A Survey of Multi-Agent Organizational Paradigms

Bryan Horling and Victor Lesser
Multi-Agent Systems Lab
Department of Computer Science
University of Massachusetts, Amherst, MA
{bhorling,lesser}@cs.umass.edu

May, 2005

Abstract

Many researchers have demonstrated that the organizational design employed by an agent system can have a significant, quantitative effect on its performance characteristics. A range of organizational strategies have emerged from this line of research, each with different strengths and weaknesses. In this article we present a survey of the major organizational paradigms used in multi-agent systems. These include hierarchies, holarchies, coalitions, teams, congregations, societies, federations, markets, and matrix organizations. We will provide a description of each, discuss their advantages and disadvantages, and provide examples of how they may be instantiated and maintained. This summary will facilitate the comparative evaluation of organizational styles, allowing designers to first recognize the spectrum of possibilities, and then guiding the selection of an appropriate organizational design for a particular domain and environment.

1 Introduction

The organization of a multi-agent system is the collection of roles, relationships, and authority structures which govern its behavior. All multi-agent systems possess some or all of these characteristics and therefore all have some form of organization, although it may be implicit and informal. Just as with human organizations, such agent organizations guide how the members of the population interact with one another, not necessarily on a moment-by-moment basis, but over the potentially long-term course of a particular goal or set of goals. This guidance might influence authority relationships, data flow, resource allocation, coordination patterns or any number of other system characteristics (Hayden et al., 1999; Carley and Gasser, 1999). This can help groups of simple agents exhibit complex behaviors and help sophisticated agents reduce the complexity of their reasoning. Implicit in this concept is the assumption that the organization serves some purpose – that the shape, size and characteristics of the organizational structure can affect the behavior of the system (Galbraith, 1977). It has been repeatedly shown that the organization of a system can have significant impact on its short and long-term performance (Carley and Gasser, 1999; Sandholm et al., 1999; Durfee et al., 1987; Horling et al., 2004; Matson and DeLoach, 2003; Barber and Martin, 2001; So and Durfee, 1998; Brooks and Durfee, 2003), dependent on the characteristics of the agent population, scenario goals and surrounding environment. Because of this, the study of organizational characteristics, generally known as computational organization theory, has received much attention by multi-agent researchers.

It is generally agreed that there is no single type of organization that is suitable for all situations (Ishida et al., 1992; Corkill and Lander, 1998; Lesser, 1998; Carley and Gasser, 1999). In some cases, no single organizational style is appropriate for a particular situation, and a number of different, concurrently operating organizational structures are needed (Gasser, 1991; Horling et al., 2003). Some researchers go so far as to say no perfect organization exists for any situation, due the inevitable tradeoffs that must be made and the uncertainty, lack of global coherence and dynamism present in any realistic population (Romelaer, 2002). What is clear is that all approaches have different characteristics which may be more suitable for some problems and less suitable for others. Organizations can be used to limit the scope of interactions, provide strength in numbers, reduce or manage uncertainty, reduce or explicitly increase redundancy or formalize high-level goals which no single agent may be aware of (Lesser and Corkill, 1981; Fox, 1981). At the same time, organizations can also adversely affect computational or communication overhead, reduce overall flexibility or reactivity, and add an additional layer of complexity to the system (Horling et al., 2004). By discovering and evaluating these characteristics, and then encoding them using an explicit representation (Fox et al., 1998), one can facilitate the process of organizational-self design (Corkill and Lesser, 1983) whereby a system automates the process of selecting and adapting an appropriate organization dynamically (Lesser, 1998; Schwaninger, 2000). This approach will ultimately enable suitably equipped agent populations to organize themselves, eliminating at least some of the need to exhaustively determine all possible runtime conditions a priori. Before this can occur, the space of organizational options must be mapped, and their relative benefits and costs understood.

These benefits and costs, and the potential advantages that could be provided by technologies able to make use of such knowledge, motivate the need to determine the characteristics of organizations and under what circumstances they are appropriate. While no two organizational instances are likely to be identical, there are identifiable classes of organizations which share common characteristics (Romelaer, 2002). Several organizational paradigms suitable for multi-agent systems have emerged from this line of research (Fox, 1981). These cover particularly common, useful or interesting structures that can be described in some general form. In this paper we will describe several of these paradigms, give some insight into how they can be used and generated, and compare their strengths and weaknesses. The vast amount of research which has been done in this field precludes a complete survey of any one technique; we hope to provide the reader with a concise description and a sample of the interesting work that has been done in each area.

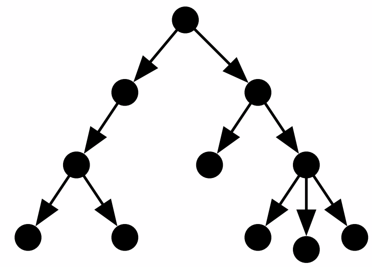

Figure 1: A hierarchical organization.

In the following sections, we will describe the origin, form, function and characteristics of a typical structure for each organizational paradigm. Example applications will be presented, along with a discussion of techniques that have been employed to create the structures. By separating these concepts, we will distinguish between the characteristics of the organization generation process and those of the organizational structure itself, independently of how it was generated.

2 Hierarchies

The *hierarchy* or *hierarchical organization* is perhaps the earliest example of structured, organizational design applied to multi-agent system and earlier distributed artificial intelligence architectures (Fox, 1979; Lesser and Erman, 1980; Fox, 1981; Davis and Smith, 1983; Bond and Gasser, 1988; Malone and Smith, 1988; Montgomery and Durfee, 1993). Agents are conceptually arranged in a tree-like structure, as seen in Figure 1, where agents higher in the tree have a more global view than those below them. In its strictest interpretation, interactions do not take place across the tree, but only between connected entities. More recent work (Mathieu et al., 2002) has explored starting with a strict hierarchy and augmenting it with cross links to allow more direct communication, which can reduce some of the latency that results from repeated traversals up and down the tree.

The data produced by lower-level agents in a hierarchy typically travels upwards to provide a broader view, while control flows downward as the higher level agents provide direction to those below (Bond and Gasser, 1988). The simplest instance of this structure consists of a two-level hierarchy, where the lower level agents' actions are completely specified by the upper, which produces a global view from the resulting information (Chandrasekaran, 1981). More complex instances have multiple levels, while data flow, authority relations or other organizationally-dictated characteristics may not be absolute.

Fox (Fox, 1979) describes several different types of organizational hierarchies. The *simple* hierarchy endows a single apex member with the decision making authority in the system. *Uniform* hierarchies distribute this authority in different areas of the system to achieve efficiency gains through locality. Decisions are made by the agents which have both the information needed to reason about the decision, and the organizational authority to do make the decision. Each level acts as a filter, explicitly transferring information and implicitly transferring decisions up the hierarchy only when necessary. *Multi-divisional* hierarchies further exploit localization by dividing the organization along "product" lines, where products might represent different physical artifacts, services, or high-level goals. Each division has complete control over their product, which facilitates the decision making and resource allocation process by limiting outside influences. The divisions themselves may still be organized under a higher-level entity which evaluates their performance and offers guidance, but is strictly separated from the divisional decision process. These more sophisticated hierarchies look very much like like holarchical organizations, which are discussed in Section 3.

2.1 Characteristics

The applicability of hierarchical structuring comes from the natural decomposition possible in many different task environments. Indeed, task decomposition trees are a popular way of modeling individual agent plan recipes (Decker, 1996); a hierarchical organization can be thought of as an assignment of roles and interconnections inspired by the global goal tree. The hierarchy's efficiency is also derived from this notion of decomposition, because the divide-and-conquer approach it engenders allows the system to use larger groups of agents more efficiently and address larger scale problems (Yadgar et al., 2003). This type of organization can constrain agents to a number of interactions that is small relative to the total population size. This allows control actions and behavior decisions become more tractable, increased parallelism can be exploited, and because there is less potentially distracting data they can obtain a more cohesive view of the information pertinent to those decisions (Montgomery and Durfee, 1993).

It is not sufficient to simply aggregate increasing amounts of information to obtain higher utility or better performance. This information must be matched with sufficient computational power and analysis techniques to make effective use of the information (Lesser, 1991). Without this, the effort to transfer the data may be wasted and the excess information distract the agent from more important tasks. Alternatively, the information can be summarized, approximated or otherwise processed on its way up the tree to reduce the information load. However, in doing so, a new dimension of uncertainty is introduced because of the potential for necessary details to be lost. In this situation, the decision making authority should be correctly placed within the structure to maximize the tractable amount of useful information that is available that retains an acceptable level of uncertainty or imprecision (Fox, 1979; Lesser and Corkill, 1981).

Using a hierarchy can also lead to an overly rigid or fragile organization, prone to single-point failures with potentially global consequences (Maturana et al., 1999). For example, if the apex agent were to fail the entire structure's cohesion could be compromised. Of course this agent could be replaced, but it may then prove costly to restore the concentrated information possessed by its predecessor. It is similarly susceptible to bottleneck effects if the scope of control decisions or data receipt is not effectively managed – consider what would happen if that apex agent received all the raw data produced by a large group of agents below it.

2.2 Formation

Although the algorithm itself does not enforce a strict hierarchy such as the one described earlier, Smith's contract net protocol (Smith, 1980; Davis and Smith, 1983) provides a straightforward mechanism to construct a series of connections with most of the same characteristics. In some of this early contract net work, the protocol was to explicitly form long-term organizational relationships, rather than the short-term contracts it has been typically used for more recently. The hierarchical structure that is produced by the process is implicitly based on the way the high-level goal can be decomposed. Upon receipt of a new task, an agent first chooses to perform the task itself, or search for agents willing to help complete the task. As part of this search process, the agent may decompose the task into subtasks or *contracts*. The agent, acting as a contractor, announces these contracts along with a bid specification to a subset of its peers who then decide if they wish to submit a bid. The bids which return to the contractor contain relevant information about the potential contractee which allows it to discriminate among competing offers. A contractee is selected and notified. Upon receipt of the new task, the contractee now faces the same question - should it perform the task itself or contract it out? Repeated invocations of this process produce a hierarchy of contractors and contractees. Because agents individually choose which contracts to bid on, and contractors choose which bids to accept, this strategy can effectively assign tasks among a population of agents without the need for a global view. The drawback to this approach is that it is myopic. Because the contracting agent does not necessarily take into account the needs of other contractors, it may bind scarce resource in suboptimal ways. For example, it may select a particular bid when viable alternatives exist, even though that particular bidder is critical to another agent (Sims et al., 2003).

As with most organizational structures, the shape of the hierarchy can affect the characteristics of both global and local behaviors. A very flat hierarchy where agents have a high degree of connectivity can lead to overloading if agent resources are both limited and consumed as a result of these connections. Conversely, a very tall structure may slow the system's performance because of the delays incurred by passing information across multiple levels. One approach to making this tradeoff is the use of agent *cloning* (Ishida et al., 1992; Decker et al., 1997; Maturana et al., 1999). An agent in such a system may opt to create a copy or clone of itself, possessing the same capabilities as the original, in response to overloaded conditions. If additional resources are available for this clone to use, this process allows the agent to dynamically create an assistant that can relieve excess burden from the original, reducing load-related errors or inefficiencies in the process. If the new agent is subordinate to the original, then a hierarchical organization will be formed in the process. Shehory (Shehory et al., 1998) discusses using cloning when other task-reallocation strategies are not viable. In this work, an agent's overall load is a function of its local processing, free memory and communication. It uses a dynamic programming technique to compute an optimal time to clone, and an appropriately idle computational node to house the new agent. The clone receives a subset of the original task(s). The clones themselves require resources, and the results they produce may require an additional hop to get to their ultimate destination, so they may also be merged or destroyed when these costs outweigh their benefits.

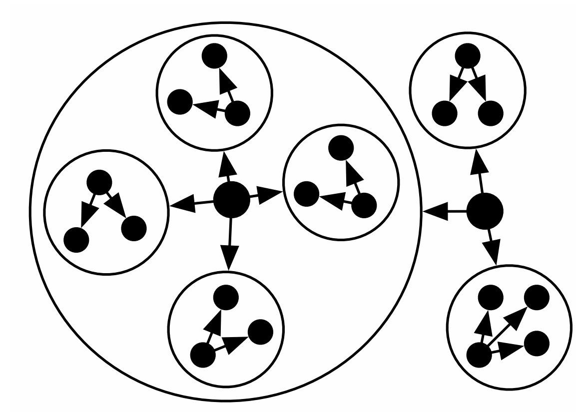

Figure 2: A holarchical organization.

3 Holarchies

The term *holon* was first coined by Arthur Koestler in his book *The Ghost In The Machine* (Koestler, 1967). In this work, Koestler attempts to present a unified, descriptive theory of physical systems based on the nested, self-similar organization that many such systems possess. For example, biological, astrological and social systems are all comprised of multi-leveled, grouped hierarchies. A universe is comprised of a number of galaxies, which are comprised of a number of solar systems, and so on, all the way down to subatomic particles. Each grouping in these systems has a character derived but distinct from the entities that are members of the group. At the same time, this same group contributes to the properties of one or more groups above it. The structure of each of these groupings is a basic unit of organization that can be seen throughout the system as a whole. Koestler called such units holons, from the Greek word *holos*, meaning “whole”, and *on*, meaning “part”. Each holon exists simultaneously as both a distinct entity built from a collection of subordinates and as part of a larger entity.

