Agent Organizations for Information Retrieval and Electronic Commerce: The Next Frontier

Les Gasser

Computational Organization Design Lab
Institute of Safety and Systems Management
University of Southern California
Los Angeles, CA 90089-0021 USA
Voice: 213.740.4046 Fax: 213.740.9732
Internet: gasser@usc.edu

Abstract

Considering organizations of agents, and addressing configuration problems of emerging network-based commerce and information management applications at organizational levels, can provide leverage. In addition, such agents will clearly be embedded and operate in rich organizational contexts. Finally, the sheer scale of global information exchange will cause segmentation and grouping of activities into patterned structures - namely organizations. We have a wide range of informal conceptual models of organization that we can bring to bear to understand these phenomena, including rational, bureaucratic, and negotiated approaches. If we draw links from these notions of organization and clearer models based on constraints and distributed search, we can begin to build more operational theories. Computational models of human and human-machine organizations will also help, and are becoming more possible at a large scale.

1 Introduction

The area of Agent-Based Systems (ABS), as both a theoretical and applications-oriented subfield of AI has undergone explosive growth over the past several years. The study and practice of Agent-Based (or Agent-Oriented) systems comprises investigations of semiautonomous, knowledge-rich, and responsible artificial entities that exist in both virtual environments (e.g. softbots) and in physical worlds (e.g. robots). New technical and infrastructural foundations such as very widespread and highly accessible internetworks, the G/NII, etc., with their burgeoning resource collections, have appeared. New techniques for modeling agents, their environments, and their interactions are coming into play, including new languages for describing agents, new communication protocols (e.g. KQML, Telescript), and new theoretical and practical research issues (e.g. security, autonomy, responsibility, etc.).

The current interest in Agent-Based Systems takes place against an existing backdrop of many years of research into Multi-Agent Systems and Distributed AI (hereinafter MAS): the branches of AI concerned with how to coordinate behavior among semi-autonomous problem-solving agents including how they can coordinate their knowledge, goals, and plans to act together, to solve joint problems, or to make individually or globally rational decisions in the face of uncertainty and multiple, conflicting perspectives.

The complementary issues of increasing the power and scope of agent-based systems, especially agent-based information management systems, and of increasing the collaborative, knowledge-sharing, coordination and emergent-behavior prospects for multi-agent systems, are coming into convergence. Agent-based and multi-agent systems are rapidly becoming practical partners in critical human problem-solving environments such as manufacturing (e.g. in recent initiatives by Allied Signal, NCMS, etc.), telecommunications management (e.g. at GTE, British Telecom, and The National Research Council of Canada), design, and enterprise coordination. Accordingly, ABS/MAS is a rapidly developing field of both application and research, with explosive growth in the US, Europe, and Japan.

2 The Information Environment

What is the situation that encourages us to consider questions of information agents/assistants, heterogeneous knowledge bases, electronic commerce, etc.? What is the overall environment here? We can roll out a by-now familiar litany of "usual suspects:"

- The emerging Global Information Infrastructure (GII). For example, in 1989, when I visited ATR in the new Kansai Science City, ATR executives told me that the long term vision of ATR was to enable communication with "anyone, anywhere, and any time." By now I'm sure that quotation has been augmented to include "using any modality" as well. Evidently, we are well on the way toward making this vision a reality, at least technologically speaking.
- The explosion of online information resources (information provision). The number of online databases, libraries, repositories, web pages, etc. currently available is staggering and growing at an enormous rate. However, most of these repositories are passive, that is, they require active retrieval and computation interventions to be useful. Moreover, the most common indexing schemes are typically fairly rigid hierarchies (i.e. they are rarely associative or statistical).
- The (expected) explosion of information demand. While the growth of online information demand has not kept pace with the supply, it is expected to grow, and possibly to outstrip the supply, creating a potential imbalance.
- Gradual development of tools for linking demand with provision. There is a real growth in the types and kinds of tools becoming available for satisfying information demand, and for linking particular demands with particular providers, including such tools as Mosaic/Netscape, information assistants, intelligent agents, softbots, gopherbots, waisbots, many information retrieval projects/initiatives, and web wanderers/spiders. The available functionalities generally fall into categories of Browsing (tools that allow for user-directed exploration), Verification (tools that verify connections or the existence of resources), Indexing (tools that automatically create network indices), Abstraction (tools that create abstracted views of networks and their contents, e.g. for cacheing and learning purposes) and Mapping (tools that construct and evolve dynamic network maps).
- Existing and rapidly developing electronic commerce: In the US there are fairly experimental and fairly formal local and regional experiments in electronic commerce such as PACT (Palo Alto Collaboration Testbed) and FAST (large-scale electronic purchasing). On a more informal scale, there are forums such as the internet shopping network, cdconnection.com with its online CD catalogues, the misc.forsale and local .forsale newsgroups with very flexible and informal mechanisms of advertisements, bids, offers, and online negotiations, even to the level of commodities. For example, the newsgroup misc.forsale.memory routinely contains postings offering memory chips and SIMMS in lots of 50,000. The Internet Society reports some 21,000 "storefronts" currently on the Internet.

