to be able to study the convergence of the exchange model.

The first step in analyzing the data collected in these four sessions was to understand what the values in the output matrices (e.g., interest, control, and equilibrium control) mean. Monte-Carlo simulation was conducted to arrive at a density function for control/equilibrium control differences in two subsequent measurements. A total of 10,000 random linear system of action models were generated, involving over 500,000 assessments of control for individual subjects.

The results of the analysis provide tentative support for hypotheses 1 and 2. There are several reasons why the evidence has to be put in some perspective. First of all, the problem at hand had a dynamic equilibrium control, in the sense that after each round, the subjects would reframe their problem perception and, therefore, partly re-iterate their interests. Second, group composition and role distributions create group dynamics not related to the policy game. Third, the strength of role identification varies over the subjects.

To provide stronger evidence for hypotheses 1 and 2, and to test hypothesis 3, we will conduct laboratory experiments with a simpler and more controlled problem setting. This implies that (i) the number of actors and issues will be reduced, (ii) the problem will be defined to have a fixed equilibrium control, (iii) the time for negotiations will be reduced to allow for a larger number of measurements, and (iv) the problem will be defined to maximize role identification.

References

Coleman J.S., 1990. Foundations of Social Choice Theory, The Belknap Press of Cambridge University Press, Cambridge, MA.

Pappi F.U. and Knoke D., 1991. "Political Exchange in the German and American Labor Policy Domains." In Marin B. and Mayntz R. (eds.), Policy Networks: Empirical Evidence and Theoretical Considerations. Campus Verlag, Frankfurt am Main, Germany, 179-208.

How to preserve the richness of interpretation frames and reasoning mechanisms in formalizing organization theory?[1]

Henk W.M. Gazendam[2] and John L. Simons[3]
Faculty of Management and Organization
Groningen University, The Netherlands

Introduction

In many organization theories, the interpretation frame is the most important part of the theory. Organization theorists use a rich variety of reasoning mechanisms, for instance reasoning about the relation between the individual level and the organizational level, and reasoning about actions using abstraction from time. In formalizing organization theories, one has to try to preserve the richness of the interpretation frames of those theories as well as the richness of the reasoning mechanisms used.

Fayol's organization theory offers an interesting example of a rich interpretation frame and a wealth of reasoning mechanisms. The specification of the interpretation frame has been done in ENBF (Gazendam, 1993). This specification comprises eight pages of ENBF-expressions. A translation of these specifications into logic, however, has not been done successfully thus far. Fayol's theory uses at least thirteen different reasoning mechanisms (Gazendam, 1993), several of which encounter severe difficulties in the process of translating them into logic. Most of them can, however be written down as rules in a notation that derives form ENBF, Prolog and Smalltalk. With respect to reasoning about the relation between the individual level and the organizational level, Fayol uses circular reasoning (individual properties define an organizational property, and this organization property influences in turn individual properties), aggregation, and network-based aggregation. Most organization theories do not explicitly use time as a parameter. They abstract from time. With respect to reasoning about actions using abstraction from time, Fayol uses task analysis (reasoning about the composition and best way to accomplish tasks), opportunity-based reasoning (based on preferences), equilibrium-based reasoning, and flow-based reasoning.

Formalization strategy

In order to represent the richness of interpretation frames, one can express the concepts in these interpretation frames as types using a suitable theory of types that enhances FOL. In such a theory of types, a type is seen as a combination of a concept with a set of values. Contributions to such a theory of types have been given by Devlin (1991) and De Brock (1995).

Reasoning about the relation between the individual level and the organizational level needs powerful instruments to handle functions, for instance those of FP, an extension of lambda calculus that uses operators for combining sets of functions and sets of objects. Using FP, problems of expressing distribution and aggregation in FOL can be solved elegantly.