True to Koestler’s intent, this notion of a hierarchical, nested structure does accurately describe the organization of many systems. This concept has been exploited, primarily in business and manufacturing domains, to define and build structures called *holarchies* or *holonic organizations* which have this dual-nature characteristic. A sample such organization is shown in Figure 2. In this diagram, hierarchical relationships are represented as directed edges, while circles represent holon boundaries. Enterprises, companies, divisions, working groups and individuals can each be viewed as a holons taking part in a larger holarchy. Fischer (Fischer, 1999), Zhang (Zhang and Norrie, 1999), and Ulieru (Ulieru et al., 2001) have each organized agent systems by modeling explicit or implied divisions of labor in real-world systems as holons. In doing so, they create abstractions of these divisions, imparting capabilities to individual holons instead of individual agents. This layer of abstraction allows other entities in the system to make more effective use of these capabilities, by reasoning and interacting with the group as a single functional unit.

The defining characteristic of a holarchy is the partially-autonomous *holon*. Each holon is composed of one or more subordinate entities, and can be a member of one or more superordinate holons. Holons frequently have both a software and physical hardware component (Zhang and Norrie, 1999; Ulieru, 2002), although this does not preclude their usage in purely computational domains. The degree of autonomy associated with an individual holon is undefined, and could differ between levels or even between similar holons at the same level. There is the presumption, however, that the level of autonomy is neither complete nor completely absent, as these extremes would lead to either a strict hierarchy or an unorganized grouping, respectively. Within the holarchy, the chain of command generally goes up – that is, subordinate holons relinquish some of their autonomy to the superordinate groupings they belong to. However, there is also the more heterarchical notion that individual holons determine how to accomplish the tasks they are given, since they are likely the locus of relevant expertise. Many holonic structures also support connections between holons across the organization, which can result in more amorphous, web-like organizational structures that can change shape over time (Fischer, 1999; Zhang and Norrie, 1999).

It would not be incorrect to conclude that a holarchy is just a particular type of hierarchy. If we relax our definition of hierarchy to allow some amount of cross-tree interactions and local autonomy, the two styles share many of the same features and can be used almost interchangeably. These richer models then begin to resemble and take on the characteristics of nearly-decomposable hierarchies (Simon, 1968), where lateral interactions are weak but still relevant. Very flat holarchies can also begin to resemble federations, which will be discussed in Section 8.

3.1 Characteristics

As with the conventional hierarchies from the previous section, holarchies are more easily applied to domains where goals can be recursively decomposed into subtasks that can be assigned to individual holons (although this is not essential). Given such a decomposition, or a capability map of the population, the benefits the holonic organizations provide are derived primarily from the partially autonomous and encapsulated nature of holons. Holons are usually endowed with sufficient autonomy to determine how best to satisfy the requests they receive. Because the requester need not know exactly how the request will be completed, the holon potentially has a great deal of flexibility in its choice of behaviors, which can enable it to closely coordinate potentially complementary or conflicting tasks. This characteristic reduces the knowledge burden placed on the requester and allows the holon’s behavior to adapt dynamically to new conditions without further coordination, so long as the original commitment’s requirements are met. A drawback to this approach is that, lacking such knowledge, it is difficult to make predictions about the system’s overall performance (Bongaerts, 1998).

3.2 Formation

The challenge in creating a holonic organization revolves around selecting the appropriate agents to reside in the individual holons. The purpose of the holon must be useful within the broader context of the organization's high-level goals, and the holon's members must be effective at satisfying that purpose. Zhang (Zhang and Norrie, 1999) uses a model of static holons along with so-called mediator holons to create and adapt the organization. The static groups consist of product, product model and resource holons, each of which corresponds to a group of physical or information objects in the environment (e.g. manufacturing device, design plans, conveyors, etc.). The mediator holon ties these together, by managing orders, finding product data and coordinating resources in a manner similar to a federation, which will be discussed in Section 8. Each new task is represented by a dynamic mediator holon (DMH), which is created by the mediator holon. The DMH is destroyed when the task is completed.

Another approach to holarchy construction uses *fuzzy entropy* minimization to guide the formation of individual holonic clusters (Stefanoiu et al., 2000; Ulieru, 2002). In this work, the collection of holons is assumed to be initially described with a set of source-plans, each of which describes a potential assignment of holons to clusters, along with a set of probabilities that describe the degree of occurrence of those clusters. From this initial uncertain information, one can derive the preferences which agents have to work with one another, and then choose the source plan which has the minimal entropy with respect to those preferences. The goal of this technique is to ensure that each holon has the necessary knowledge and expertise needed to perform its task. The preference that one agent has for another represents this knowledge or expertise requirement, so the minimally fuzzy set will satisfy this goal by clustering agents which have common preferences. In (Ulieru, 2002), Ulieru adds a genetic algorithm approach to this scheme to help explore the space of possible clustering assignments.

4 Coalitions

The notion of a *coalition* of individuals has been studied by the game theory community for decades, and has proved to be a useful strategy in both real-world economic scenarios and multi-agent systems. If we view the population of agents A as a set, then each subset of A is a potential *coalition*. Coalitions in general are goal-directed and short-lived; they are formed with a purpose in mind and dissolve when that need no longer exists, the coalition ceases to suit its designed purpose, or critical mass is lost as agents depart. Related research has extended this to longer-term agreements based on trust (Breban and Vassileva, 2001) and to the iterative formation of multiple coalitions in response to a dynamic task environment (Mérida-Campos and Willmott, 2004). They may form in populations of both cooperative and self-interested agents.

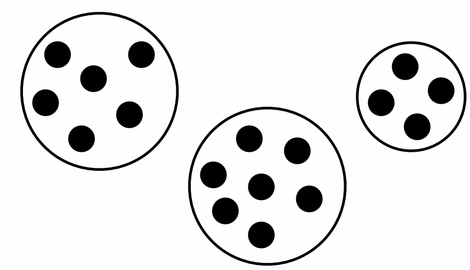

Figure 3: A coalition-based organization.

A population of agents organized into coalitions is shown in Figure 3. Within a coalition, the organizational structure is typically flat, although there may be a distinguished "leading agent" which acts as a representative and intermediary for the group as a whole (Klusch and Gerber, 2002). Once formed, coalitions may be treated as a single, atomic entity. Therefore, although coalitions have no explicit hierarchical characteristic, it is possible to form such an organization by nesting one group inside another. Overlapping coalitions are also possible (Shehory and Kraus, 1998). The agents in this group are expected to coordinate their activities in a manner appropriate to the coalition's purpose. Coordination does not take place among agents in separate coalitions, except to the degree that their individual goals interact. For example, if one coalition's goal depends on the results of another, these two groups might need to agree upon a deadline by which those results are produced. In this case, it would be the leading or representative agents forming the commitment, not arbitrary members of the coalition.

In addition to the problem of generating coalition structures, one must also determine how to solve the goal presented to the coalition. If the population is self-interested, a division of value to be apportioned to participants once that goal has been satisfied must also be generated and agreed upon (Sandholm and Lesser, 1997).

4.1 Characteristics

The motivation behind the coalition formation is the notion that the value of at least some of the participants may be superadditive along some dimension. Analogously, participants' costs may be subadditive. This implies that utility can be gained by working in groups – this is the same rationale behind buying clubs, co-ops, unions, public protests and the "safety in numbers" principle. For instance, in an economic domain, a larger group of agents might have increased bargaining strength or other monetary reward (Tsvetovat and Sycara, 2000). In computational domains we might expect more efficient task allocation, or the ability to solve goals with requirements greater than any single agent can offer (Shehory and Kraus, 1998). In physically-limited systems, coalitions have been used to trade off the scope of agent interactions with the effectiveness of the system as a whole (Sims et al., 2003). This last application directly affects the coordination costs incurred by the system; we will see that this capability and purpose are shared by congregations in Section 6.

One could argue that all agents in the environment should always join to form the all-inclusive *grand coalition*. Indeed, under certain circumstances this is appropriate, since the structure would have the resources of all available agents at its disposal, which theoretically would provide the maximum value. There are costs associated with forming and maintaining such a structure however, and in real world scenarios this can be both an impractical and unnecessarily coarse solution (Sandholm and Lesser, 1997). Therefore, the problem of coalition formation becomes one of selecting the appropriate set(s) $S \subseteq A$ which maximizes the utility (value minus costs) that coalition v_S can achieve in the environment. The value and cost of the coalition are generic terms, which may in fact be functions of other domain-dependent and independent characteristics of the structure.

4.2 Formation

The complexity of the coalition formation task depends on the conditions under which the coalitions will exist, and the types of coalitions which are permitted. As with all organizations, operating in dynamic environments will be harder to maintain than in static ones. Additional complexity is also incurred if the partitioning of agents is not disjoint; that is, agents can have concurrent membership in more than one coalition. Uncertain rewards, self-interested agents and a potential lack of trust while coordinating add further obstacles to the process.

Sandholm (Sandholm et al., 1999) analyzes the worst case performance of forming exhaustive, disjoint coalitions over a static agent population from a centralized perspective. They show that by searching only the two lowest levels of a complete coalition structure graph, an a-approximate value solution can be found to the partitioning problem, where $a = |A|$. Although the search of 2^{a-1} possible allocations still grows exponentially with a, the fraction of coalition structure needing to be searched approaches zero. They also present an anytime algorithm which can meet tighter bounds given additional time. Later work empirically evaluates the average-case performance of three anytime search techniques (Larson and Sandholm, 2000). The algorithms' performances varied by domain characteristics; and no single technique was best in all conditions.

Shehory (Shehory and Kraus, 1998) has studied how coalitions may be used to enable task achievement by a group of agents. In their scenario, a set of interdependent (precedence) tasks must be accomplished, some of which require multiple agents to perform. The agents are cooperative and potentially heterogeneous in their capabilities. The strategy they employ draws on techniques used by Chvatal's greedy set covering algorithm (Chvatal, 1979), which tries to find the minimum set of subsets that together contain each member of a target set. The initial values of all possible size-bounded coalitions are first computed and then iteratively refined in a distributed manner by the agents, taking into account task ordering and capability requirements. Once computed, the highest valued coalitions, either disjoint or overlapping depending on the selection algorithm, are instantiated. This algorithm was also augmented to support dynamically arriving tasks. A drawback to this addition is that, in the worst case, the organization process needs to be redone for each task, incurring a significant communication cost. Also limiting the potential scalability of this approach is the need for each agent to have full knowledge of the available agents and tasks.

Lerman (Lerman and Shehory, 2000) presents a scalable strategy where coalitions are formed between self-interested agents based only on local decision making. In this work agents operate in an electronic marketplace consisting of a number of extant purchase orders, with the objective of forming or joining a coalition of buyers that satisfied a need at the lowest price. Coalitions form around purchase orders, where agents form or join a coalition by adding a purchase request to an order, and can leave that coalition by removing their request. Agents in the system can move at will between purchase orders, searching for the one which offers the best value (lowest cost). An analysis based on differential equations shows that this strategy reaches equilibrium (later work (Lerman and Galstyan, 2001) expands on these mathematical techniques to analyze other distributed behaviors). It also has low communication and computational requirements. However, it does not provide guarantees on the achievable value or convergence rate, which would be affected by scale, and does not have a notion of deadlines on the purchase orders.

Soh (Soh et al., 2003) presents a technique where coalitions are dynamically created in response to the recognition of tracking tasks in a distributed sensor network. In this work, agents are assumed to have incomplete, uncertain knowledge and must respond to events in real time for goal achievement to be possible. As such, coalitions are formed in a saticificing, rather than optimal manner. An agent initiates coalition formation by first using local knowledge to select a subset of candidate partners that it believes will satisfy its requirements, both in terms of capabilities and willingness to cooperate. Next, it sequentially engages these candidates, in utility-ranked order, in argumentative negotiation, where offers and counteroffers are exchanged. This proceeds until satisfactory membership is decided, or the candidate list is exhausted. Agents are cooperative, so during this negotiation process agents explicitly decide what coalition(s) they are willing to join based on perceived gains in utility. This approach does not make any guarantees about coalition value, or even that a satisfactory coalition will be found, but given the relatively short time in which an allocation must be made it would seem to be a reasonable strategy. In addition, reinforcement learning is used over the course of events to estimate candidate utility more accurately and select the most beneficial negotiation strategy, which should improve coalition value in the long run for reasonably stable environments. By storing preferences over multiple episodes, this learning also implicitly adds longevity to coalitions, giving organizational structures produced by this technique an interesting mix of dynamic and long-term characteristics.

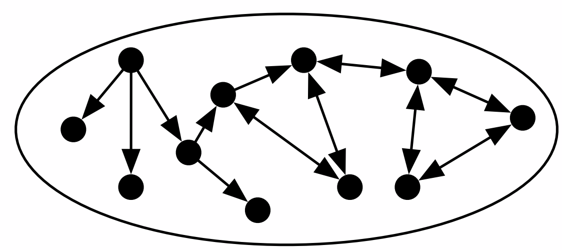

Figure 4: A team-based organization.