3 Some Key Issues in the Information Environment

In this environment, there is a number of complex issues that cross technological and organizational boundaries, including the difficult issues of (re)organizing highly-flexible, rapid-exchange remote (as versus co-located) markets, such as:

- Transaction management issues such as security (e.g., for cash/credit transactions).
- Principal-agent issues such as assuring that interacting parties meet commitments, avoiding fraud, recourse methods to resolve problems, purchase DOA ("dead-on-arrival"), enforcement of guarantees, etc.
- Improvised and alternate-medium escrow services (e.g. Postal service COD (cash-on-delivery) for third-party escrow of both merchandise and funds).
- Policy and quasi-legal issues, including which agents should have access to which information on technical grounds (See, e.g., the "robot exclusion standard" which is an attempt to provide technical means to enforce policy decisions about limiting robotic access to network sites and files, primarily in response to heavy resource demands placed by automated information retrieval and searching [Bowman, 1994].)
- Information market organization. Information seekers will be asking for information at many levels of aggregation and analysis. Similarly, Information providers have potential capabilities to provide information at many levels of aggregation. We need to understand what are the organizational issues inherent in this sort of market, and how to think about solving them.
- Market segmentation issues (e.g., the appropriate division of markets and services among modality and infrastructure providers such as telephone, cable, and multimedia organizations).
- Content issues. The actual content of information made available and gathered is a key dimension. As Karen Coyle has pointed out, "I find it particularly puzzling that as move into this new `information age' our efforts are focused on the machinery of the information system, while the electronic information itself is being treated like just so much more flotsam and jetsam; this is not a democratization of information, but a devaluation of information." She goes on to note some "social responsibilities" with respect to information, namely Collection, Selection, Preservation, Organization, and Dissemination of information. How, when, and where these are done are matters of social-organizational choice, that are embedded in the emerging properties of information organizations and information agents.
- Issues of technology or modality transition for evolving industries. For example, the Los Angeles Times newspaper has aggressively moved into alternative modalities of information delivery including bidirectional automated-telephone, fax, and internet services, as well as the traditional unidirectional hard-copy newspaper delivery medium. At the same time, the types of automation carried out in certain information-intensive industries is not necessarily keeping pace with expectations. One investigation showed that 79% of the (US) newspapers surveyed had computer graphics capability, only 29% had a computerized library, and even fewer used information-gathering tools such as CD-ROM databases [Cable, 1994].
- Evolution and change: the continuing evolution and rapidly-dynamic nature of the GII is itself an issue. The GII is currently characterized by:
  - Rapid evolution, in terms of technology being ever more widely available, widely usable, and widely used.
  - Rapid evolution of complexity and scale of information available.
  - Rapid construction of indexing and retrieval schemes, and rapid evolution and entrenchment of retrieval/indexing schemes. Typically these are:
    - Fundamentally web-like, which means pointer-driven (e.g. as versus table-driven).
    - Fundamentally hierarchical at early stages (e.g., WWW and gopher indices, file systems.)

• Fundamentally rigid and becoming ever-more entrenched.

• Rapid evolution of social arrangements on the networks (e.g. etiquette, norms).

In a sense, though fluid, GII is also duplicating many characteristics of existing information universes (e.g. libraries) but (at present) without the additional conventionalization and social understanding (i.e., organization) that makes commonly-used information sources and repositories books and libraries so effective. What happens, for example, if we have highly dynamic indexing? To some extent, all indexing schemes are dynamic, because the uses (i.e., interpretive contexts) of the indexed information are dynamic. However, we have little understanding of how time scales of change affect interpretability of information. Suppose, for example, that indexing schemes become very dynamic - how does that affect the possibilities for commerce and organization (i.e., action patterns)?

Information organization: there is an ongoing "cultural clash" in the network world between hierarchies and webs, which will be (is already being) won by webs. That is, any particular set of information is being pointed to by many distributed pointers, and there is much duplications (mirror sites, etc.) But at the same time, there are very few backpointers--for example, where is the Science Citation Index-like backpointer structure on the web? Thus an information management concept like a web is fundamentally and organizational approach: that is, it is organized around "multiplicity of representation."

4 Why We Need To Deal With Organizations

All of this provides a perspective on the context into which we are discussing putting information management agents, and all of these will affect and shape the way we must think about information management agents/ environments and already are doing so. In particular, this perspective also begins to illustrate some of the dimensions of the very fuzzy boundaries between the technological characteristics in the design and capabilities of agents, and the organizational and social contexts in which these agents work. Indeed, certain technical capabilities of agent-based or network-oriented information management systems and transport/communications infrastructures have very strong implications for social and organizational policy that are even now causing tremendous debate. These include, for example, identification-capture technologies such as Caller ID (CNID), Automatic Vehicle Identification technologies in proposed automated highway systems, and the HTTP-FROM variable found in recent version of WWW browsers such as, Lynx, Emacs-W3, OmniWeb, Mozilla. (See, for example, the "Identification Protocol," [St. Johns, 1993], dealing with TCP connection logging.)

The overall point here is that the agents we are building and the heterogeneous databases we are connecting are already becoming parts of a huge fabric of information and activity, that is to say, a society (or a society of organizations) in its own right. This society or macro-organization is a collection of active entities that will be "doing things together" in Howie Becker's terminology [Becker, 1986]. I believe that this doing-together will be come more and more organized over time, simply because, first, the existing social organizations within which they are embedded are already highly organized (e.g. public and private enterprises of all sorts and the rules, procedures, and interpretations that structure action among them)---new active, semiautonomous technologies will have to fit in. Second, organization is a necessary response to resource limitation---organizations can augment the capabilities of individuals, and in fact are necessary to accomplish large-scale tasks that are beyond the capabilities of individuals. In addition, organization provides a certain degree of stability and predictability to the actions and interactions of individuals--that is, organization constrains and channels the activity of individuals within a set of (hopefully) useful limitations. To summarize, when thinking about designing and building heterogeneous agent-based systems, we need to understand and manage organizations:

- Because they are there.
- Because they are levers.
- Because they can provide a balance between redundancy and efficiency.
- Because they can provide a balance between stability and flexibility.

5 Conceptual Modeling of Organizations

Organization theory in general is concerned with two classes of research problem. One is how to build conceptual models of organizations that will be useful for analysis, theory-building, and prediction. The main question of this sort of research is what are the useful set of concepts and relations for describing organizations and their behaviors, and how can they be operationalized? The second class of problem is how to actually use these conceptual models to explain and predict phenomena of organizational life. Both classes of problem are also of central concern to researchers in multiagent systems, for implementing and for reasoning about collections of automated agents.