The formalization of reasoning about actions using abstraction from time needs an analysis in which the strategy for abstracting from time is made explicit. Task analysis can be seen as a strategy to split the time-dependent description of tasks into a time-independent part, a program-like description of action sequences, on the one hand, and on the other hand a system of triggers that can described as for instance schedules or as recurrent behavior patterns in time and space. Opportunity-based reasoning distinguishes the more or less time-independent preferences of actors, from the decision situations that occur in time, in which decisions are actually made and executed. These decision

[1] Contribution to the Computational and Mathematical Organization Theory Workshop, May1st and 2nd, 1999, Cincinnati.
[2] Prof.dr. Henk W.M. Gazendam is professor of Information Systems in the Public Sector at the Faculty of Public Administration at Twente University and associate professor of Information Strategy at the Faculty of Management and Organization at Groningen University (P.O.Box 800, NL-9700-AV Groningen, The Netherlands, tel +31-50-3637090, email h.w.m.gazendam@bdk.rug.nl).
[3] Prof.dr.ir. John L. Simons is professor of Information Systems at the Faculty of Management and Organization at Groningen University (P.O.Box 800, NL-9700-AV Groningen, The Netherlands, tel +31-50-3633864, email j.l.simons@bdk.rug.nl).

situations may be the result from the system of triggers mentioned earlier, or arise in a more random way. The equilibrium concept is dependent on aggregation, and can be formalized using the instruments mentioned earlier for handling types and functions. Flow-based reasoning uses steady state parallel and sequential processes in which recurrent behavior patterns in space and time are structured. Most flows can be seen as programs describing action sequences that are triggered by stable systems of recurrent triggers.

Interpretation frames

Interpretation frames can be described using ENBF. An example taken from the description of the interpretation frame of Fayol's theory is given below.

<Fayol's organization> ::=
{<corps social>, <material organism>, <contingency factor>+, <symbol structure>+, <person&task relation>+, <person&action relation>+, <management process>+}.

<corps social> ::=
{<person>+, <person&person quality>+, <corps social quality>+, <corps social network quality>+}.

(After these eight lines of specification, another 8 pages of specifications follow in order to capture Fayol's interpretation frame.)

Capturing an interpretation frame in logic requires the use of a theory of types. A type can be seen as a structured object consisting of a concept and a corresponding range of values. The concept is defined as a basic concept or more general concept fulfilling a set of logical constraints. Devlin (1991) elaborates the concept side of a type. De Brock (1995) gives us an elaboration of the value side of a type, which is basically a set of set-valued functions.

The relationship between micro and macro characteristics

Circular reasoning

A general problem in the formalization of organization theories is the way in which relationships between micro characteristics of people and the macro characteristics of an organization should be handled. An example of this problem area is within the concept of discipline. Discipline can be defined as the state of an organization in which all person behave according to the rules and obey their superiors. This means that discipline, a macro characteristic, is derived from the micro characteristics of people. On the other hand, one can state that people will behave disciplined when discipline is a characteristic of the organization. This means that micro characteristics of people are dependent on the macro characteristics of the organization they are member of. A possible solution to this problem is to distinguish discipline as a state of the organization from discipline as a culture characteristic of the organization. This solves the loop problem but creates another: how is the cultural characteristic of discipline created, from which phenomena is it dependent? A similar problem exists between structures or states and actions. On the one hand, actions create states and structures; on the other hand, specific states or structures are the preconditions for taking specific actions.

Aggregation

We can distinguish two types of reasoning that derive qualities of the whole from qualities of the parts: simple aggregation and network-based aggregation. *Simple aggregation* determines a quality of the whole by a simple sum, average, or product of the underlying properties of the parts. An example is the determination of organizational productivity. *Network-based aggregation* is based on the determination of topological properties of the agent network. Deduction of network properties is, for instance, used in the principles of unity of command and centralization: each person has to have one and only one superior, and there is only one topmost superior or directive part.

Abstraction from time

Many organization theories do not explicitly represent time. This is also the case in Fayol's theory. Fayol uses four mechanisms for coping with actions without the necessity of explicitly representing time: task analysis, opportunity-based reasoning, equilibrium-based reasoning, and flow-based reasoning.