5 Teams

An agent *team* consists of a number of cooperative agents which have agreed to work together toward a common goal (Fox, 1981; Tambe, 1997; Beavers and Hexmoor, 2001). In comparison to coalitions, teams attempt to maximize the utility of the team (goal) itself, rather than that of the individual members. Agents are expected to coordinate in some fashion such that their individual actions are consistent with and supportive of the team's goal. Within a team, the type and pattern of interactions can be quite arbitrary, as seen in Figure 4, but in general each agent will take on one or more roles needed to address the subtasks required by the team goal. Those roles may change over time in response to planned or unplanned events, while the high-level goal itself usually remains relatively consistent (although exception handling may promote the execution of previously dormant subtasks).

This description of agent teams is quite general, and nearly any cooperative agent system has characteristics that are similar to these, if only implicitly. However, systems that maintain an explicit representation of their teamwork or joint mental state are differentiated in their ability to reason more precisely about the consequences of their teamwork decisions (Jennings, 1995; Grosz and Kraus, 1996; Tambe, 1997). For example, they will typically have representations of shared goals, mutual beliefs and team-level plans. This type of representation provides flexibility and robustness by allowing the agents to explicitly reason about team-level behaviors, where a less explicit system may rely on a set of assumptions that ultimately make the system brittle in the face of unexpected situations.

5.1 Characteristics

The primary benefit of teamwork is that by acting in concert, the group of agents can address larger problems than any individual is capable of (Grosz and Sidner, 1990). Other potential benefits, such as redundancy, the ability to meet global constraints, and economies of scale can also be realized (Hexmoor and Beavers, 2001). However, it is the ability of the team (members) to reason explicitly about the ramifications of inter-agent interactions which gives the team the needed flexibility to work in uncertain environments under unforeseen conditions. The drawback to this tighter coupling is increased communication (Parker, 1993), so the team and joint goal representations, domain characteristics and task requirements are frequently used to determine what level of cooperation (and therefore communication) is needed (Pynadath and Tambe, 2002).

Jennings (Jennings, 1995) describes an electricity transportation management system which employs teamwork to organize the activities of diagnostic agents. Lacking such structure, the agents were prone to incoherent and wasteful activities, since they did not always share useful behavior information or propagate important environmental knowledge. By providing agents with an explicit representation of shared tasks and the means by which cooperation should progress, the agents were able to accurately reason about and resolve these interactions by employing team-level knowledge. Similarly, in (Tambe, 1997), teamwork is used to provide the structure and coordination needed by agents to address interdependent goals in dynamic environments, such as tactical military exercises and competitive soccer games. These works demonstrate how pathological, but hard to predict failures can be addressed if the plans are backed up by a general model of teamwork.

5.2 Formation

The challenges associated with team formation involve three principal problems: determining how agents will be allocated to address the high-level problem, maintaining consistency among those agents during execution, and revising the team as the environment or agent population changes (Jennings, 1995; Marsella et al., 2001; Tidhar et al., 1998).

The selection and role-assignment of agents that will work on the high-level problem depends on the goal's requirements, the capabilities of the candidate agents, and the knowledge of the selecting process itself (Tidhar et al., 1996; Beavers and Hexmoor, 2001). Initially, the process or agent performing the team construction must be aware of the agents which could potentially form the team. In the case of a static, reasonably sized agent population this can be done off-line as part of the system design, or the members can be dynamically discovered and assessed. This latter technique can be accomplished using well-known discovery mechanisms such as the contract net protocol (Smith, 1980) or matchmaker intermediaries (Sycara et al., 1997). Once a suitable pool has been found, the capabilities and preexisting responsibility of those agents must be evaluated relative to the needs of the goal. Typically, agents are each denoted to have a set of capabilities, while the goal's subtask(s) are of a particular type. If an agent's capabilities include that subtask's type, it can perform the task (Tidhar et al., 1996; Fatima and Wooldridge, 2001). The discovery mechanisms may include an implicit ranking technique, such as the bidding process employed in contract net, which makes the selection process relatively straightforward. Tidhar (Tidhar et al., 1996) suggests a different technique where the agent characteristics are derived at compile time, either through designer input or automatic analysis of the agent's plan library. Candidate teams comprised of a subset of those agents may also be specified, which also are marked with their characteristics. At runtime, these characteristics are matched with the goal requirements as part of the team allocation search. By including these characteristic labels, the number of possible team combinations can be greatly reduced.

Tambe's STEAM (Tambe, 1997) architecture provides a flexible method for representing and adapting team behaviors. It is based on the joint intentions framework (Levesque et al., 1990), which formally defines how agents should reason over joint commitments and shared goals, and SharedPlans theory (Grosz and Kraus, 1999), which provides a formal way to encode and reason about joint plans, intentions and beliefs. Together, these help ensure a consistency of belief, or a desire to enact such a belief, across all team members. The commitments formed through the joint intentions process provide the explicit structure needed to reason about and monitor performance on a team level. Team plans are represented using a hierarchical decomposition tree, with nodes representing tasks for both teams and individuals, with associated preconditions, application and termination rules. Agents may simultaneously take part in several different tasks, and corresponding roles. The team's cohesion is derived primarily from the joint intentions created as part of executing the team plans. Upon selecting a team task, agents first broadcast this intention to affected agents, and wait until a commitment to that task has been established between all participants. The existence of this commitment directs agents to propagate changes whenever the task is perceived to be achieved, unachievable or irrelevant, before taking local action itself. This trades off the potential reaction speed of the team and the cost of communication with group conformity. A decision theoretic approach is used to guide communication acts, which explicitly trades off the costs of communication with those of inconsistent beliefs. Nair (Nair et al., 2003a) has also explored the possibility of using simulated emotions to provide the motivation to enforce team-level behaviors.

In STEAM, monitoring and repair of the team is accomplished with the use of role constraints (Tambe, 1997). Team members are assigned a role, based on the particular task they are working on. These roles are further constrained such that some particular combination of them (e.g. and, or) are needed to accomplish the task. One can then monitor if a task is achievable by monitoring the health of the individual agents, and using that information to evaluate the satisfiability of the role constraints. Such monitoring can be performed through explicit queries, environmental observations or by eavesdropping on communication, which can reduce the increased communication usually associated with teams. Kaminka (Kaminka et al., 2002) has demonstrated that the latter technique can perform well when coupled with a plan-recognition algorithm. Failures can thus be detected, and potentially resolved through an appropriate role-substitution, or the task abandoned if no substitution is possible. Alternately, one could use a diagnosis system (Jennings, 1995; Horling et al., 2001) to more precisely identify the root cause of the failure. Interestingly, this repair operation can itself be cast as a team task, so mutual agreement that a repair is necessary must be achieved before potentially drastic measures are taken. Nair (Nair et al., 2003b) shows how an MDP incorporating team and role-allocation knowledge can improve the system's response in cases of multiple role failure. In this case, a suitable locally optimal policy for the reallocation decision problem can be found by analyzing the team's plans, and then used to guide online responses to failures. This work showed that such policies can provide improved performance versus more heuristic and analytic techniques. A similar technique was also shown in that work to improve initial role allocation.

Tidhar (Tidhar et al., 1998) uses a similar hierarchical plan representation to represent teamwork in a tactical air mission scenario. Team membership and role assignment are performed by matching agent capabilities to one or more role's requirements. As in STEAM, teams can be broken down into sub-teams, and agents may use both implicit (observation) and explicit (messaging) forms of coordination.

The Generalized Partial Global Planning (GPGP) framework also employs techniques that allow agents to act using team semantics (Decker and Lesser, 1992; Lesser et al., 2004). Where a STEAM-driven system will typically organize in an explicit, controlled fashion in response to a perceived goal, a GPGP-team is created in a more dynamic, emergent fashion. GPGP agents are provided with a set of individual plans which model a range of alternative ways that goals may be achieved. The subgoals modeled in these plans may affect or be affected by other agents in the environment, although this may not be initially recognized. By communicating with one another and exchanging plans and schedules, these *non-local interrelationships* between tasks may be recognized. For example, the results from one agent's activity may be a strict prerequisite for another agent's task. They may alternately be a facilitating, but not required input to a

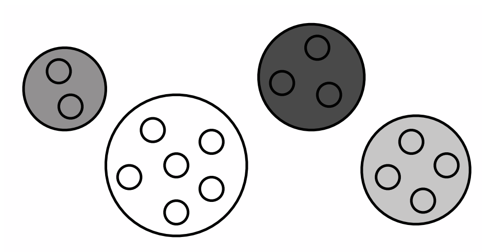

Figure 5: Congregations of agents.

task. By recognizing these interrelationships, and sharing knowledge of what goals are being pursued, agents gradually build an internal model of how their actions may affect others. This knowledge is similar to that created by the more formal joint intentions of STEAM, and allows agents to influence local behavior and communicate results as if they were members of a common team.

6 Congregations

Similar to coalitions and teams, agent *congregations* are groups of individuals who have banded together into a typically flat organization in order to derive additional benefits. Unlike these other paradigms, congregations are assumed to be long-lived and are not formed with a single specific goal in mind. Instead, congregations are formed among agents with similar or complementary characteristics to facilitate the process of finding suitable collaborators, as modeled in Figure 5. The different shadings in this figure represent the potentially heterogeneous purpose behind each grouping, in comparison to the typically more homogeneous coalitions in Figure 3. Individual agents do not necessarily have a single or fixed goal, but do have a stable set of capabilities or requirements which motivate the need to congregate (Brooks et al., 2000; Griffiths, 2003). Analogous human structures include clubs, support groups, secretarial pools, academic departments and religious groups, from which the name is derived.

Congregating agents are expected to be individually rational, by maximizing their local long-term utility. Group or global rewards are not used in this formalism (Brooks et al., 2000). It is this desire to increase local utility which drives congregation selection, because it is the utility that can be provided by a congregation's (potential) members that determine how useful it is to the agent. Agents may come and go dynamically over the existence of the congregation, although clearly there must be a relatively stable number of participants for it to be useful. Agents must also take enough advantage of the congregation so that that the time and energy invested in forming and finding the group is outweighed by the benefits derived from it. Since congregations are formed in large part to reduce the complexity of search and limit interactions, communication does not occur between agents in different congregations, although the groups are not necessarily disjoint (i.e., an agent can be a member of multiple congregations). The net result of the congregating behavior is an arrangement that can produce greater average utility per cycle spent computing or communicating (Brooks and Durfee, 2002).

6.1 Characteristics

Although congregations can theoretically share many of the same benefits of coalitions, their function in current research has been to facilitate the discovery of agent partners by restricting the size of the population that must be searched. As a secondary effect these groupings can also increase utility or reliability by creating tighter couplings between agents in the same congregation, typically by imposing higher penalties for decommitment or increasing information sharing among congregating peers. The downside to this strategy is that the limited set may be overly restrictive, and not contain the optimal agents one might interact with given infinite resources. So, in forming the congregation, one is trading off quality and flexibility for a reduction in time, complexity or cost. If an appropriate balance can be found, this will result in a net gain in utility.

This hypothesis is borne out in the experiments from an information economy domain (Brooks and Durfee, 2002). This work varied the number of congregations that agents were allowed to form. Since the population size was static, the average congregation size decreased as the number of congregations increased. The accumulated quality decreased proportionally because of less flexibility in agent interactions. However, these smaller congregations also incurred lower overhead, and thus had less cost. A median point was discovered in the space which produced maximum value.

6.2 Formation

Like coalition formation, congregation formation involves selecting or creating an appropriate group to join, and suffers from similar complexity problems as the agent population grows. Because congregations are more ideologically or capability driven, and there is usually no specific goal or task to unite them, one must first define how these groups may be differentiated. In (Brooks and Durfee, 2003) Brooks proposes using labels to address this problem. A label is a suitably descriptive tag assigned to each congregation which serves to both distinguish it from other groups and advertise the characteristics of its (desired) members. Assuming that agents have an ordered preference for such labels, the congregators' action is simply to move to the congregation for which it has the highest preference. The problem is then to create a number of logical points where agents may congregate and then decide upon the labels each congregation point will have; these labels help determine the makeup of the population which gathers there. Each agent was placed into one of several affinity groups, and a congregation is stable if and only if it contains only members of the same affinity group. Different numbers of labelers were then added which could attach labels to the congregation points. As with the congregators, the labelers were stable if and only if the congregation they provided the label to was homogeneous. The experimental and analytic results demonstrated that by increasing the number of labelers the system converged more quickly.

Brooks (Brooks and Durfee, 2002) presents a variation of this formation technique used in an information economy which also takes into account the costs associated with congregation size. In this scenario there are a set of buyers and sellers. Each buyer has an information preference, and each seller may choose what type of information to offer. The buyer's preference is soft – they have an optimal type, but are also willing to purchase related information, where similarity determines how much they are willing to pay. Instead of explicitly labeling congregation points, agents freely move through the system seeking groups that provide acceptable utility. The scenario is episodic, where during each episode agents elect to stay in place or randomly move to a new congregation. At the end of each episode an auction takes place from which buyers and sellers obtain their utility. The utility is based on the price of the goods bought and sold, combined with the costs incurred during the auction. This cost, divided uniformly among the congregation members, is proportional to the complexity of the auction, which is itself determined by the number of participants. Satisfied agents remain, while those which do not obtain enough utility move. This process results in an emergent population of congregations that trades off utility for computation time.