There are numerous conceptual modeling perspectives on organizations in the organization theory literature. It has been clear for some time in multiagent systems circles that organization is a powerful concept for thinking about how to structure the interactions of collections of problem solvers. Here we join the conceptual models into three categories, namely treating organizations as rational production processes, as bureaucratic decisionmaking units, and as negotiated orders in multiple social worlds. These three categories are very roughly parallel to the models of organizational decisionmaking presented by Allison [Allison, 1971].

5.1 Organizations as Rational Production Processes

One prevalent model of organization since the work of Taylor and Barnard, (see [Perrow, 1972]) in the early twentieth century, has been to treat organizations as rational production processes. In this model, organizational goals are set by strategic planning units, and carried out by functional action units. Decisions are made in the most rational manner, to maximize organizational returns. The primary activity of an organization, at any level of abstraction, is to transform raw materials or information. The actions of organizational members are rational to the extent they conform to and contribute to the overall organizational goals. The structure of organizations ideally ought to correspond to the structure that maximizes efficiency of resource transformation. The focus is on doing things that create measurable (and hence rationalizable) results, not on decisionmaking.

Individual actors in the organization, in this model, are completely rational agents, with enough knowledge to make the most rational decisions.

Some basic orienting hypotheses of this perspective would be that appropriate structure and rational processes lead to effective overall performance, that rationality is globally understood, and that evaluation criteria and values are shared.

5.2 Organizations as Bureaucratic Decisionmaking Units

Following the work of Simon and colleagues (e.g., [Simon, 1945]), another perspective has emerged: that of organizations as bureaucratic information processing units. Organizations again are resource transformers, but the structure of organizations is linked to contingencies in the environment and in the decisionmaking processes within the organization. Organizational processes are seen as structures with buffers between key, complex operations. Slack resources provide greater adaptive ability in uncertain environments, by freeing each unit to take its own corrective actions. Slack resources and buffers and decouple the impacts of local decisions across the organization, and reduce the impact of uncertainties in the organization's environment or production technologies. However, under optimal conditions of prediction and control, zero slack is possible.

The individual participants in organizations, in this view, are limited agents with bounded rationality. These agents react to pressures and uncertainties in their environments, yet exert little influence over the environment. Planning provides longer-term perspective on action, and is more possible the lower the uncertainty in the environment or technology.

Some basic orienting hypotheses of this perspective would be that appropriate degrees of resources, coupled with structure matched to environmental and process uncertainties and to limited rationality of agents, leads to effective overall performance, and that ability to share rationality, perspectives, and values is limited by processing capacity of participants.

Garbage Can Models

One particular model in the bureaucratic school is the garbage can model of Cohen, March, and Olsen [Cohen, et al., 1972], operationalized in and AI model by Masuch and LaPotin [Masuch and Lapotin, 1989]. In this bureaucratic model, procedures are limited by the information, problem mix, and available solutions at hand, at any given time. In effect, the organization's participants throw information, solutions, and problems into a "garbage can" and pick out combinations that are both available and workable. Thus organizational behavior is very opportunistic; what seems deterministic may be the simple product of circumstances.

5.3 Organizations as Constructions, Negotiated Orders and Social Worlds

From this perspective, organizations are complex and highly permeable social orders, held together by webs of mutual commitment among participants. Organizations are patterns of behavior - these behavioral patterns are be driven by the technical and information infrastructure. The boundaries between organization and environment are difficult to establish, and the activities of organizations directly affect and are affected by their environments, making directions of causality for certain actions difficult to establish. Agents participate in many social worlds (and indeed in many organizations) simultaneously, continually balancing conflicting demands and commitments in their different arenas.

As in all social orders, the ongoing norms, rules, procedures, and patterns of activity that comprise an organization are in some sense negotiated [Strauss, 1978]. As Strauss and his colleagues have put it, "Order is something at which members of any society, any organization, must work. For the shared agreements, the binding contracts---which constitute the grounds for an expectable, non-surprising, taken-for-granted, even ruled orderliness---are not binding for all time....In short, the bases of concerted action (social order) must be reconstituted continually; or as remarked above, `worked out'" [Strauss, 1963, pg. 148]. This working-out and negotiation takes place at multiple metalevels simultaneously. For example, the meaning and applicability of a rule may be negotiated in a situation---but in addition the authority of an agent to apply the rule may also be negotiated, and in fact the authority of an other agent to question the authority of the applying agent may be negotiated (cf [Gasser, 1993].

Some basic orienting hypotheses of this perspective would be that continuous multilevel negotiation of appropriate actions, rules, procedures, etc. lead to locally usable results, that overall rationality of activity is dependent on perspective, political interest, and power, and that organizational values, rules, and environments are "selected," "enacted," constructed outcomes of organizational action in an ecology of other organizations and multiple overlapping social worlds.

6 A Problem-space Model of Multi-agent Action and Organization

One useful model of multiagent and organizational systems that begins to capture some of the different aspects of organization discussed above is to see them as distributed search processes [Lesser, 1991]. To see multi-agent systems in this way, we need to build the concept of multi-level search spaces. At the "lowest," or target level (call it L0), a set of domain-problem states are connected by a set of domain-problem operators that transform one state into another. The set of available operators structures the overall space as a graph with a particular topology. A problem is a specification of states and operators, along with a distinguished start state and a set of constraints (partially) describing a goal state. Solving a problem means applying the available operators in such a way as to place the problem-solving system into a state that meets the goal-state constraints. Hence two key abilities of a problem-solver in this model are the ability to select and apply operators to take actions, and the ability to match its current state with the goal constraints to know when a problem has been solved.

The target level space may also be structured hierarchically into several abstraction levels (Call them L1 to Ln).

In this case, an operator-caused transition from one state to the next at some abstraction level Li may actually entail a collection of state transitions (or possibly an entire search process) at a more detailed descriptive level Li-k (as in SOAR [Laird, et al., 1985]). (For the moment we may ignore the metaproblem of constructing, describing, and revising the search space itself).