Task analysis
Task analysis states that to perform a task, one has to perform certain subtasks, or that a task should be done in a particular way. For instance, the establishment of social order presupposes the following subtasks according to Fayol: (1) continuous observation of the human resources of the organization, and continuous estimation of the human requirements of the organization, (2) deciding about the necessary positions, and (3) selecting personnel for these positions.
Opportunity-based reasoning
Opportunity-based reasoning uses tendencies or attitudes to describe what people would do when a certain situation would arise. For instance, if people get the chance to gain authority, they will do that, but they will try to get rid of responsibility:

> ". . .authority is sought after, while responsibility is feared”

Equilibrium-based reasoning
Equilibrium-based reasoning uses the picture of an equilibrium that must be maintained by management while people tend to take actions that disturb this equilibrium. This is, for instance, the case in the principle of subordination of individual interest to general interest, and the maintenance of social order. In both cases, selfish interests or impulses such as ambition, ignorance or simply human passions lead to actions that disturb the equilibrium; management has to prevent and counteract these disturbing actions.
Flow-based reasoning
Flow-based reasoning uses flows to describe action patterns. Flows are used to describe the actions of persons that are dependent on each other by passing information or objects. For instance, he distinguishes a command flow from the top of the organization to the bottom, and an information flow from the bottom to the top. This picture is refined further by distinguishing lateral flows of communication.

AGENTS AND LEARNING

Why not Multiple Solutions: Agent-Based Social Interaction Analysis

Authors: T. Terano, S. Kurahashi, U. Minami
Graduate School of Systems Management, The University of Tsukuba, Tokyo
address: 3-29-1 Otsuka, Bunkyo-ku, Tokyo 112-0012, Japan
tel.: +81-3-3942-6855
fax.: +81-3-3942-6829
email: terano@gssm.otsuka.tsukuba.ac.jp

Topic areas:

Computational models of the evolution of social networks;
Modeling organizational adaptation and learning;

ABSTRACT:

This paper proposes a new method: Inverse Simulation for analyzing emergent behaviors of agents in artificial societies, which aims at modeling social interactions in the electronic mediated communication.

In conventional computational society models, the simulation is executed straightforwardly: Initially, many micro-level parameters and initial conditions are set, then, the simulation steps are executed, and finally the macro-level results are observed. Unlike these approaches, Inverse Simulation executes these steps in the reverse order: set a macro-level objective function, evolve the worlds to fit to the objectives, then observe the micro-level agent characteristics. Thus, we must solve very large inverse problems. So far, it has been considered difficult to adopt such an inverse approach to social system simulation studies, however, here we succeeded by utilizing a Genetic Algorithms to evolve the societies by changing the predetermined parameters to optimize macro-level socio-metric measures, which can be observed in such real societies as e-mail oriented organizations and electronic commerce markets.

Another unique point of our approach is that, since GAs we developed utlizes tabu search techniques in Operations Research literatures, the proposed method is able to optimize multi-modal functions. This means that, from the same initial conditions and the same objective function, we can evolve different results, which we often observe in real world phenomena.

The proposed method is implemented in a simulator TRURL, which evolves artificial worlds of multiagents to socially interact with each other. The micro-level agent activities are determined by both predetermined and acquired parameters. The former parameters have constant values during one simulation cycle, however, the latter parameters change during the interactions. Unlike conventional CMOT models, TRURL evolves the societies by changing the predetermined parameters to optimize macro-level socio-metric measures, which can be observed in such real societies as e-mail oriented organizations and electronic commerce markets. Thus, using TRURL, we automatically tune the parameters up and observe both micro- and macro-level phenomena grounded in the activities of real worlds.

This paper reports experimental results on the model with multiple solutions of electronic mediated communication. The experiments are concerned with the behaviors of agnets with conforming attidudes. There are often cases where very subtle environmental changes cause a
radical change of public opinions. To analyze the situation, we evolve a computational society where the agents conform to the opinions of a single "strong agent" or a leader of the society. The results suggest that the conforming society evolve two different "strong agents": the one who has some speciality on the topic to persuade the other agents, the other who has wide knowledge on the topic and always talks to the other agents. Therefore, we believe the proposed method is very powerful to analyze organizational behaviors of agents from the CMOT viewpoints, if we could design appropriate macro-level objective functions to be optimized.