Griffiths' notion of a *clan* closely parallels the definition of a congregation (Griffiths, 2003). He presents a technique where clans are formed as part of a self-interested activity to increase local utility or decrease the probability of failure. If a motivating factor is exhibited by the agent, such as a desire to increase information gain or decrease commitment failure, clan formation may be initiated. Clan formation begins with the agent identifying how large a clan it wishes to create, which is based on the competing utility (in value added) and cost (in computational complexity) that grow in proportion to clan size. A trust value is then used to determine what agents it could invite, while the perceived capabilities or benefits of those individual agents are used to determine the appropriately sized subset that it will invite. In lieu of a negotiation process or explicit reward, invitation recipients determine if they will accept the invitation based first on their trust in the sender, and second on the perceived local gain they would receive by joining. The sender includes information about itself in the invitation as a sort of capability advertisement to facilitate this determination. If a sufficient number of agents agree, the clan is formed, otherwise the attempt is abandoned.

Although it does not strictly deal with congregating agents, Sen's work on reciprocal behavior (Sen, 1996) has some of the same characteristics. In this system, agents become more inclined to cooperate or assist another agent when it has a favorable history with that other agent. Specifically, agents track if others have cooperated with it in the past, or if it has cooperated with them, along with the approximate costs of those experiences. If an agent has a favorable balance of cooperation, it will be more inclined to give or receive assistance. The cooperation decision process is stochastic, enabling reciprocal relationships to be created or promoted even when a strictly positive balance does not exist. Weak groups may form between agents using this strategy who have complementary capabilities, which is similar to the notion of con-

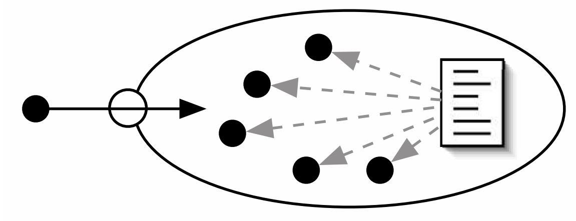

Figure 6: An agent society.

gregations we have presented. Because agents will more likely communicate with those that will help it, interactions can become implicitly confined within the group. These groupings are not formalized or well-defined, however, and communication is not necessarily restricted by the approximate boundaries that form. Sen showed that, among a group of self-interested agents operating in a package delivery domain, a population containing reciprocal agents outperformed a selfish population.

7 Societies

Drawing from our own experiences with biological societies, a *society* of agents intuitively brings to mind a long-lived, social construct. Unlike some other organizational paradigms, agent societies are inherently open systems. Agents of different stripes may come and go at will while the society persists, acting as an environment through which the participants meet and interact. A canonical example of this paradigm is the electronic marketplace (discussed in more detail in Section 9), consisting of buyers and sellers striving to maximize their individual utility (Wellman and Wurman, 1998; Artikis, 2003). A more ambitious example is the "agent world", a permanent operating environment for agents that in some ways parallels our own (Dellarocas and Klein, 1999; Willmott et al., 2001). Agents will have different goals, varied levels of rationality, and heterogeneous capabilities; the societal construct provides a common domain through which they can act and communicate. Societies are also more ephemeral constructs than others paradigms we have seen so far. They impose structure and order, but the specific arrangement of interactions can be quite flexible. Within the society, agents may be sub-organized into other organizations, or be completely unrelated.

A second distinguishing characteristic of societies is the set of constraints they impose on the behavior of the agents, commonly known as *social laws*, *norms* or *conventions*. This arrangement is shown abstractly in Figure 6, where the agents within the society have been provided with a set of specified norms. These are rules or guidelines by which agents must act, which provides a level of consistency of behavior and interface intended to facilitate coexistence. For example, it might constrain the type of protocol(s) agents can use to communicate, specify a currency by which they can transfer utility, or limit the behaviors the agent can exhibit in the environment. Penalties or sanctions may also exist to enforce these laws.

The set of laws embedded in a society must strike a balance among objectives (Fitoussi and Tennenholtz, 2000). It must be sufficiently flexible that goals are achievable, but not so much so that the beneficial constraints provided by the laws are lost. It must also be fair, such that the goals of one class of individuals are not incorrectly valued higher than those of another. These issues arise naturally in any structured, multiple participant system; Moses argues that most multi-agent systems have some form of social laws in place, if only implicitly (Moses and Tennenholtz, 1995).

7.1 Characteristics

In (Shoham and Tennenholtz, 1995), Shoham presents a grid world where robots must move from one location to another in accordance with a set of dynamically arriving tasks. Conflicts can arise when two or more agents attempt to occupy the same location at the same time along their chosen paths. They argue that a centralized solution is untenable, because of the potentially large number of interactions that must be continuously reasoned over in the heterogeneous population. Neither is a fully decentralized solution appropriate, because of the number of negotiation events that would need to take place at each time step. This motivates the need for "traffic laws", a type of social law which does not eliminate such interactions, but should minimize the need for them. The traffic laws in this research are computed offline, and constrain the robots' movement patterns in such a way that collisions do not occur, and destinations are reachable within a bounded amount of time. Vehicular traffic laws serve the same purpose in human societies. When driving a car there is no central authority which determines when and where we should go, and neither is there a free-for-all on the roads where one must talk to every other driver before proceeding. The challenge then is to design a set of laws that minimizes conflicts and encourages efficient solutions.

Although social laws were used to provide efficiency benefits in the work above, the purpose of an agent society is not always as quantitatively-driven as other organizational constructs. Indeed, most research on agent societies is more concerned with how the concepts they embody can be used to facilitate the construction of large-scale, open agent systems in general. For example, Moses (Moses and Tennenholtz, 1995) argues that social laws can provide a formal structure upon which more complex inter-agent behaviors can be built. By limiting and enforcing these restrictions, agents can make simplifying assumptions about the behavior of other agents, which can make interaction and coordination more tractable.

In additional to formalizing normative behaviors, mechanisms may also be established to ensure or encourage that such laws are respected. One approach accomplishes this through explicit representations of reputation or trust (Mui et al., 2002; S. D. Ramchurn and Jennings, 2004; Sabater and Sierra, 2002). An agent's behavior and interactions are observed by its peers and evaluated in the context of the norms it has agreed to. Deviation from those norms will result in a worsening reputation. This decreased reputation can in turn affect the utility the agent obtains, through increased decommitment penalties or competition from more reputable peers. In a rational agent this will serve as a deterrent to violating conventions. A different, but complementary approach instantiates and enforces social laws using social institutions provided in the environment (Dellarocas and Klein, 1999; Colombetti et al., 2004). Agents are expected to formalize their interactions using contracts, which are independently verified by these institutions, thereby relocating some of the traditionally agent-centric complexity into a service available to the population as a whole. This reduces the burden placed on agent designers, and provides a mechanism where systemic (non-localized or long-term) failures may be detected more readily. This more rigorous enforcement of social laws also helps address the problem of unreliable, dishonest or malicious agents operating in the open environment.

Huhns (Huhns and Stephens, 1999) provides similar motivation for common communication languages, shared or interoperable ontologies and coordination and negotiation protocols, all of which may be specified as part of the society's structure. These beliefs can be supported by our own experiences in real life. It should be clear that complex human societies are founded upon the ability to interact with one another. Mutually understood and respected norms simplify many aspects of day-to-day existence. These principles can be used to the same effect in agent societies.

7.2 Formation

There are two aspects to the society formation problem. The first is to define the roles, protocols and social laws which form the foundation of the society. Given such a definition, the second problem is to implement the more literal formation of the society, by determining how agents may join and leave it.

If the society is to be an open and flexible system, its structure must be formally encoded so that potential members may analyze it and determine compatibility. This description can be as simple as a set of common interfaces that must be implemented, or a complex description of permissible roles, high-level objectives and social laws. Dignum (Dignum, 2003; Dignum et al., 2002) presents a three-part framework, consisting of organizational, social and interaction models. The organizational model defines the roles, norms, interactions and communication frameworks that are available in the environment. The social model, instantiated at run-time, defines which roles agents have taken on. The interaction model, also created at run-time, encodes the interactions between agents that have been agreed-upon, including the potential reward and penalties. The latter two models are supported by contracts between the relevant entities. This formalism is similar to that proposed by Artikis (Artikis, 2003), which provides additional details describing operators that can be used to encode social laws, roles and normative relations. Because the society is intended to be open, these structures do not involve the internal implementation of agents, but describe only the intended or expected externally observable characteristics of the participants and environment.

Assuming it is possible to encode the social laws in a way that makes them intelligible to agents, one still faces the challenge of determining what conventions should be enacted. Fitoussi (Fitoussi and Tennenholtz, 2000) presents a notion of minimal social laws, where he argues that one should choose the smallest and simplest set of norms that address the needs of the society. This is consistent with the tradeoff between flexibility and complexity mentioned above. Work has also been done exploring the dynamic emergence of norms, for when social laws cannot specified off-line or if there is a desire for the corpus to be responsive to changing conditions (Axelrod, 1986; Hewitt, 1986). Walker and Wooldridge (Walker and Wooldridge, 1995) propose and evaluate a number of ways that a group of agents can reach norm consensus based on locally available information.

Dellarocas defines the act of an agent entering a society to be the *socialization* process (Dellarocas and Klein, 1999). In that work, they suggest this can be accomplished through an explicit negotiation process between the agent and a representative of the society, as shown in the left side of Figure 6. This exchange results in a *social contract*, or an explicit agreement made between the agent and the society indicating the conditions under which the agent may join that society. This allows the possibility of capable agents dynamically learning, and potentially negotiating over, the rules it must abide by in that society. López y López (y López et al., 2004), present a framework in which the facilities for norm reasoning needed to support these behaviors can take place. A similar view is taken by Glaser in (Glaser and Morignot, 1998), with the additional stipulation that the joining agent must increase the utility of the society. This naturally extends to multi-society environments, where an agent's skills and goals define how good a fit it is with a particular society. Some of the challenges associated with operating in multi-society environments seem to be comparable, though larger in scale, to those encountered during coalition or congregation formation.

Because of their inherent flexibility, a great deal of additional complexity may be associated with social organizations. Sophisticated legal systems, communication bridges, ontologies, exception handling services, directories may all be part of the society model (Dellarocas and Klein, 1999; Dignum, 2003; Klein et al., 2003). Some or all of these may be directly instantiated by trusted agents taking on so-called facilitation roles (differentiated from the operational roles taken on by worker agents). Of course, agents acting in the society must have a certain level of sophistication to know how and when to use such services. An interesting almost-paradox exists in this relationship. Although the society exists in part to reduce the complexity burden imposed on the participants, the participants must raise their level of complexity to take advantage of these benefits. In the case where interactions with some or all social services are mandatory (e.g. legal or arbitration services), this additional complexity is similarly unavoidable and can act as a barrier to entry.

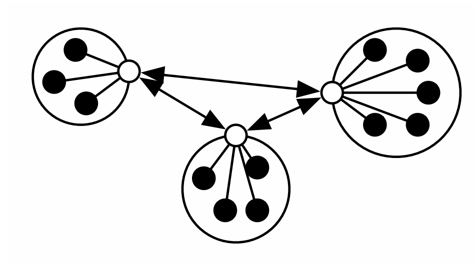

Figure 7: An agent federation.

8 Federations

Agent *federations*, or *federated systems*, come in many different varieties. All share the common characteristic of a group of agents which have ceded some amount of autonomy to a single delegate which represents the group (Wiederhold et al., 1992; Genesereth, 1997). This organizational style is modeled on the governmental system of the same name, where regional provinces retain some amount of local autonomy while operating under a single central government. The delegate is a distinguished agent member of the group, sometimes called a facilitator, mediator or broker (Sycara et al., 1997; Hayden et al., 1999). Group members interact only with this agent, which acts as an intermediary between the group and the outside world, as shown in Figure 7. In that figure each grouping is a federate, and the white agent situated at the edge of each federate is the delegated intermediary. Typically, the intermediate accepts skill and need descriptions from the local agents, which it uses to match with requests from intermediaries representing other groups. In this way the group is provided with a single, consistent interface. This level of indirection is similar to that seen in holons, and provides some of the same benefits.

8.1 Characteristics

The capabilities provided by the intermediary are what differentiate a federation from other organizational types. The intermediary functions on one hand by receiving potentially undirected messages from its group members. These may include skill descriptions, task requirements, status information, application-level data and the like. These will typically be communicated using some general, declarative communication language which the facilitator understands (Genesereth, 1997). Outside of the group, the intermediary sends and receives information with the intermediaries of other groups. This could include task requests, capability notifications and application-level data routed as part of a previously created commitment. Implicit in this arrangement is that, while the intermediary must be able to interact with both its local federation members and with other intermediaries, individual normal agents do not require a common language as they never directly interact. This makes this arrangement particularly useful for integrating legacy or an otherwise heterogeneous group of agents (Genesereth, 1997; Shen and Norrie, 1998).