Action in a problem-solving system is carried out by computational processes usually called ``agents." Any problem-solving agent that operates in this problem space has local decisions to make about how to generate and explore the search space most effectively. For example, the agent may have to choose which operator to apply (and therefore which next state to visit) from a collection of options. We term the choice of what action to take next a control decision, and note that any control decision is a problem to be solved: given a collection of potential operators, select the one best operator to apply. Any knowledge brought to bear to solve such a control problem is termed control knowledge; reasoning about control in this way is often called metalevel reasoning or metareasoning.

As a problem, the control decision can also be treated using a search process. The state-space of control decisions can be represented using a model of the cause-effect relationships among problem-solving (i.e. domain operator) actions, at some abstraction level. This metalevel representation is typically called a goal structure. Many AI treatments of problem solving treat the notion of an agent's "goals" intuitively, linking them implicitly with a sort of intentional position of the agent. However, here we take a more concrete and somewhat less mentalistic position, treating what is normally termed a goal of an agent as simply a kind of data structure that the agent uses as part of a control decisionmaking process. Goals can be connected explicitly into a graph that models the problem-solving process, by using several types of edges. First, goals may be connected using and-or goal-subgoal relationships, which indicate abstraction levels in the cause-effect description of the problem-solving process. Second, goals can be linked with constraint relationships, that indicate how values of parameters in one goal state place limitations on values of other states. In particular, AND relationships among goals require several goals to be achieved together, necessitating these constraints. For example, these constraints may indicate that certain goals cannot be jointly achieved, possibly due to shared but undersupplied resources. Third, goals may be connected with other types of edges such as semantic or class-subclass links that provide additional descriptive knowledge at the control level. The structure of goals, goal-subgoal and constraints may be generated in advance, or may be generated on-the-fly as part of an opportunistic, reactive, or environment-driven problem-solving process.

Knowing which goals to address enables an agent to select problem-solving operators (namely the operators that are known to best achieve those goals), but the agent still has the problem of selecting among alternative possible goals (subgoals connected with or relationships). An agent must make decisions about which goals to address next, which is a "reasoning-about-control" or metacontrol decision, and which can be treated at a still higher (e.g., strategic) level, again using a search process. This upward metalevel recursion can be continued to arbitrary levels (cf. [Hayes-Roth, 1985; Maes, 1988].)

Into any of these search spaces we can construct agents, each of which takes on some subpart of the task of exploring the entire space searching for goal states. That is, since each of the domain and control level problem-solving processes can be represented as a search process, each can be addressed as a multi-agent problem, by partitioning the search space among multiple agents. The introduction of agents requires partitioning a search space into subspaces and allocating each of them to one or more agents. The problem of partitioning the space into subspaces is the problem of decomposition (also a problem of organization design), and the issue of which agent(s) to assign to which subspace search is the problem of task allocation. (Each of these problems can also be treated as search, and hence each can be treated as a multiagent, multi-agent problem.) If different agents search partial spaces that have states in common, their activity is said to be redundant. (Redundancy may be useful, for failure tolerance, or problematic, for reasons of efficiency).

6.1 Uncertainty

As Lesser has pointed out [Lesser, 1991], the very activity of partitioning the problem-space introduces uncertainty about what control choices any agent should make. This is because with partitioning, each agent has limited knowledge of what choices other agents are making. The uncertainty becomes an important issue, because control choices interact, due to interacting goal constraints. One agent's control choice may eliminate the possibility for, make redundant, or assist, another agent's action. If control is defined as next-action choice, we can define control decision uncertainty as ambiguity in that choice. That is, the greater the ambiguity in control choice, the greater the control decision uncertainty. Control decision uncertainty can be characterized as the size of the set of highest-valued next-states whose values cannot be distinguished. (We are assuming that heuristic knowledge that allows a provider to distinguish among the expected values of next-states is, on average, accurate; more highly valued next-states actually are better performers on the average.) However, this localized notion of control decision uncertainty is insufficient for optimal control because simply considering the ambiguity in an individual decision doesn't capture the effects of that decision on other successive decisions. In this way, individual control decisions have impacts on the future control uncertainty that an individual problem-solver faces. Thus control decision uncertainty may be better characterized as uncertainty over a range of related decisions in a subtree, than over a single decision, because the range of decisions accounts for both the individual decision uncertainty and its future impacts. The difficulty is in characterizing the future impacts, because they are, after all, in the solver's future.

Many kinds of uncertainty are introduced by the partitioning of a search space, including:

- Uncertainty about the existence and character of remote states. search space
- Uncertainty about goal relationships with and among remote states.
- Uncertainty about actions of other agents in the search space.

V. Lesser has presented another treatment of the types and impacts of uncertainty in distributed problem solving [Lesser, 1991]. His uncertainty types include solution uncertainty (uncertainty over whether a local solution will be incorporated into a top-level goal) and control uncertainty (uncertainty over goal-ordering, problem-solving context for a goal, and how much effort to expend toward solving a goal.

6.2 Reducing Uncertainty and its Impact

Most coordination approaches involve combinations of two strategies: reducing the degree and/or reducing the impact of uncertainty in local-level or network-level decisions. The impact of control uncertainty is felt in the arbitrariness of control decisions, and its effects are related to the density of goals in the search space. On average, for constant goal density, greater control uncertainty would be expected to increase search effort.

The impact of control uncertainty can be reduced by reducing coupling between the activities of separate providers. Coupling can be reduced by reducing common dependencies that providers share. Two kinds of common dependencies are logical dependencies and resource dependencies. Two tasks are logically interdependent when the ability to solve one task depends upon the solution reached by another task. AND-subgoaling and logical contradiction are examples this kind of dependency.

In a sense, resource dependencies can give rise to a kind of logical dependency. For example, the logical relationship between two mutually exclusive goals might arise because they both require a consumable resource that is in short supply. If the solution to one of them consumes the resource, solving the other will be impossible, and vice-versa. In this case, it is possible to restructure the goal relationships to remove the dependency by supplying more of the scarce resource. When there is enough resource so that so that both goals can be solved, the dependency no longer exists. Such a reduced dependency reduces the impact of individual control uncertainty, freeing each agent to act more independently.