Mood and Learning in Dynamic Decision Tasks

Bud Gibson

Abstract

Many boundary-spanning jobs such as sales or credit collections urge employees to adopt and display emotions in their efforts to influence customers. The major activity of these jobs might be characterized as a dynamic decision task where the employee is attempting to achieve goals with the customer under time pressure. In this talk I will discuss my efforts to understand how pleasant and unpleasant moods with different degrees of arousal affect novices' choices of actions in dynamic decision tasks. Past evidence suggests that decision makers in more positive moods will choose more pleasant actions while those in more negative moods will choose more unpleasant actions, perhaps because brain systems that favor these choices are engaged as part of the emotional experience. Additionally, higher levels of arousal in the emotional experience may limit attentional resources that the decision maker brings to bear on the task. Together, these prior results suggest that valence of the emotion will bias decision makers' selection of actions, and high arousal may limit decision makers' exploration of alternatives, potentially damaging their learning. From the perspective of what it might take to develop a computational model of emotions' effects on learning, I will discuss my efforts to develop laboratory manipulations to distinguish the causes of performance in dynamic tasks where the adoption and display of emotions is deemed an important component.

Introduction

Call-center workers such as telesales representatives and credit card collectors interact directly with customers and are encouraged to use emotional display to influence them. For instance, credit collectors use an overall urgent tone with delinquent bankcard members as they issue warnings concerning possible consequences of continuing not to pay (Gibson & Fichman, submitted; Hochschild 1983; Sutton 1991). The combination of tone and message are expected to frighten or at least arouse collectors' targets into paying before the end of the call. Telesales representatives evoke fear of lost opportunity for similar aims.

Despite these widespread practices, there is little empirical support for the effectiveness of emotional display in influencing targets. Sutton and Rafaeli (1989) found a negative correlation between courteous emotional display, expected to raise store sales, and actual store sales. This pattern resulted from customers in busier stores pushing clerks to be more efficient at the expense of courtesy. Gibson and Fichman (submitted) found no consistent relation between collectors' use of warning (fear-evoking) messages and delinquent debtors' willingness to pay. In light of indications of dubious effectiveness, it is worthwhile considering potential drawbacks in the side-effects of emotional display.

One of the consistent side effects of emotional display is that it biases the displayer's mood toward the displayed emotion (Izard, 1993). More directly, organizational agents are frequently encouraged to adopt the emotion they are displaying (Hochschild, 1983). Beyond motivating agents to display emotion, do agents' moods in other ways bias their behavior? Are these effects potentially negative enough to warrant discouraging the use of emotional display by organizational agents? Finally, how might information about the effects of mood on agent behavior be incorporated into agent models that could then offer predictions about manipulations of work structure and conditions on performance? Such models are important for dynamic tasks because the complexity of interactions is frequently sufficient to defy descriptive verbal modeling.

Moods and Their Possible Effects on Dynamic Decision Making

Moods are a form of longer-term emotion. In general, emotions are considered to be a state of mind containing elements of: motivation, action readiness, action tendency, perceptual selectivity, cues for cognition and action, and a conscious feeling state. For instance, when faced with an unpleasant task such as calling a delinquent bankcard member, collectors may be motivated to avoid the task or be particularly aggressive. They may only focus on threat words and punitive offers. The collector may only become aware of feeling unpleasant emotion well into the onset of these other effects (LeDoux, 1996).

Examining these elements, it appears that emotions may initially affect perceptions and action tendencies. These initial effects may then bias cognitive assessments of what to do in a situation to achieve goals and how to evaluate the results of actions. In attempting to gain an understanding of how these biases function, researchers have observed that between 50% and 75% of variance in people's ratings of their own emotions can be accounted for by two factors, frequently labeled pleasantness and arousal (Russell, Weiss, & Mendelsohn, 1989, Watson, Clark, & Tellegen, 1988). Unpleasant emotions such as fear orient subjects toward defensive, fight or flight behavior (Lang, 1995). Pleasant emotions may orient toward more pleasant behaviors or possibly broadening of perspective (Schwarz & Bohner, 1996). High arousal leads to more focused attention, limiting subjects' ability to entertain multiple sources of information and making them quicker to react. Lower arousal tends to free up attentional resources.

In dynamic decision tasks such as credit collections or telesales, decision makers must consider multiple alternatives and come up with an effective action in a series of interdependent decisions. A valence/strength hypothesis based on the effects of emotion pleasantness and arousal suggests the following main effects of mood on learning in dynamic decision tasks:

1. Unpleasant mood in an agent such as a collector activate defensive behaviors. These agents are more likely to use negative warnings and may be more likely to demand more punitive terms.