The intermediary itself can function in many different capacities. It may act as a translator, perform task allocation, or monitor progress, among other things. An intermediary which accepts task requests and allocates those tasks among its members is known as a broker or a facilitator. As part of the allocation, the broker may decompose the problem into more manageable subtasks. This allows agents to take advantage of all the capabilities of the (potentially changing) federation, without requiring knowledge of which agents perform a task or how they go about doing it. This reduces the complexity and messaging burden of the client, but also has the potential of making the broker itself a bottleneck (Hayden et al., 1999) (a possibility common to all intermediaries). An intermediary acting as go-between among agents is known variously as a translator, embassy or mediator depending on its specific characteristics. Embassy agents provide a layer of security for members of their federation, by having the ability to deny communication requests. Mediator agents store representations of all related parties, reducing their individual complexity by providing a layer of abstraction. This capacity can be further exploited to arbitrate conflicts (Mailler and Lesser, 2004). Intermediaries which provide the ability to track the state of one or more of its participants are known as monitors. For example, result information can be automatically propagated to interested parties. Of course, one or more of these roles may be combined into a single intermediary which offers several types of services.

8.2 Formation

Singh and Genesereth in (Singh et al., 1995; Genesereth, 1997) describe how a general federated system would work. All agents are expected to communicate using an Agent Communication Language (or ACL, a somewhat-generic term used by many researchers to describe their agents' communication protocol), which in this work is a combination of the first-order predicate calculus KIF with the KQML agent messaging language. Knowledge and statements sent between agents are encoded as KIF statements, which are then wrapped in KQML to provide a standard mechanism for specifying the sender, receiver, intent, and so forth. This provides a common language and set of behavioral constraints that will allow the various agents to interact. Not all agents must implement the entire class of concepts in the ACL, but the aspects they do use must be correct with respect to the ACL's specification. In addition, although they speak the same language, not all agents must use the same vocabulary to describe a particular situation, although to interact there must be an intermediary capable of translating the vocabularies. The system is initialized with a set of intermediaries called facilitators, which serve many of the roles outlined above, notably brokering. Agents connecting to the system start by sending their capabilities to the local facilitator. Implicit in this communication is the notion that the agent is willing to use those capabilities in service of requests posed by the facilitator. Needs are similarly routed to the facilitator, which then attempts to find other facilitators that can service that need. Each facilitator provides a yellow pages function which supports this search. Khedro's Facilitators (Khedro and Genesereth, 1995) and the jointly developed PACT project (Cutkosky et al., 1997) have pro-

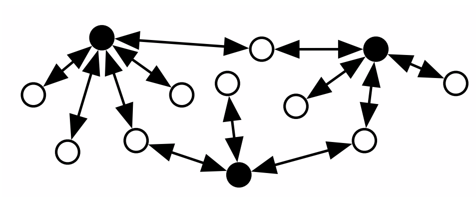

Figure 8: A multi-agent marketplace.

duced very similar systems that also use a common ACL and a community of intermediaries to produce a robust and dynamic task decomposition and allocation scheme among a group of heterogeneous participants.

The MetaMorph I (Maturana et al., 1999) and II (Shen and Norrie, 1998) architectures described by Maturana and Shen demonstrate a federated agent system for use in intelligent manufacturing. In this domain, agents are used to drive aspects of product design and manufacturing, contending with heterogeneous resources, dynamically changing conditions, and hard and soft constraints on behavior. MetaMorph's name is derived from the fact that the system can continuously change shape, adapting to new conditions as they are perceived. This is accomplished in part through the use of intermediaries called mediators, which are responsible for brokering, recruiting and conflict resolution services. The recruiting service is similar to brokering, but is differentiated by the fact that the intermediary can remove itself from the relationship once the partners have been discovered. This weaker form of federation provides efficiency gains at the cost of less flexibility, both due to the loss of the layer of abstraction that exists in the brokered approach. The federations themselves are dynamically created in response to new task arrivals or requests from other groups using a contract net (Smith, 1980) approach, or are statically created from agents in a common subsystem (e.g. tools, workers, etc.).

9 Markets

In a *market*-based organization, or *marketplace* as shown in Figure 8, buying agents (shown in white) may request or place bids for a common set of items, such as shared resources, tasks, services or goods. Agents may also supply items to the market to be sold. Sellers (shown in black), or sometimes designated third parties called *auctioneers*, are responsible for processing bids and determining the winner. This arrangement creates a producer-consumer system that can closely model and greatly facilitate real-world market economies (Wellman, 2004). These latter systems fall into the more general category of *agent-mediated electronic commerce* (Guttman et al., 1998). Because of this similarity, a wealth of research results from human economics and business can be brought to bear on agent-based markets, creating a solid theoretical and practical foundation for creating such organizations (Wellman, 1993; Wellman and Wurman, 1998; Corkill and Lander, 1998).

Markets are similar to federated systems in that a distinguished individual or group of individuals is responsible for coordinating the activities of a number of other participants. Unlike a federation, market participants are typically competitive. In addition, participants do not cede operational authority to those distinguished individuals, although they do trust the entities managing the market and abide by decisions they make. It is also common for markets to operate as open systems (Wellman, 2004), allowing any agent to take part so long as it respects the system's specified rules and interface. As such, they share some of the benefits and drawbacks of societies.

When using the terms "buyer" and "seller", one may implicitly assume that an artifact will eventually be transferred in exchange for some form of compensation (Chavez and Maes, 1996; Tsvetovatyy et al., 1997). Although this paradigm is common, it is not always the case, and market-based organizations have been used in various projects to accomplish less obvious goals. For example, Wellman (Wellman et al., 1998) proposes using a market-based approach to perform decentralized factory scheduling. In this work, each factory job is associated with a duration, deadline and value. The factory itself, acting as the seller, has a reserve price associated with the time slots it has available. Agents bid on a set of slots that have sufficient total time to satisfy the job duration and do not exceed the deadline, using the job value as a maximum bid price. Market forces will cause agents to seek out the most cost-effective time slots, while higher-valued jobs will naturally take precedence over lower ones. This should lead to an efficient allocation of (time) resources, while maximizing the factory's overall utility. Bussman (Bussmann and Schild, 2000) has developed an auction-based manufacturing control system with a similar purpose, where agents are used to represent workpieces, transportation conveyors and machines. In this work, machines bid for the right to work on workpieces, which act as sellers, by relating an expected time to completion. When a machine's bid is accepted, a series of additional negotiations between the workpiece and the conveyors move the piece to the appropriate location. Yet another example is the Mariposa distributed database system (Stonebraker et al., 1996), which uses market-based techniques to optimize query processing. Individual nodes buy and sell fragments of information. Queries inserted into the system are associated with a biding profile, indicating how much the user is willing to pay. A brokering process takes the query and requests bids from relevant nodes. who then submit bids in an effort to win the rights to process the query.

More generally, Wellman proposes the notion of *market-oriented programming* (Wellman, 1993), which uses the marketplace paradigm as a general programming methodology that can efficiently address multi-commodity flow and resource allocation problems. His WALRAS framework that implements this concept has been used to create solutions for transportation logistics, product design and distributed information services. Many other marketplace frameworks have also been developed for general use (Chavez and Maes, 1996; Rodriguez et al., 1997; Collins et al., 1998; Collins and Lee, 1998; Cuni et al., 2004); Kurbel and Loutchko provide a comparative analysis of structure and function (Kurbel and Loutchko, 2003).

9.1 Characteristics

Markets excel at the processes of allocation and pricing (Wellman and Wurman, 1998). If agents bid correctly (i.e. make truthful bids according to their perceived utility gain if they win), the centralized arbitration provided by the auctioneer can result in an effective allocation of goods. The Kasbah system (Chavez and Maes, 1996) is an example of an agent-based marketplace that demonstrates many of the typical characteristics of this type of organization. Agents in Kasbah are segregated into two categories: buyers and sellers. Both types indicate the type of object they are interested in (buying or selling) with a feature vector, along with a desired price, a threshold price (lower or upper bound), and a negotiation strategy that controls how their offered price changes over time. A sale occurs when a seller's price matches what a buyer is willing to pay. The objects being sold in this system represent the targets of the allocation process, and the price is determined dynamically according to supply and demand. The mechanism that is employed in Kasbah corresponds to an intuitively fair way to allocate among competitors, at least from a self-interested point of view: all agents gradually compromise, and the agent willing to meet the seller's price first wins.

The behaviors embodied in a marketplace, namely the existence of buyers and sellers, a potential multitude of goods, and competition among participants, make such organizations intrinsically linked with the properties of *auctions*. Kasbah is an example of a *two-sided* auction, because both sides compromise. If one of the two parties maintained a fixed price, it would be *one-sided* auction. Many other types of auctions exist to service the different needs of different communities, each with their own characteristics (Wurman et al., 2001; Kurbel and Loutchko, 2003). For example, in a *combinatorial* auction, participants bid on collections of goods, rather than single objects. In an *reverse* auction, sellers bid rather than buyers. In *sealed-bid* auctions, the participants do not see competing bids while the auction is is progress. In *continuous* auctions, a pool of items exist, exchanges occur as soon as two compatible bids are made, and the bidding process continues uninterrupted. The particular type of auction which is employed dictates the manner in which the participants interact. Much of the complexity involved in designing an effective market and marketplace agent revolves around understanding the subtleties of the auction's characteristics, and crafting an appropriate strategy based on that knowledge.

There are two drawbacks to market-based organizations. The first is the potential complexity required to both reason about the bidding process and determine the auction's outcome. The former computation may require a detailed approximation of competitors' beliefs, a practice known as *counterspeculation*, especially in single-shot or sealed bid auctions (Tsvetovatyy et al., 1997). The latter computation, also known as *clearing* the the trade, can be particularly difficult in the case of combinatorial auctions. This is known to be a NP-complete problem (Sandholm, 2002), although solutions have been devised that have good performance in practice (Sandholm, 2005). The second is security; in addition to the practical network-related security issues inherent in any open system, one must also be able to verify the validity of the auction approach itself. For example, the bidding strategy used in the Kasbah system is vulnerable to a form of cheating known as *collusion*. If two or more bidders in the system agree to reduce their rate of compromise, they have a chance to artificially lower the final sale price. It is also important that the bidding process does not reveal information about the participants. For example, if a seller could determine the threshold prices of some of its buyers, it could simply wait until the maximum such price is reached, thereby artificially increasing the sale price. Some of these issues can be resolved by selecting an appropriate auction type. The Vickrey auction's structure (Vickrey, 1961), where the highest bidder wins but pays the second highest bid price, promotes truthful bidding and discourages counterspeculation. Enforcing anonymity and secure communication channels can also help avoid many common pitfalls.

9.2 Formation

As is the case of many open systems, marketplaces are frequently static, pre-existing entities that do not require a formal creation process beyond starting the actual market process (if any) and allowing agents to connect. The well-known Trading Agent Competition market (Wellman and Wurman, 1999) operates in such a fashion, albeit for a limited amount of time. They may have certain barriers to entry, such as respecting a defined programming interface, implementing a particular transaction language, and respecting the rules of the market's auction type. These entry conditions are similar to those discussed earlier in the context of societies, although there is generally no formal negotiation or socialization process involved. Wellman (Wellman, 2004) outlines a number of other practical characteristics that should be exhibited for a marketplace to be successful. They must maintain *temporal integrity*, meaning that the outcome of an auction depends on the arrival sequence of bids, and is independent of any delays internal to the market itself. Transactions performed by the market must be *atomic*, that is, they have no effect if they fail or are canceled prior to completion. As noted above, they also require attention to security risks, so that participant information is adequately protected and the auction process itself is kept safe from conventional attacks, particularly if there is an actual exchange of goods, information or currency in the market. Markets may also incorporate product discovery services, banking services, brokering middle-agents and negotiation support, to reduce the burden placed on the participants (Tsvetovatyy et al., 1997; Guttman et al., 1998).

Other works have explored dynamic formation of markets. As mentioned in Section 6, Brooks has used the notion of congregations to dynamically form markets within a group of agents (Brooks and Durfee, 2002). Recall that congregations are groups of agents which have banded together because of some common long-term interest or goal. In this work, that long term goal is the cost-effective exchange of goods or services. In a large population, it can be difficult to directly find suitable trading partners, and expensive to contact or broadcast to all possible partners. A suitably formed congregation serves to limit the scope of this search or broadcast, which in turn facilitates the marketplace creation.

A relatively new concept being exploited in both human (Mowshowitz, 1997) and agent (Ahuja and Carley, 1999; Foster et al., 2004; Cardoso and Oliveira, 2004) organization research is the *virtual organization* (VO). A virtual organization is one that has a fixed purpose (e.g., to provide a set of services) but a potentially transient shape and membership. The key characteristics of a VO are that they are formed by the grouping and collaboration of existing entities, and there is a separation between form and function that precludes the need to rigidly define how a behavior will take place. This provides flexibility in how a particular goal is satisfied, by allowing the system to adapt the set of participants to meet resource availability and service demand. The concept is similar to the coalition and congregation paradigms discussed earlier, and has many of the same benefits as a federation, although a virtual organization can generally be thought of as an entity in and of itself more so than an empty coalition or congregation.