Dependencies influence control decisions. For example, because one agent's set of viable control alternatives depends on another agent's actions, uncertainty may be increased or decreased by the control choices and actions of the other agent. In a sense, in the context of goal distribution and asynchronous activity, dependencies may introduce meta-uncertainty (uncertainty about uncertainty) into the control decision. This control meta-uncertainty is uncertainty about what is the actual level of control uncertainty (i.e. what is the true composition of the set of viable next-action alternatives).

Obviously, communication plays an important role here, both in reducing this meta-uncertainty, and in establishing the actual degree of control uncertainty. Thus the combination of communication and dependency (constraint) propagation can be used together to modulate the degrees of control uncertainty and control meta-uncertainty. This type of uncertainty reduction is an instance of how contextual information can be used to reduce uncertainty by reducing the degrees of freedom in control choices.

Contextual information can also improve coordination by influencing control decisions. Imagine that two agents have been allocated to different subspaces that are related by a mutual exclusion constraint. One agent (Agent 1) may be able to deduce information about the probability of being able to solve its goal (e.g. by looking ahead and deciding that its goal will be difficult or impossible for it to solve, or maybe by actually solving it). This agent's new knowledge can help both it's own control decisionmaking and that of the other agent (Agent 2). Perhaps Agent 1 decides to work on an alternative goal because it discovers that the mutually-constrained goal has lower probability of solution. This information, if communicated to Agent 2, can influence Agent 2's ranking of its own goals since Agent 2 now knows that its goal is more likely to be solvable if the other agent is not working on it. Alternatively, suppose that Agent 1 actually solves the mutually-constrained goal. When communicated, this knowledge can reduce Agent 2's actual control uncertainty, because Agent 2 then knows that the mutually-constrained goal must be removed from its set of next-action alternatives. This knowledge also reduces Agent 2's control meta-uncertainty because now Agent 2 is completely certain that the mutually-exclusive goal is no longer a member of its set of next-action alternatives. Prior to the communication, Agent 2 was uncertain about the status of its interacting goal, and was therefore (meta-) uncertain about its own control uncertainty.

Overall, then, the distributed search formulation of multi-agent problem solving places several requirements on the structure and capabilities of agents. Agents, from a multi-agent perspective, typically comprise the following characteristics:

- Knowledge-based: They must have large bodies of problem-solving knowledge related to the goal structure of the problem subspace they are investigating; they are complex problem-solvers.
- Interactive: They must have some communication and interaction capability.
- Complex, adaptive internal control: They must have flexible internal control structures, and must be able to make control choices (possibly at multiple metalevels) as well as domain-level problem-solving choices.
- Ability to model others: They must incorporate some explicit mechanism of reasoning about other agents in their world, and in particular of incorporating knowledge of other agents into their own local control choices. Thus knowledge of other agents is useful just because it provides one kind of control knowledge.

6.3 Organizations and Uncertainty

Finally, uncertainty has long been recognized as playing a key role in organizations [Galbraith, 1973; Perrow, 1972]. Organizations themselves are uncertainty-reduction mechanisms in that they create and are based upon stable relationships among participants. In addition, they may act to create predictability in their environments. Particular loci of power in organizations have been linked to the ability to control important uncertainties facing the organizations. Thus, in a multiagent system in which control effectiveness/efficiency is a key determinant of performance, and in which information can be controlled, one would expect that participants who are better at managing control uncertainty both locally and remotely would become centers of power.

Conceptualizations of the idea of organization, and the available approaches to implementing organizations in multi-agent systems, provide the ground upon which to build computational mechanisms of organizational structure, adaptation, and design. Several conceptual approaches to organization have been introduced in the multi-agent literature, including treating organization as 1) a long term, strategic load-balancing technique [Corkill, 1982l Durfee, et al., 1987], 2) a structural set of control and communication relationships among agents [Malone, 1987], 3) sets of interaction patterns among agents [Gasser, 1986; Hewitt, 1977], 4) sets of commitments and expectations among agents, [Bond, 1988; Gasser, 1991; Hewitt, 1991], or 5) collections of settled and unsettled questions about knowledge and action [Gasser, et al, 1989].

Conceptual foundations for organization in multi-agent systems may also be characterized along a spectrum with two poles [Gasser, 1990, Gasser, 1991]. Purely individualist approaches (generally the standard in multi-agent systems) build an organization out of individual, pre-existing agents with relatively fixed internal structure and one locus of action. These agents interact with each other under some set of internally or externally-enforced constraints, and it is the constraints that provide the organization. In purely social-interactionist approaches, neither the structure of the individual agents nor the nature of the organization is necessarily fixed (see, e.g., [Ishida, et al., 1992]. Instead, here agents and organization are both treated as flexible constructions, carved out of a fabric of distributed interactions. Relationships between problem-solving knowledge, resources, and the loci of action are variable. Agents might thus be distributed and concurrent entities, and their boundaries and contents might change. Organization consists of emergent patterns of interaction, and is relative to an observer's viewpoint.

Approaches to implementing organization have generally included two sorts. Structure-based organizations use fixed interaction structures or capability restrictions to configure actions (e.g., by establishing roles among identical problem solvers using capability constraints, as in [Durfee, et al., 1987]). Structure-based organizations are changed by changing the structural properties of the organization, such as the number or types of agents, their interaction structures (e.g., interagent connections or communication channels) or by modifying agent capability constraints (e.g., allowing an agent access to a different set of its own decision procedures). In knowledge-based organizations, the particular distribution and use of knowledge configures actions (e.g., flexible networks of default knowledge proposed in [Gasser, 1989]). Modifying the knowledge that agents have---e.g., about the beliefs, goals or capabilities of another agent---or changing the distribution of knowledge in the group, causes changes in the possible and actual patterns of action, and thus changes the organization. (Cf. the discussion of negotiated order approaches to organization, above.)