2. Pleasant moods activate more pleasant behavior which may be shown in the form of encouragement and less punitive terms.

3. High arousal leads to less variety in actions. This reduction in search may lead to less exploration of the space of possible actions, thereby damaging learning as measured by the ability to take the best action in any given contingency.

4. Lower arousal (but not asleep) leads to more variety in actions. This increase in search may lead to more exploration of the space of possible actions, thereby improving learning.

These main effects suggest that, in interactions between valence and strength, higher arousal will lead to less deviation from the action set potentiated by the mood's pleasantness while lower arousal will allow for more deviation from this set.

Method and Implications for Agent Modeling

Individual learning under the influence of emotions has not been well studied (Bower, 1991), and there is little or no data on which to calibrate a learning model. Models providing a principled account of individual agent learning would seem to have to include a mechanism for *affective* information processing in addition to more cognitive mechanisms (LeDoux, 1996). For the models to be predictive, the hypothesized affective system must indicate how affect biases action selection and stimulates arousal without simply hard coding in observed effects.

I am currently at the testing stages for an empirical study of the effects of mood on behavior in dynamic decision tasks. Manipulations in this study are motivated by the valence/strength hypothesis and are designed to generate data for the more intensive theorizing implied by a modeling effort. The manipulations are operationalized in two experiments using a credit collections task. Experiment 1 is a correlational study where subjects' naturally occurring mood is measured along pleasantness and arousal as they undertake 20 computer-simulated contacts (Subjects will be deceived into thinking they are dealing with real peole on the other end). A correlational study is used to eliminate possible experimenter demand explanations for why subjects induced into particular moods might engage in their observed behaviors. In addition to their bids and offers, subjects will be required to use, at their choice, either pleasant or unpleasant messages about the cardmembers' status. Half of the cardmembers will be randomly assigned as responsive to pleasant messages and the other half to unpleasant messages.

Each contact will consist of 12 speaker turns during which subjects make offers and deliver messages to cardmembers. Measures will include the degree to which subjects learn to choose the action configurations which lead to the highest likelihood of attaining the maximum achievable payment within the specified number of days. Additionally, subjects will be measured on their exploratory behavior, operationalized as using different action configurations from speaker turn to speaker turn. The cardmember responses are probabilistic and non-obvious requiring some search to find the best configuration.

Experiment 2 induces mood for two levels of arousal crossed by two levels of valence. While mood induction potentially introduces experimenter demand effects, inducing specific moods permits distinguishing between the effects of specific moods and the more general dimensions of valence and arousal, thereby helping to distinguish between different classes of potential models. The message effectiveness manipulation and performance measures are the same as those used in Experiment 1.

Conclusion

This study is designed to shed light on how mood might affect learning and performance in dynamic decision tasks where workers are expected to display emotion as a means of influence. Mood effects are important because workers are encouraged to adopt moods to help their emotional display, and the display itself may lead to changes in mood. The study broadly examines predictions from a valence/strength hypothesis that predicts different effects for mood pleasantness and arousal. Although broadly conceived, this study should provide data that can be used to distinguish between broad classes of effects as well as initial data to be used in more detailed modeling efforts. The ultimate benefit of these efforts is that they can be used predictively to evaluate possible interventions to improve performance.

By the time of the conference, I should have data from both Experiments 1 and 2 as well as some more developed ideas about how to model these data.

References

Bower, G. H. (1992). How might emotions affect learning? In The handbook of emotion and memory: Research and theory. Hillsdale, NJ: Lawrence Erlbaum and Associates, Inc.

Gibson, F. P. and Fichman, M. (submitted). Fear appeals: What happens when the credit collector calls?

Hochschild, A. R. (1983). The managed heart: Commercialization of human feeling. Berkeley: University of California Press.

Izard, C. E. (1993). Four systems for emotion activation: Cognitive and noncognitive processes. Psychological Review, 100 (1), 68-90.

Lang, P. J. (1995). The emotion probe: Studies of motivation and attention. American Psychologist, 50 (5), 372-385.