The CONOISE project has explored the dynamic creation of virtual organizations within a larger marketplace environment (Norman et al., 2003). In this context, the creation of a VO can be thought of as the creation of a new market entity (buyer or seller) from a group of exist-

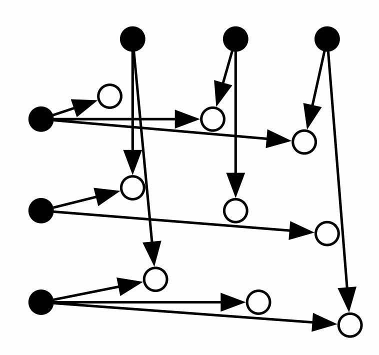

Figure 9: A matrix organization.

ing participants. This can give those participants greater leverage, efficiency or reliability as they combine their producing or consuming power. The members of a VO may remain distinct when outside of the marketplace, but within the market they act as a single unit. For example, two producers might combine to offer a new joint product. Two consumers might combine to obtain greater buying power. In responding to bids, a VO will then be able to offer the union of services or goods over all its members. VOs may also split when the relationship is no longer beneficial or if levels of trust or reputation have been sufficiently degraded. In all cases, the shape of the market is affected as these changes are made, and thus the market as a whole will evolve over time based on the needs and capabilities of the participants, and the corresponding consolidation decisions they make.

10 Matrix Organizations

We have seen that the strict hierarchical organization method is based on a tree-like structure of control. Agents or agent teams report to a single manager, which provides the agents with goals, direction and feedback. *Matrix organizations* relax the one-agent, one-manager restriction, by permitting many managers or peers to influence the activities of an agent. This forms a mixed-initiative environment, where successful agents reason about the effects their local actions can have on multiple entities. This is in some sense a closer approximation to how humans exist. A person may receive guidance or pressures from their manager, co-workers, spouse, children, colleagues, etc. Even in a purely business setting one might have to report to an immediate supervisor, project managers, vendors, and peers at cooperating businesses. Interrelationships can come from many directions, each with its own objectives, relative importance and pertinent characteristics (Wagner and Lesser, 1999).

The term matrix organization comes from a grid based view of the participants. One can place managers (black) around a group of "worker" agents (white), and use a directed edge to indicate authority, as in Figure 9. Alternately, agents are the rows and managers the columns (these sets may overlap), and a check is used to denote where an authority relationship exists. Like the hierarchy's tree, the matrix provides a graphical way to depict which managers can influence the activities of each agent.

10.1 Characteristics

Matrix organizations provide the ability to explicitly specify how the behaviors of an agent or agent group may be influenced by multiple lines of authority (Decker et al., 1995). In this way, the agent's capabilities may be shared, and the agent's behaviors (hopefully) influenced so as to benefit all. This is particularly important if the agents themselves are viewed as a functional, limited resources. For example, if a particular skill is needed by two separate tasks, the agent can be used to address both, provided it has sufficient computational power. In the case where the agent has multiple ways of performing a task, it can also choose the method which best satisfies its peers.

This sharing come as a price, however, because the shared agent becomes a potential point of contention. If its managers disagree, the agent's actions may become dysfunctional as it is pulled in too many directions at once (Schwaninger, 2000; Romelaer, 2002). To operate effectively, the agent must have a commitment ranking mechanism and sufficient autonomy to resolve local conflicts, or the ability to promote conflicts to a higher level where they may be resolved (Mailler et al., 2003). Wagner's *motivational quantities* framework (Wagner and Lesser, 1999) is one approach that addresses this problem. In that work, task valuation is performed by combining both the local intrinsic worth of the task with the perceived or specified worth that task will have on other entities. This valuation is quantified through the expected production and consumption of different motivational quantities (MQs), which act as a virtual resource or medium of exchange. The preference for particular MQs is specified with a set of utility curves that together determine the agent's overall utility. By coupling the production of different types of MQs with the tasks associated with different managers, the framework is able to capture the quantitative motivation behind a particular course of action. This explicitly represents the type and state of the relationships the agent has with those managers, which can enable it to correctly balance its behavior in a matrix organization.

10.2 Formation

Decker (Decker et al., 1995) describes the MACRON organizational architecture, in which agents form a matrix organization. The domain for their system is cooperative information gathering, where multiple agents search for relevant data in response to a user's query. Individual agents are separated into predefined functional groups that contain agents able to access a particular type of information. These groups are under the control of a functional manager, who assigns agents to query tasks as they arrive. User query agents generate those query tasks, and therefore use the functional managers to dynamically select agents to satisfy their own goals. Individual gathering agents report to two agents: a static functional manager, and a query manager which changes depending on the user's actions. This has the effect of assigning the minimal needed set of agents to the query, increasing efficiency when compared to a system employing a set of static teams where particular team members might go unused, depending on the query characteristics. At the same time, this approach uses fewer resources than one lacking functional groups, which would have to search through all available agents for each query.

In (Horling et al., 2003), Horling and Mailler describe a distributed sensor network application where a matrix organization is used to address a resource allocation problem. In this case, the sensors themselves were limited resources, since their heterogeneous locations and orientations made each one unique. The tracking process for each target was controlled by a different track manager, which was responsible for discovering and coordinating with the sensors needed to track its target. When multiple targets came in close proximity to the same sensor, a matrix organization is dynamically formed as the relevant managers interact with that sensor. At the same time, that sensor may have previously been given tasks by a regional manager responsible for detecting new targets. The result is an individual which may be under contention by three or more managers, and which must then decide how best to meet those demands. This was done using a combination of a predefined ranking scheme (tracking has higher priority than scanning for new targets), local autonomy (round robin scheduling) and conflict elevation (track managers negotiate directly once aware of the conflict).

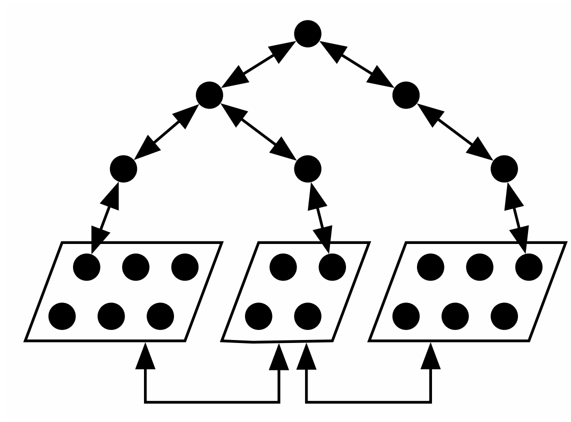

Figure 10: A compound organization.

11 Compound Organizations

Not all organizational structures fit neatly into a particular category, and some architectures may include characteristics of several different styles. A system may have one organization for control, another for data flow, a third for discovery, and so on. For example, Durfee's PGP (Durfee and Lesser, 1991) incorporates one organization for interpretation, and another separate structuring of the same agents to manage coordination problems. Compound organizations can be overlapped, operating as virtual peers at the same conceptual level, or be nested, so that some subset of agents in a group are organized in a potentially different way within the larger context. A sample such organization is shown in Figure 10, which combines a hierarchy with a set of coalitions. As with singular organizations, they may be created or adapted over time, or they may be instantiated as part of a transient form while a population shifts between organizational styles. Ideally, these compound architectures can use the most effective structure for the particular goal at hand, without limiting options that might be used elsewhere in the system. The tradeoff in this situation is usually one of complexity. Because an individual agent might take on different roles in response to different organizational demands, the agent itself must have sufficient sophistication to act efficiently and asynchronously in all those roles.

Some of the organizational paradigms which have been discussed so far are more amenable to coexistence than others. In much of the teamwork research, for example, a loose hierarchy of control was created among the agents after the team had formed (Tambe, 1997; Tidhar et al., 1996). Hierarchical structures for interpreting and consolidating raw data are also a popular mechanism for handling scale that can augment a preexisting or lower-level structure (Yadgar et al., 2003). Societies frequently have an internal organizational structure within the larger context defined by the social laws and norms (Dellarocas and Klein, 1999; Dignum, 2003). In other cases, researchers have exploited the characteristics of one type of organization to create another. Congregations, for example, have been used to facilitate the dynamic formation of markets (Brooks and Durfee, 2002), while both markets (Lerman and Shehory, 2000) and hierarchies (Abdallah and Lesser, 2004) have been used to efficiently create coalitions. Societies can also be viewed as a common "pool" of agents, from which a range of other organizations can be constituted. In this type of compound organization, the society may exist in support of other, more dynamic structures created to address particular tasks (Sichman and Demazeau, 2001). This begins to touch on the notion of organization longevity, which will be addressed in Section 13.

11.1 Characteristics

The positive and negative characteristics of a compound organization are derived primarily from its constituent parts. However, the interplay between organizations can lead to unexpected consequences. For example, if the distinguished intermediary in a federated system plays a key role in a separate overlay organization, it may be unable to fulfill both roles adequately. Similar to a matrix organization, agents may be faced with conditions where it is not clear which of two competing objectives it should satisfy (Romelaer, 2002). Conversely, its knowledge of the requirements of both organizations may enable it to make more globally effective decisions. The possible interactions and formation strategies among arbitrary coexisting organizations are difficult to characterize in a general manner; instead we will proceed with a discussion of example systems employing this technique.

11.2 Example Compound Organizations

The distributed sensor network solution described by Horling and Mailler (Horling et al., 2003) uses several different overlapping organizational techniques. Agents are first partitioned into federations, called sectors, where membership is based on their geographic proximity. A distinguished member of each group is given the role of sector manager, who provides a form of recruiting service to other agents in the environment. This recruiting service supports the activities of track managers, who must discover and use the appropriate sensors as part of their tracking task. In forming the federations, the search time is reduced because only a subset of the population (the sector managers) needs to be interacted with, and communication requirements requirements are reduced because only the necessary subset of sensors will be returned. As discussed in Section 10, both the sector and track managers provide tasks to individual sensors, forming a matrix organization in the process. This arrangement facilitates resource sharing by allowing the sensors to guide their local activities based on the needs of potentially several interested parties, but can also lead to conflicts caused by over-demand. Because the sensor is a finite resource, a cloning technique like the one discussed in Section 2 cannot be used to address the conflict. Instead, a loose peer-to-peer relationship between track managers allows them to negotiate directly, alleviating the conflict through demand relaxation or by using alternate sensors. This resource allocation scheme employs a second, weaker form of federation through its use of mediators (Mailler and Lesser, 2004). The conflicts, which may be potentially multi-linked and far-reaching, are partially centralized by a mediator agent which acts on the part of the relevant agents to find a suitable solution. In (Horling et al., 2004) the quantitative effects of these interactions are demonstrated through a set of experiments that vary the shape of the organizational structure.

Yadgar (Yadgar et al., 2003) describes a different approach in a distributed sensor environment. Groups of geographically-related sensors are first formed into sampler groups, which are essentially federations with a single agent called the sampler group leader acting as the intermediary. These groups then form the lowest level of a data aggregation hierarchy that exists above them. This arrangement is similar to the example organization shown in Figure 10. The sampler group leader collects raw data from the members of its group, and passes the data to its parent agent in the hierarchy, known as a zone leader. It is this zone leader's responsibility to interpret the sensor data to the best of its ability, by building motion equations and combining data perceived to be from the same target. This more abstract view is then passed to the next level of the hierarchy, where the process repeats. This will eventually terminate at the apex agent which should be able to reconstruct a global view from the abstract pieces it receives. The hierarchy itself is strict, and communication is only permitted between connected agents, which reduces the level of sophisticated needed by the agents. The experimental results showed that this solution could scale to thousands of sensors and targets. The tradeoff they discovered was that shorter hierarchies produced more accurate results, because the fragmentation of the area was minimized, which in turn reduced the number of fusion processes data must survive before it is incorporated. Conversely, taller hierarchies dramatically reduced the computational load placed on any one agent, because the area each agent was responsible for became relatively small. By weighing these characteristics against the domain requirements one can select an appropriate structure to use.

12 Other Organizational Topics

In this survey we have focused entirely on particular organizational paradigms. However, there are a number of other topics related to organizational design which we will not cover in detail, but are sufficiently important to warrant mention. These are outlined below:

1. *Global Organizational Representation* Implicit in the concept of an intentional organizational design is an explicit representation of its structure. This is of use to designers, as a means of specification and exploration, and to the agents themselves, as a template and diagnostic tool. A number of general modeling representations have been proposed, notably by Fox (Fox et al., 1998), Tambe (Tambe et al., 1999), Hübner (Hübner et al., 2002), Pattison (Pattison et al., 1987), Dignum (Dignum, 2003), Sims (Sims et al., 2004), Horling (Horling and Lesser, 2005) and Vázquez-Salceda (Vázquez-Salceda et al., 2004).

2. *Local Organizational Representation* The organization's global view is not always the most appropriate vehicle to guide agents' behaviors. It can be too coarse in granularity, too qualitative or simply too large to be of practical use. Agents require a well-defined, quantitative mechanism that can be used to select appropriate local actions while respecting global organizational specifications. This process was originally described as *local elaboration* by March and Simon (March and Simon, 1958), where the activities performed by an agent are first constrained by its position in the organization, and then selected using local information and capabilities. The social consciousness model suggested by Glass and Grosz (Glass and Grosz, 2000), Decker's TÆMS language (Decker and Lesser, 1993b), Shoham's social laws (Shoham and Tennenholtz, 1995), and Wagner's MQ framework (Wagner and Lesser, 1999) provide ways to accomplish this.