Numerous organizational morphologies, including many variants of market-like and hierarchical forms, have been proposed in efforts to structure behavior in multi-agent systems [Bond, 1988; Davis, 1983; Fox, 1981; Gasser et al., 1989]. Comparative information processing performance of rigid organization structures has been studied by Malone [Malone, 1987]. However, since no single organization is appropriate in all situations, organization self-design (OSD) has been proposed to allow an organization of problem solvers to adapt itself to dynamically changing situations [Corkill, 1982; Ishida, et al., 1992].

7 Possibilities for a Unified View of Organizations

Finally, many of the difficulties in unifying multi-agent systems research and organization theories arise because of the limited scope of current problem-solving models in multi-agent, and because of very limited notions of the nature of organization and the number and types of concepts and variables useful for characterizing organization. Organization theory has such conceptual variablility, and richness, even in the most conservative and well-accepted approaches, that multi-agent models would currently be hard-pressed as comprehensive modeling frameworks. Nonetheless, the advantages of explicit knowledge representations and clear procedural methods seem to be encouragements for increased attention to multi-agent techniques as modeling tools in organization theory. Multi-agent approaches can provide clarity, conceptual modeling techniques, and simulability not now found in organization theory research.

Similarly, the knowledge bases of organization theory can have direct impact in multi-agent systems, by 1) giving stronger and clearer explanations for phenomena observed in multi-agent systems (e.g. by supplying conceptual and comparative frameworks useful for multi-agent systems, such as those of Williamson, Perrow, Simon, or Strauss), 2) by giving empirically derived, predictive or causal knowledge useful for reasoning about organizational adaptation and design, and 3) by presenting a set of standards and empirical modeling challenges, which multi-agent research can aspire to meet.

8 An Integrative Organizational Model

For several years, we have been working toward an integrative organizational model that is applicable along a spectrum of organizations including those comprising people, people and machines, and machine-only collectives. In the model we have built, integration of an organization refers to the degree to which the organization functions well when viewed from a collection of different perspectives at one time. That is, when the organization is seen to be working well when viewed from alternative perspectives, it is said to be well-integrated along those perspectives.

We can model effective organizational integration using the open-systems (OS) perspective of matching and coordinating numerous organizational features (rather than optimizing just one of them), and of optimizing this match to fit organizational objectives and environmental circumstances. The quality of an integrated organization design can be measured as the degree of congruence among organizational system characteristics, the ability of an organization to provide attributes deemed to be needed for a particular strategic configuration, the ability to achieve organization objectives, and the ability to control or eliminate variances.

We have implemented a variant of this theory for human organizations, specifically for production areas of discrete parts manufacturing organizations. This version of the theory captures several thousand relationships among some 500-600 integrative features, from two sets of categories. First, it models how organizational objectives and strategies support and constrain one another. The seven organizational objectives modeled thus far in the theory include:

- Minimizing throughput time. (The goal of minimizing the cumulative time it takes to carry out a task or set of tasks efficiently while meeting quality and cost goals.)
- Maximizing output (product) quality. (The goal of maximizing the quality of the process by meeting output specifications the first time through the process).
- Maximizing agent flexibility. (The goal of minimizing the time lost from core operations due to the overhead cost of adding capabilities to agents (e.g., "training," "reskilling") or the goal of maximizing the number of different tasks that agents can effectively perform).
- Maximizing product responsiveness. (The goal of quickly developing and bringing out new organizational products.)
- Maximizing process responsiveness. (The goal of quickly changing or modifying a core organizational process for a new or modified organizational product, or making significant improvements to the organization's core process in order to improve quality, cost, or schedule.)
- Maximizing changeover responsiveness. (The goal of quickly changing the organizational process to accommodate different existing products through a portion of the overall process.)
- Maximizing producibility of designs. (The goal of optimizing the compatibility of the design of the organizational product with the core organizational processes).

Second, the integrative theory models how each of these objectives is related to a detailed set of integrative factors from a set of other categories, as well as how these integrative factors support and constrain one another. These the integrative factor categories include:

- Specific technical characteristics of the production system (e.g. tools available to agents for performing work.).
- Specific characteristics of the production process (dependencies, information and resource flows, etc.).
- Information, tool, and technology resources and their attributes.
- Variances in the internal and external organizational environments (e.g., reliability of agents, information quality or certainty).
- Task group features (activities in a task group, resource and focus requirements of a task group, task group loads, etc.).
- Organizational unit structure features including reporting relationships and control relationships.
- Features of coordination among task groups and organizational subunits.
- Capability ("skill") needs and opportunities.
- Types of authority to make particular sets of decisions.
- Measurement, feedback, and other performance management systems characteristics.
- Supplier- and consumer-involvement variables.
- Evaluation criteria utilized by decisionmaking agents.

Both theoretically and in its software realization, the open systems model treats each integrative factor as a variable in a large network of constraints. The organization theory specifies the variables, their potential values (value domains), and the mutually reinforcing or constraining relationships among them. The purpose of this model is to allow the model to be used to (1) capture and track a very large number of features and their interactions, (2) to discover specific points of congruence and misalignment among features, as a basis for redesign or trade-off analysis for optimizing specific organizational objectives, and (3) to flexibly revise and explore alternative model inputs and outputs and their impacts.

Relationships of support and constraint are modeled at two levels of abstraction , called summary and detailed levels. Many of the same concepts are modeled at each level, but the specificity and complexity varies across levels. At the detailed level, detailed hierarchical relationships among TOP features are modeled using cascades of up to seven-way relationships in a large graph structure, called a theory minimodel. These minimodels are tree-like in structure.

At the summary level, pairwise relationships between factors are modeled, typically by using theoretical matrices. The cells of a matrix indicate which relationships are supportive or constraining. The summary-level theory employs 23 such matrices, modeling relationships between objectives and variances, as well as between each of these and the organizational features such as information resources, skills, feedback system characteristics, loci of decisionmaking discretion etc mentioned above.