LeDoux, J. (1996). The emotional brain: The mysterious underpinnings of emotional life. New York: Simon and Schuster.

Russell, J. A., Weiss, A., and Mendelsohn, G. A. (1989). Affect grid: A single-item scale of pleasure and arousal. Journal of Personality and Social Psychology, 57 (3), 493-502.

Schwarz, N. and Bohner, G. (1996). Feelings and their motivational implications: Moods and the action sequence. In The psychology of action: Linking cognition and motivation to behavior., New York, NY: Guilford Press; 1996, 119-145.

Sutton, R. I. (1991). Maintaining norms about expressed emotions: The case of bill collectors. Administrative Science Quarterly, 36,
245-268.

Sutton, R. I., and Rafaeli, A. (1988). Untangling the relationship between displayed emotions and organizational sales: The case of convenience stores. Academy of Management Journal, 31 (3), 461-487.

Watson, D., Clark, L. A., and Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. Journal of Personality and Social Psychology, 54 (6),
1063-1070.

Rumor, Deception and Cooperation in Organizational Learning by Boundedly Rational Agents

Michael Prietula
University of Florida
prietula@ufl.edu

Kathleen Carley
Carnegie Mellon University
Kathleen.Carley@CENTRO.SOAR.CS.CMU.EDU

Abstract

Our work focuses on understanding the relationship among boundedly rational agents, tasks, the social situations in which they are engaged, and the collective impact on organizational learning. In this work we explore a type of organizational learning that involves agents incorporating a strategy to reduce their task effort – building a coalition of trusted information sources. Agents that have trusted information sources can reduce task (and organizational) effort. However, certain agent-based propensities to trust (and deceive) and environmental turbulence may lead to difficulties with this strategy. From this, we seek to establish the elemental basis of social behavior and group phenomena, and make predictions about them.

Summary

Our work focuses on understanding the relationship among agents, tasks, and the social situations in which they are engaged from a computational perspective (Prietula, Carley & Gasser, 1998). From this, we seek to establish the elemental basis of social behavior and group phenomena, and make predictions about them. Our guide for this effort is the Model Social Agent matrix.

The Model Social Agent matrix (Carley & Newell 1994) provides a two-dimensional categorization scheme that specifies the kind of *knowledge* required by the agent(s) (in terms of increasingly complex social situations), and the kind of information *processing* capabilities required by the agent(s) to operate with that knowledge, in order to exhibit various kinds of individual and collective phenomena. The scheme is summarized in Figure 1 (modified to included the Boundedly Rational Emotional Agent).

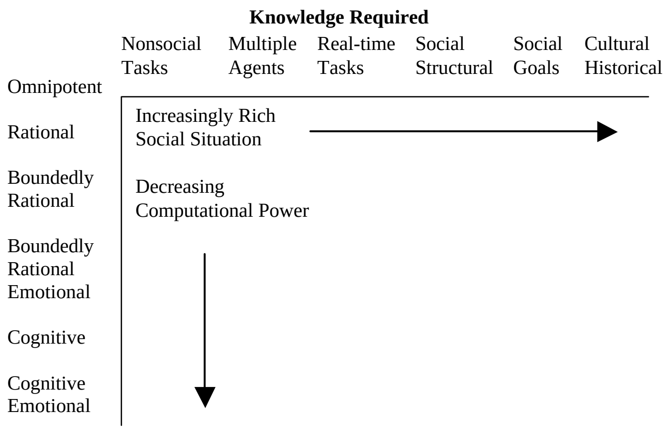

The type of agent we are at investigating in this work is the Boundedly Rational Emotional agent. These agents are restricted in both their processing capacity (what they can know) and how they use it. They rely on internal (perhaps flawed) representations of the social and task situations, but are goal-driven and bring to bear knowledge in service of those goals. Our goal for the project is to systematically simulation these agents throughout the range of social situations.

In this work we explore the impact of emotional components on (boundedly rational) agents in a social setting. As one might well imagine, there are a host of definitions for “emotion.” However, as we focus on the cognitive influences of agent affect, emotion, and behavior, our view of *emotion* is most similar to that articulated by Ortony, Core and Collins (1988) as “valenced reactions to events, agents, or objects, with their particular nature being determined by the way in which the eliciting situation is construed (p. 13).” In fact, their theory forms the basis for the elicitation structure of emotion in our model. Under certain conditions, Agents can have emotional responses to events. Basic affective reactions are differentiated with respect to cognitive constraints (as conditions) defining a fundamental set of emotional types.