3. *Organizational Performance* Other researchers have taken a different approach by creating formal analytic or statistical models that focus on the activities or behaviors of the organization, rather than representing the organization as a whole (Malone and Smith, 1988; Decker and Lesser, 1993a; Montgomery and Durfee, 1993; So and Durfee, 1996; Lerman and Galstyan, 2001; Shen et al., 2004; Gnanasambandam et al., 2004; Horling and Lesser, 2005; Schmitt and Roedig, 2005). These typically more quantitative representations can provide insights into organizational performance that are largely absent from purely descriptive or logical

representations. A different approach is to use experimental or simulation studies, which can offer a more general-purpose approach to analyze organizational performance that may not be amenable to modeling (Lesser and Corkill, 1983; Lin and Carley, 1995; Sierra et al., 2004). The drawback to using empirical analysis is the time required to run such tests, which is usually much greater than that needed for analytic techniques. Conversely, analytic models may require simplifying assumptions to be tractable, or otherwise fail to take into account the complexity real-world behaviors. Parunak (Parunak et al., 1998) provides further discussion on the tradeoffs between these approaches. However they are obtained, such predictions can play a critical role in the search and evaluation process, by allowing the designer to directly compare alternative organizational strategies before implementing a design. This can provide the foundation for a more proscriptive organizational tool.

4. *Generative Paradigms* In each section, we have presented different ways in which organizations may be formed. We have not, however, presented a unified discussion of specific generative paradigms – a classification of the techniques that may be used to produce organizations. These may be broadly separated into at least three classes: scripted, controlled and emergent. The first includes organizations that are produced from statically predefined instructions, possibly from an external third party or during startup. The second includes those that are explicitly applied to a population by an individual or group of individuals in response to perceived conditions. The third captures techniques which have no central or global direction, but are instead self-directed or grown organically through the individual actions of agents. In practice, it may be difficult to clearly classify particular techniques. For example, congregations emerge from individual agent decisions using the technique described by Brooks (Brooks and Durfee, 2003). However, the fact that it uses heuristics intended to simulate a controlled decision, along with agents which provide labels to guide the formation, gives the appearance of a controlled process.

5. *Organizational Adaptation* Although we have briefly touched on adaptation previously, an organization's ability to adapt is a general concept that is critical in any dynamic environment. The organization must have the ability to detect and react to changes in a timely manner in realistic, open domains (Carley, 1998; Barber and Martin, 2001; Horling et al., 2001). Any organizational change which occurs at runtime will have associated costs. These costs may be observed in direct consumption of resources, such as bandwidth or processing power, or indirectly because of inefficiencies or opportunities missed while in an intermediate state. The ability to adapt an organization depends on first recognizing potential problems, evaluating the costs and benefits of candidate solutions, and then implementing the selected changes. Related to adaptation is the notion of social pathologies, which occur when an organization adapts inappropriately (Turner, 1993; Jensen and Lesser, 2002).

6. *Coordination and Negotiation* Many of the organizational styles that we have covered assume some that some sort of interaction or coordination will take place between agents. This is seen in the authority relationships of hierarchies, the joint intentions of teams, data routing protocols in federations, and negotiations of society members. The characteristics provided by these interactions are critical to the effective qualities of these paradigms. For example, aggregating nodes and managers in hierarchies and intermediaries in federations frequently take on responsibilities related to coordination, by assigning tasks or routing information in such a way that interrelationships among their subordinates can be avoided (Galbraith, 1977). Argumentative negotiation has been shown to be effective in resolving conflicts in team settings (Jung et al., 2001). The techniques that are used can heavily influence the interactions and behaviors exhibited by the group, ultimately affecting the performance of the organizational structure. Work by Prasad Prasad and Lesser (1999), Lesser (Lesser et al., 2004) and Toledo (Excelente-Toledo and Jennings, 2004) have also explored the dynamic selection of coordination strategies, which in this context can be considered a form of organizational adaptation.

7. *Autonomy* The manner in which an agent behaves, and in particular how its motivations are determined, is intimately related to its position within the organization. Agents may be externally directed, self-directed or some combination of the two (Lesser and Corkill, 1981). For example, we have seen that agents in hierarchies, federations and matrix organizations all generally have manager-supervisor relationships, implying that local actions are partially or completely decided by an external entity. Conversely, agents operating in markets are typically more autonomous, independently deciding how and when to bid. Like other characteristics, the level of autonomy can affect the performance of the system as a whole. Authoritarian structures can exploit centralization to make good decisions, while an organization of more autonomous entities offers better balance and parallelism. Because the needs and constraints exhibited by participants change over time, it can also be beneficial to dynamically adapt agents' levels of autonomy in response to changing events (Barber and Martin, 2001; Scerri et al., 2002; Zhang et al., 2003).

8. *Human Organizational Analogues* For much of the time that multi-agent organizations have been researched, attempts have been made to draw upon the large body of work that has been done on human organizations. The fields of sociology, anthropology, biology, economics, business management and formal organization theory (among others) contain a wealth of analytic and case study information describing how human organizations are structured and perform (Fox, 1981; Gasser, 2001). Although on the surface much of this work is intimately tied to the human experience, attempts to extract concepts and abstractions have met with some success.

9. *Diversity* Although role assignment clearly plays a critical role in an organizational specification, the notion of agent diversity is rarely treated as or reasoned about as a first-class characteristic. As with stock portfolios, animal populations and security techniques, diversity can play an important role in agent systems susceptible to failure. Enforcing agent diversity through heterogeneous roles, agent types or division of labor, can impart semantic and capability fault-tolerance on the system as a whole (Corkill and Lesser, 1983; Reed and Lesser, 1980; Corkill and Lander, 1998; Lybäck, 1999). Diversity can be embedded in the organizational design to encourage such characteristics.

13 Discussion

In this article we have presented a number of methods by which a multi-agent system could be organized. A brief comparison of the potential benefits and drawbacks of each strategy is summarized in Figure 11. A more complete depiction of the range of relevant organizational characteristics in general has been compiled by Carley and Gasser (Carley and Gasser, 1999), while Malone and Smith (Malone and Smith, 1988) provide a focused comparison of the characteristics of hierarchy and marketplace designs. It should be clear from this discussion that no single approach is necessarily better than all others in all situations. The selection made by a designer should be dictated by the needs imposed by the system's goals, the resources at hand, and the environment in which the participants will exist. That said, if one looks at the depth of available research and how frequently their concepts have been applied, it is the case that hierarchical, team-centric, coalition-based organizations and marketplaces have proved to be most popular among multi-agent researchers. These four paradigms seem to offer the most in terms of flexibility, ease of implementation and their innate ability to produce demonstrable, positive effects. Hierarchies are effective at addressing issues of scale, particularly if the domain can be easily decomposed along some dimension. Teamwork can be critical when working on large-grained tasks that require the coordinated capabilities of more than one agent. Coalitions allow agents to take advantage of economies of scale, without necessarily ceding authority to other agents. Markets take advantage of competition and risk to decide allocation problems in a fair, utility-centric manner. We also feel that if the broad vision of an agent-connected or agent-facilitated world that many proponents of multi-agent technology describe is to be realized, many of the characteristics of the agent society paradigm must be incorporated (Gasser, 2001).

A popular approach not mentioned thus far is the

Paradigm	Key Characteristic	Benefits	Drawbacks
Hierarchy	Decomposition	Maps to many common domains; handles scale well	Potentially brittle; can lead to bottlenecks or delays
Holarchy	Decomposition with autonomy	Exploit autonomy of functional units	Must organize holons; lack of predictable performance
Coalition	Dynamic, goal-directed	Exploit strength in numbers	Short term benefits may not outweigh organization construction costs
Team	Group level cohesion	Address larger grained problems; task-centric	Increased communication
Congregation	Long-lived, utility-directed	Facilitates agent discovery	Sets may be overly restrictive
Society	Open system	Public services; well defined conventions	Potentially complex, agents may require additional society-related capabilities
Federation	Middle-agents	Matchmaking, brokering, translation services; facilitates dynamic agent pool	Intermediaries become bottlenecks
Market	Competition through pricing	Good at allocation; increased utility through centralization; increased fairness through bidding	Potential for collusion, malicious behavior; allocation decision complexity can be high
Matrix	Multiple managers	Resource sharing; multiply-influenced agents	Potential for conflicts; need for increased agent sophistication
Compound	Concurrent organizations	Exploit benefits of several organizational styles	Increased sophistication; drawbacks of several organizational styles

Figure 11: Comparing the qualities of various organization paradigms.

(sparsely) connected graph structure, sometimes called a *network organization* or *adhocracy* (van Alystyne, 1997; Borgatti and Foster, 2003), where agents interact because of particular role-based requirements but no overarching design principle is explicitly applied. The connection pattern superficially resembles a team, but without a team's strong interaction semantics. Some aspects of the structure may be statically defined, but a more emergent, dynamic construction is more typical. If there is an absence of explicit control over the organizational structure, the set of of interactions may change in response to every newly recognized goal. The network design is also a common basis for compound organizations in a manner similar to societies, where individual entities in the network are entire sub-organizations. These approaches can be effective and cost-efficient, but as the environment scales or the agent population becomes more dynamic a more structured organization can provide additional framework to address the more demanding context. Corkill and Lander (Corkill and Lander, 1998) enumerate several other factors which motivate the need for explicit organization, including scarce resources, the potential for collaboration and the amount of repetition of work.

Other conditions may in fact preclude the use of particular paradigms. For instance, it can be difficult to generate optimal coalition or congregation structures when there is either limited time or a large population. When individual agent resources are constrained, particular instances of organizations which suffer from bottleneck effects, such as hierarchies, federations and holarchies, can become inefficient. We have also previously noted how some types of structures, such as matrices, societies, and certain compound organizations, require a somewhat higher level of sophistication of the participating agents. As above, the operating context will guide, or in this case restrict, the choice of organizational design.

As research progresses in these areas, typically by adding features and relaxing assumptions, it can become difficult to precisely categorize a particular approach. For example, we noted how hierarchies and holarchies are closely related, as are coalitions and congregations. To a certain extent, we have focused on the extreme or most constrained examples of organizations in this paper to better delineate discrete classes, and it is frequently the case that the "rules" of a particular paradigm as we have presented them have been broken in an attempt to broaden its abilities or applicability. While this might frustrate one's attempt at categorization, our opinion is that the convergent evolution of these strategies towards a common form lends additional credence to the applicability of that form.

A somewhat more elusive goal is to define what exactly constitutes an organization in general. At what level of abstraction in the system's design should the influence of the organization diminish and more transient "operational" decisions become more important? Must a structure exist for some period of time or some number of iterations before it is considered an organization? We have looked at strategies that are generally short-lived, such as coalitions, while societies may outlast the lifetime of any of its participants. Teams may exist to satisfy only a single goal, while federations see a continuous stream of different tasks. In each of these cases, the pattern of interactions between the agents is a defining characteristic, influencing the behaviors and qualities exhibited by the system. If this same pattern exists in two different circumstances, is one an organization and the other not? To a certain extent, this is just a matter of semantics, and we could just as easily name it a "pattern of interactions" and leave it at that. However, maintaining a broad and flexible concept of organization allows one to more easily recognize that commonalities may exist between these architectures. In particular, characteristics observed in superficially different circumstances may be derived specifically from these interactions. Thus, we propose that under all circumstances this pattern can be interpreted as an organizational design. The fact that it may exist for a single moment or a single task certainly impacts its performance and construction, but much of the underlying purpose and qualities of the structuring remain the same, and should be recognized as such.

Whatever they are called, the type of short and long term patterns of interaction we have described in this article will become increasingly important as multi-agent technology is used to address more complex, real-world problems. Scale, real-time constraints and bounded rationality all conspire to create challenging environments to operate in. Because of their ability to regulate the increased complexity of the local problem solving process required in such domains, organizations should be a critical part of any comprehensive, multi-agent solution. By recognizing and understanding organizational paradigms such as those we have presented, we hope that the use of explicit organizational design is encouraged and facilitated.

14 Acknowledgments

The authors would like to thank Sherief Abdallah, Dan Corkill, Frank Dignum, Mark Sims and all our reviewers for their helpful critiques and comments on this paper.

This effort has been sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreements number F30602-99-2-0525, F30602-03-C-0010 and DOD DABT63-99-1-0004. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. This material is also based upon work supported by the National Science Foundation under Grant No. IIS-9812755. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory or the U.S. Government.

References

S. Abdallah and V. Lesser. Organization-Based Cooperative Coalition Formation. *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Techonology, IAT*, pages 162–168, September 2004.

M. Ahuja and K. M. Carley. Network structure in virtual organizations. *Organization Science*, 10(6):741–757, 1999.

A. Artikis. *Executable Specification of Open Norm-Governed Computational Systems*. PhD thesis, Department of Electrical & Electronic Engineering, Imperial College London, November 2003.

R. Axelrod. An evolutionary approach to norms. *The American Political Science Review*, 80(4):1095–1111, 1986.

K. S. Barber and C. E. Martin. Dynamic adaptive autonomy in multiagent systems: Representation and justification. *International Journal of Pattern Recognition and Artificial Intelligence*, 15(3):405–433, 2001.

G. Beavers and H. Hexmoor. Teams of agents. In *Proceedings of the IEEE Systems, Man, and Cybernetics Conference*, pages 574–582, 2001.

A. H. Bond and L. Gasser. An analysis of problems and research in DAI. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 3–35. Morgan Kaufmann, 1988.

L. Bongaerts. *Integration of Scheduling and Control in Holonic Manufacturing Systems*. PhD thesis, Katholieke Universiteit Leuven, Belgium, 1998.