Knowledge used for summary analysis and design is abstracted from the knowledge used for detailed analysis and design. For example, the organizational goal of minimizing overall computation (e.g., throughput) time is linked to the detailed concept PERCENTAGE OF PROCEDURES THAT ARE RELIABLE through a dense network of concepts and relationships in the throughput minimodel. However, in the summary analysis, this complex relationship has been simplified, and is represented as a pairwise constraint in the summary relationship matrix. In this way, summary relationships directly reflect underlying detailed theory, in abstracted form.

At both the summary and detailed levels, theoretical variables and their interactions are modeled, in mathematical terms, as constraint satisfaction problems [Gasser et al., 1993b]. This constraint-based theoretical model has the representational advantages of (1) directly reflecting the open-systems orientation, and (2) accounting for equifinality---the possibility that alternative organization configurations accomplish similar objectives. In addition, it has the mathematical advantage that it can be used to reason in any direction. Thus, by using organizational objectives as inputs (i.e., independent variables), the constraint model can be used (when augmented by a user-guided search process) as a design structure---it can generate characteristics of appropriate organizations such that, to the extent possible, these characteristics jointly conform to the theoretical model's constraints and achieve the business objectives. Alternatively, if organizational characteristics are input (treated as independent variables), the propagation of resulting constraints through the network provides a way of evaluating the quality of an organization with respect to the theory. For example, degree of congruence can be measured as a function of the number and type of violated constraints.

In either direction (design or evaluation) if the entire theory network is underconstrained (basically, no violations and some remaining degrees of freedom), the interpretation is that there are several solutions, all equifinal with respect to the knowledge encoded in the theory. The more specific the theory, the fewer the equifinal outcomes in a given circumstance. Note that underconstrained situations focus the researcher's or designer's attention on how to refine the theory or organization design for specific circumstances, so as to narrow the number of equifinal outcomes.

Theory Structure: The structure of a theory is the basic architecture of relationships among its key elements. For example, theories are generally built as a set of propositional relationships among elements of a collection of (possibly taxonomized) concepts. Some theories are structured as a collection of predictive rules. Others are structured as systems of constraints. The structure and representation of a theory are important because they impact complexity, comprehension, applicability, and testability of the theory.

The abstract structure and representation machinery for our theory is is the multilevel network of constraints. It is built from a) a set of conceptual element types that describe the classes of concepts used in the theory (e.g., theory variables), and b) a set of constraint relationships among the element types and their instances.

We map a particular theory's content into this general structure by specifying particular element and constraint types relevant to the theory's domain (in this case, organizations), and by creating instances of the element types specific to refinements of the theory domain. Finally these instances are linked together with theoretical propositions about how they affect one another, represented as constraint relationships.

For example, our current organization theory includes, as element types, several organizational strategy elements (e.g., organizational objectives, variance control strategies), and eleven major kinds of descriptive features for organizations (the features mentioned above). These element types allow the theory to apply to many kinds of organization, both automated and human.

Specific instances of these element types refine the theory to a specific class of organization. For example, our current theory currently includes instances such as "Process Changeover Responsiveness" (one of seven instances of "Organizational Objective"), "Proactively Eliminating Incoming Information Quality Variance" (one of 66 instances of "Variance Strategy"), and "Interpreting Output Quality Specifications" (one of 30 instances of the "Capability/Skill" element type). Constraint relationships link these. In this example case, one theoretical proposition (as a constraint) relates the strategy of controlling information quality variances with the need for the skill of interpreting quality specifications.

In addition to this structure, our representational machinery for theories includes a collection of measurement functions and comparative functions. The measurement functions aggregate or measure attributes of the overall theory structure (such as how many different information sources are predicted to be needed to implement a particular strategic choice, or how many reciprocal dependencies exist in a task structure). The comparative functions generate comparisons between measurements of actual values for theory elements, and theoretical predictions (e.g., how many needed information sources are actually provided in the organization). Finally, the representational machinery contains facilities to actively propagate effects of constraints through the theory structure, yielding a theory-driven "qualitative simulation" of the effects of different organizational or strategic configurations. The model aggregates the specific elements into an overall assessment of a specific design. In this way, representational tools include a model of the (qualitative) likelihood of successfully achieving a given business objective given a particular set of design elements.

This entire set of representational, analysis, and application machinery is domain-independent, and new specific theories can be "plugged in" easily using data files. This pluggability has been verified by plugging in some 30 different versions of the our domain theories, as well as via preliminary experiments with several other organizational types, including various types of human organizations.

9 Conclusions

The need to treat organizations of agents in the context of emerging network-based commerce and information management applications is clear, because of the leverage gained by addressing problems at organizational levels, because of the clearly organizational contexts into which agents will be embedded, and because the sheer scale of global information exchange will necessarily lead to segmentation and grouping of activities into patterned structures - namely organizations. The conceptual models of organizations that we have available cover a range from purely rational, through bureaucratic and rule-driven, to highly flexible and negotiated patterns based on ongoing, multilevel interactions. We can begin to draw links between these relatively informal notions of organization, and some more clearly-structured models based on formulations of distributed search and networks of constraint. We can begin to build operationalizable theories of organization that account for degrees of organizational integration, equifinality, and structural integrity by combining insights from human organizations studies, computational modeling and analysis methodologies, and experimentation.

10 Acknowledgements:

I thank Toru Ishida and Takahira Yamaguchi for helping to make possible my attendance at this workshop. I gratefully acknowledge the insights provided by discussions on these issues with Dan Oleary, Vic Lesser, and Takahira Yamaguchi. Ann Majchrzak and Bryan Borys provided helpful comments on certain sections. Phil Agre's Red Rock Eater internet news service (rre@weber.ucsd.edu) has been an invaluable source of raw information and insight.

11 References

[Allison, 1971] Graham T. Allison, The Essence of Decision. Little, Brown and Company, Boston, 1971.