In our model, however, we must account for four kinds of additional phenomena not addressed by their theory. First, the ET model provides guides for the elicitation functions, but we must incorporate some qualitative or quantitative representations that provide the specification of the underlying scales onto which differential values can be mapped.

Second, we must handle *multiple* events. The Agents in our model define a repetitive game with memory and consequence. Agents experience sequences of events (some originating from interacting with the world, some originating from Agent communications) that may (under the Ortony et al. theory), demand multiple (repeated and perhaps diverse) emotional responses. Agents alter their affective views of the world – they adapt. Thus, a theoretical elucidation of “emotional change and integration” over time must be made. What is the effect of multiple events on the eliciting conditions and subsequent emotional types triggered?

Third, a substantive result of an *emotional* response (as defined by the eliciting conditions) is often a subsequent *behavioral* response to the eliciting conditions (in the contexts of those conditions and emotional state). As emotional states are influenced by events, they too influence event choices by the Agent. Given the situation and emotional state of an Agent, what is the likelihood of specific subsequent behaviors?

To address these first three phenomena (scaling, multiple events, behavioral responses) we turn from a cognitive model to a model from sociology – affect control theory. Affect Control Theory (Heise 1979, 1987) provides the theoretical substrate and mathematical articulations of how events, agents and objects are perceived (socially) and how those perceptions influence and are influenced by social interactions (or descriptions of interactions). Furthermore, Affect Control Theory affords the mechanisms to predict affective changes (adaptation) and behavioral responses (activity) to such interactions.

Finally, the last phenomenon we explore is the *collective effect* of the previous three phenomena over time. Though informal considerations of "group emotions" have often been discussed, little empirical or theoretical guidance is provided for rigorously defining and predicting collective emotional states and behavior as we are considering them. To this end, we are striving to define such models, metrics and theoretical apparatus that can subsume the set of phenomena and explanations woven in the models we are crafting and the theoretical stance we are adopting.

The tasks we model explore the mutual effects of events, behaviors, and objects under varying agent properties, event types, group sizes, and task stability. The tasks defined within this model are simple (Prietula & Carley, in press):

> Agents seek specific items in an external search space (e.g., the Internet, a warehouse) and may cooperate (give and receive advice) as to item locations or the quality of other Agents' advice.

What complicates matters is how task and agent properties impact choices, events, affect, and emotions. As this task is communication-based (i.e., all interaction between Agents are communications), then *trust* (*in* other Agents' advice) is a natural cognitive construct to explore within the setting. In our model, trust has both cognitive and emotion-based components. As this task is also about communicating collectives, then *cooperation* (giving advice about *locations*) and *rumor* (giving advice about other Agents' *advice*) are natural organizational constructs to explore within the setting. In our model, generating advice and rumor are behavioral responses that also have both cognitive and emotional components.

In this presentation we first describe the three theoretical contributions to the paper. We then present the details of the model and how emotion, trust, cooperation, and deception fit together. We conclude with a report on the results of an initial set of simulation experiments.

References

Carley, K. & A. Newell (1994). The Nature of the Social Agent, *Journal of Mathematical Sociology*, 19(4), 221-262.

Heise, D. (1987). Affect control theory: Concepts and model. *Journal of Mathematical Sociology,* 13(1-2), 1-33.

Heise, D. (1979). *Understanding Events: Affect and the Construction of Social Action.* Cambridge, England: Cambridge University Press.

Ortony, A., G. Clore & A. Collins (1988). *The Cognitive Structure of Emotions*. Cambridge, England: Cambridge University Press.

Prietula, M. & K. Carley (in press). Exploring the effects of agent trust and benevolence in a simulated organizational task. *Applied Artificial Intelligence*.

Prietula, M., K. Carley & L. Gasser (Eds.) (1998). *Simulating Organizations: Computational Models of Institutions and Groups.* Cambridge,MA: AAAI/MIT Press.