[Becker, 1986] H.S. Becker, Doing Things Together, Northwestern University Press, 1986.

[Bond 88] Bond88a Alan H. Bond and Les Gasser, Readings in Distributed Artificial Intelligence, San Mateo, CA: Morgan Kaufman, 1988.

[Bowman, et al. 1994] C.M. Bowman, P.B. Danzig, U. Manber, and M.F. Schwartz, Scalable Internet Resource Discovery: Research Problems and Approaches." Communications of the ACM, 37:8, pp. 98--107, August,

[Cable, 1994] "Technology in the Newsroom" Cable & Broadcasting, 10/31/94 p.26.

[Cohen, et al., 1972] M.D. Cohen, J.J. March, and J.P. Olsen. "A Garbage-Can Model of Organizational Choice," Administrative Science Quarterly, 17, pp. 1--25, 1972

[Corkill 82] D.D. Corkill, "A Framework for Organizational Self-Design in Distributed Problem Solving Networks," PhD. Dissertation, Technical Report COINS-TR-82-33, University of Massachusetts, 1982.

[Coyle, 1994] Karen Coyle, "ACCESS: Not Just Wires," Talk given at the 1994 CPSR Annual Meeting, San Diego, CA, 8 October, 1994, (Available from kec@stubbs.ucop.edu).

[Davis 1983] Randall Davis and Reid G. Smith, "Negotiation as a Metaphor for Distributed Problem Solving," Artificial Intelligence, Vol. 20, pp. 63-109, 1983.

[Durfee, et al., 1987] E. H. Durfee, V. R. Lesser, and D.D., Corkill, "Coherent Cooperation among Communicating Problem Solvers," IEEE Transactions on Computers, Volume C-36, pp. 1275-1291, 1987.

[Fox 1981] Fox, M., "An Organizational View of Distributed Systems," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-11, January, 1981.

[Galbraith, 1973] J. Galbraith, Designing Complex Organizations, Addison-Wesley, 1973.

[Gasser 86] Les Gasser, "The Integration of Computing and Routine Work," ACM Transactions on Office Information Systems, 4:3, pp., 205-225, July, 1986.

[Gasser, 1989] Les Gasser, Nicholas Rouquette, Randall Hill and Jon Lieb, "Representing and Using Organizational Knowledge in DAI Systems," in Les Gasser and Michael N. Huhns, Editors, Distributed Artificial Intelligence, Volume II, London:Pitman, pp. 55-78, 1989.

[Gasser, 1990] Les Gasser, "Conceptual Modeling in Distributed Artificial Intelligence," Journal of the Japanese Society for Artificial Intelligence, Vol. 5:4, July, 1990.

[Gasser, 1991] Les Gasser, "Social Conceptions of Knowledge and Action," Artificial Intelligence, January, 1991.

[Gasser, 1993] L. Gasser, "Heterogeneity in Practice," in Proceeedings of the International Joint Conference on Artificial Intelligence (IJCAI-93), Chambery, France, 1993.

[Gasser, et al., 1993] Les Gasser, Ingemar Hulthage, Brian Leverich, Jon Lieb, and Ann Majchrzak, "Organizations as Complex, Dynamic Design Problems," in Miguel Filgueiras and Luis Damas, eds. Progress in Artificial Intelligence: Lecture Notes in Artificial Intelligence 727, Springer Verlag, 1993.

[Hayes-Roth, 1985] B. Hayes-Roth, "A Blackboard Architecture for Control," Artificial Intelligence, Vol. 26, pp. 251-321, 1985.

[Hewitt, 77] Carl Hewitt, "Viewing Control Structures as Patterns of Passing Messages," Artificial Intelligence, 1977.

[Hewitt, 1991] Carl Hewitt, "Open Systems Semantics for Distributed Artificial Intelligence," Artificial Intelligence, January, 1991.

[Ishida et al., 1992] Toru Ishida, Les Gasser, and Makoto Yokoo, "Organization Self-Design of Distributed Production Systems," IEEE Transactions on Data and Knowledge Engineering, 4:2, pp 123-134, April, 1992.

[Koster, 1994] Martijn Koster, A Standard for Robot Exclusion, URL http://web.nexor.co.uk/mak/doc/robots/robots.html, 1994.

[Laird, et al., 1985] Jon Laird, Allen Newell, and Paul Rosenbloom, "SOAR: An Architecture for General Intelligence," Artificial Intelligence, 1985.

[Maes, 1988] Pattie Maes and Daniel Nardi, Editors, Metalevel Architectures and Reflection, North-Holland, Amsterdam, 1988.

[Malone 87] Malone87 T. W. Malone, "Modeling Coordination in Organizations and Markets," Management Science, Vol. 33, No. 10, pp. 1317-1332, 1987.

[Masuch, 1989] Michael Masuch and Perry LaPotin, "Beyond Garbage Cans: An AI Model of Organizational Choice," Administrative Science Quarterly, 33:4, March, 1989.

[Lesser, 1991] Victor R. Lesser, "A Retrospective View of the FA/C Architecture: DAI as Distributed Search." IEEE Trans. Systems, MAn, and Cybernetics, November, 1991.

[Perrow, 1972] C. Perrow, Complex Organizations: A Critical Essay (2nd Edition). Scott, Foresman, 1972.

[Simon, 1945] Herbert A. Simon, Administrative Behavior. MacMillan, New York, 1945.

[St. Johns, 1993] M. St. Johns, RFC1413: Identification Protocol, URL URL http://www.cis.ohio-state.edu/htbin/rfc/rfc1413.html, February, 1993.

[Strauss, 1963] A.L. Strauss, L. Schatzman, R. Bucher, D. Ehrlich, and M. Sabshin, "The Hospital and its Negotiated Order," in E. Friedson, ed., The Hospital in Modern Society. Free Press, New York, 1963.

[Strauss, 1978] A.L. Strauss, Negotiations. Jossey-Bass, San Francisco, 1978.