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

S. Breban and J. Vassileva. Long-term coalitions for the electronic marketplace. In *Proceedings of Canadian AI Workshop on Novel E-Commerce Applications of Agents*, pages 6–12, 2001.

C. Brooks and E. Durfee. Congregating and market formation. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 96–103. ACM Press, 2002.

C. Brooks and E. Durfee. Congregation formation in multiagent systems. *Journal of Autonomous Agents and Multiagent Systems*, 7(1-2):145–170, 2003.

C. Brooks, E. Durfee, and A. Armstrong. An introduction to congregating in multiagent systems. In *Proceedings of the Fourth International Conference on Multiagent Systems*, pages 79–86, 2000.

S. Bussmann and K. Schild. Self-organizing manufacturing control: An industrial application of agent technology. In *Proceedings of the 4th International Conference on Multi-Agent Systems (ICMAS 2000)*, pages 87–94. IEEE Computer Society, July 2000.

H. L. Cardoso and E. Oliveira. Virtual enterprise normative framework within electronic institutions. In *Proceedings of the 5th Int. Workshop on Engineering Societies in the Agents World (ESAW 04)*, October 2004.

K. Carley. Organizational adaptation. *Annals of Operations Research*, 75:25–47, 1998.

K. M. Carley and L. Gasser. Computational organization theory. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 299–330. MIT Press, 1999.

B. Chandrasekaran. Natural and social system metaphors for distributed problem solving: Introduction to the issue. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1):1–5, Jan. 1981.

A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. In *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96)*, pages 75–90, London, UK, 1996. Practical Application Company.

V. Chvatal. A greedy heuristic for the setcovering problem. *Mathematics of Operations Research*, 4(3), August 1979.

J. Collins, M. Tsvetovat, B. Mobasher, and M. Gini. MAGNET: A multi-agent contracting system for plan execution. In *Proceedings of Workshop on Artificial Intelligence and Manufacturing: State of the Art and State of Practice*, pages 63–68. AAAI Press, August 1998.

J. C. Collins and L. C. Lee. Building electronic marketplaces with the ZEUS agent tool-kit. *Agent Mediated Electronic Commerce*, 1571:1–24, 1998.

M. Colombetti, N. Fornara, and M. Verdicchio. A social approach to communication in multiagent systems. In J. A. Leite, A. Omicini, L. Sterling, and P. Torroni, editors, *Declarative Agent Languages and Technologies*, volume 2990 of *Lecture Notes in Artificial Intelligence*, pages 191–220. Springer-Verlag, May 2004.

D. Corkill and V. Lesser. The use of meta-level control for coordination in a distributed problem solving network. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 748–756, August 1983.

D. D. Corkill and S. E. Lander. Diversity in Agent Organizations. *Object Magazine*, 8(4):41–47, May 1998.

G. Cuni, M. Esteva, P. Garcia, E. Puertas, C. Sierra, and T. Solchaga. MASFIT: Multi-Agent System for FIsh Trading. In *Proceedings of the 16th Eureopean Conference on Artificial Intelligence, (ECAI'2004)*, pages 710–714. IOS Press, August 2004.

M. R. Cutkosky, R. S. Englemore, R. E. Fikes, M. R. Genesereth, T. R. Gruber, W. S. Mark, J. M. Tenenbaum, and J. C. Weber. PACT: An experiment in integrating concurrent engineering systems. In M. N. Huhns and M. P. Singh, editors, *Readings in Agents*, pages 46–55. Morgan Kaufmann, San Francisco, CA, USA, 1997.

R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20 (1):63–109, Jan. 1983.

K. Decker. TAEMS: A Framework for Environment Centered Analysis & Design of Coordination Mechanisms. In *Foundations of Distributed Artificial Intelligence, Chapter 16*, pages 429–448. G. O'Hare and N. Jennings (eds.), Wiley Inter-Science, January 1996.

K. Decker and V. Lesser. Generalizing the Partial Global Planning Algorithm. *International Journal on Intelligent Cooperative Information Systems*, 1(2):319–346, June 1992.

K. Decker and V. Lesser. An Approach to Analyzing the Need for Meta-Level Communication. *International Joint Conference on Artificial Intelligence*, 1, January 1993a.

K. Decker, V. Lesser, N. Prasad, and T. Wagner. MACRON: An Architecture for Multi-Agent Cooperative Information Gathering. *Proccedings of the CIKM Workshop on Intelligent Information Agents*, December 1995.

K. Decker and V. R. Lesser. Quantitative Modeling of Complex Environments. *International Journal of Intelligent Systems in Accounting, Finance and Management. Special Issue on Mathematical and Computational Models and Characteristics of Agent Behaviour.*, 2:215–234, January 1993b.

K. Decker, K. Sycara, and M. Williamson. Cloning for intelligent adaptive information agents. In C. Zhang and L. D, editors, *Multi-Agent Systems*, pages 63–75. Springer Verlag, 1997.

C. Dellarocas and M. Klein. Civil agent societies: Tools for inventing open agent-mediated electronic marketplaces. In *Agent Mediated Electronic Commerce (IJCAI Workshop)*, pages 24–39, 1999.

V. Dignum. *A model for organizational interaction: based on agents, founded in logic*. PhD thesis, University of Utrecht, Utrecht, The Netherlands, 2003.

V. Dignum, J.-J. Meyer, H. Weigand, and F. Dignum. An organizational-oriented model for agent societies. *Proceedings of International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA'02)*, July 2002.

E. Durfee and V. Lesser. Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5):1167–1183, September 1991.

E. Durfee, V. Lesser, and D. Corkill. Coherent Cooperation Among Communicating Problem Solvers. *IEEE Transactions on Computers*, C36(11):1275–1291, November 1987.

C. B. Excelente-Toledo and N. R. Jennings. The dynamic selection of coordination mechanisms. *Autonomous Agents and Multi-Agent Systems*, 9(1-2):55–85, July - September 2004.

S. S. Fatima and M. Wooldridge. Adaptive task and resource allocation in multi-agent systems. In J. P. Müller, E. Andre, S. Sen, and C. Frasson, editors, *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 537–544, Montreal, Canada, 2001. ACM Press.

K. Fischer. Agent-based design of holonic manufacturing systems. *Journal of Robotics and Autonomous Systems*, 27(1-2):3–13, 1999.

D. Fitoussi and M. Tennenholtz. Choosing social laws for multi-agent systems: minimality and simplicity. *Artificial Intelligence*, 119(1-2):61–101, 2000.

I. T. Foster, N. R. Jennings, and C. Kesselman. Brain meets brawn: Why grid and agents need each other. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 8–15. IEEE Computer Society, August 2004.

M. Fox, M. Barbuceanu, M. Gruninger, and J. Lin. An Organizational Ontology for Enterprise Modeling. In M. J. Prietula, K. M. Carley, and L. Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*, pages 131–152. AAAI Press / MIT Press, 1998.

M. S. Fox. Organization structuring: Designing large complex software. Computer Science Technical Report CMU-CS-79-155, Carnegie-Mellon University, December 1979.

M. S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1):70–80, Jan. 1981.

J. Galbraith. *Organization Design*. Addison-Wesley, Reading, MA, 1977.

L. Gasser. Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence*, 47(1-3):107–138, 1991.

L. Gasser. Perspectives on organizations in multi-agent systems. In *Mutli-agents systems and applications*, pages 1–16. Springer-Verlag New York, Inc., 2001.

M. R. Genesereth. An agent-based framework for interoperability. In J. Bradshaw, editor, *Software agents*, pages 317–345. MIT Press, 1997.

N. Glaser and P. Morignot. Societies of autonomous agents and their reorganisation. In W. Tschacher and J. Dauwalder, editors, *Dynamics, Synergetics, Autonomous Agents — Nonlinear Systems Approaches to Cognitive Psychology and Cognitive Science*. World Scientific, 1998.

A. Glass and B. Grosz. Socially conscious decision-making. In C. Sierra, M. Gini, and J. S. Rosenschein, editors, *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 217–224, Barcelona, Catalonia, Spain, 2000. ACM Press.

N. Gnanasambandam, S. Lee, N. Gautam, S. R. T. Kumara, W. Peng, V. Manikonda, M. Brinn, and M. Greaves. Reliable MAS performance prediction using queueing models. In *Proceedings of the IEEE Multi-agent Security and Survivability Symposium (MASS)*, 2004.

N. Griffiths. Supporting cooperation through clans. In *Cybernetic Intelligence, Challenges and Advances – Proceedings of the 2nd IEEE Systems, Man and Cybernetics, UK & RI Chapter Conference*, 2003.

B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

B. Grosz and S. Kraus. The Evolution of SharedPlans. In *Foundations and Theories of Rational Agencies*, pages 227–262. Kluwer Academic Publishers, 1999.

B. Grosz and C. Sidner. Plans for discourse. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication*, pages 417–444. MIT Press, 1990.

R. H. Guttman, A. G. Moukas, and P. Maes. Agent-mediated electronic commerce: a survey. *Knowl. Eng. Rev.*, 13(2):147–159, 1998.

S. C. Hayden, C. Carrick, and Q. Yang. A catalog of agent coordination patterns. In *AGENTS '99: Proceedings of the third annual conference on Autonomous Agents*, pages 412–413, New York, NY, USA, 1999. ACM Press.

C. Hewitt. Offices are open systems. *ACM Transactions on Office Information Systems*, 4(3):271–287, July 1986.

H. Hexmoor and G. Beavers. Towards teams of agents. In *Proceedings of the International Conference in Artificial Intelligence (IC-AI'2001)*. CSREA Press, 2001.

B. Horling, B. Benyo, and V. Lesser. Using Self-Diagnosis to Adapt Organizational Structures. *Proceedings of the 5th International Conference on Autonomous Agents*, pages 529–536, June 2001.

B. Horling and V. Lesser. Analyzing, Modeling and Predicting Organizational Effects in a Distributed Sensor Network. *Journal of the Brazilian Computer Society, Special Issue on Agents Organizations*, 2005.

B. Horling, R. Mailler, and V. Lesser. A Case Study of Organizational Effects in a Distributed Sensor Network. In *Proceedings of the International Conference on Intelligent Agent Technology (IAT 2004)*, pages 51–57, Beijing, China, September 2004.

B. Horling, R. Mailler, J. Shen, R. Vincent, and V. Lesser. Using Autonomy, Organizational Design and Negotiation in a Distributed Sensor Network. In V. Lesser, C. Ortiz, and M. Tambe, editors, *Distributed Sensor Networks: A multiagent perspective*, pages 139–183. Kluwer Academic Publishers, 2003.

J. F. Hübner, J. S. Sichman, and O. Boissier. A model for the structural, functional, and deontic specification of organizations in multiagent systems. In *Proceedings of the Brazilian Symposium on Artificial Intelligence (SBIA'02)*, pages 118–128, 2002.

M. N. Huhns and L. M. Stephens. Multiagent systems and societies of agents. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 79–120. The MIT Press, Cambridge, MA, USA, 1999.

T. Ishida, L. Gasser, and M. Yokoo. Organization self-design of distributed production systems. *IEEE Transactions on Knowledge and Data Engineering*, 4(2): 123–134, 1992.

N. Jennings. Controlling cooperative problem solving in industrial multi-agent systems. *Artificial Intelligence*, 75(2):195–240, 1995.

D. Jensen and V. Lesser. Social Pathologies of Adaptive Agents. In M. B. . H. Guesgen, editor, *Safe Learning Agents: Papers from the 2002 AAAI Spring Symposium.*, volume TR SS-02-07. AAAI Press, August 2002.

H. Jung, M. Tambe, and S. Kulkarni. Argumentation as distributed constraint satisfaction: applications and results. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 324–331, New York, NY, USA, 2001. ACM Press.

G. A. Kaminka, D. V. Pynadath, and M. Tambe. Monitoring teams by overhearing: A multi-agent plan-recognition approach. *Journal of Artificial Intelligence Research*, 17:83–135, 2002.

T. Khedro and M. Genesereth. Facilitators: A networked computing infrastructure for distributed software interoperation. In *Working Notes of the IJCAI-95 Workshop on Artificial Intelligence in Distributed Information Networks.*, 1995.

M. Klein, J. Rodriguez-Aguilar, and C. Dellarocas. Using domain-independent exception handling services to enable robust open multi- agent systems: The case of agent death. *Journal of Autonomous Agents and Multi-Agent Systems*, 7(1-2):179–189, 2003.

M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17 (3):42–47, 2002.

A. Koestler. *The Ghost In The Machine*. Hutchinson Publishing Group, Arkana, London, 1967.

K. Kurbel and I. Loutchko. Towards multi-agent electronic marketplaces: What is there and what is missing? *The Knowledge Engineering Review*, 18(1):33–46, 2003.

K. Larson and T. Sandholm. Anytime coalition structure generation: An average case study. *Journal of Experimental and Theoretical AI*, 11:1–20, 2000.

K. Lerman and A. Galstyan. A general methodology for mathematical analysis of multiagent systems. Technical Report ISI-TR-529, University of California, Information Sciences Institute, 2001.

K. Lerman and O. Shehory. Coalition formation for large-scale electronic markets. In *Proceedings of the International Conference on Multi-Agent Systems (ICMAS'2000)*, 2000.

V. Lesser. A Retrospective View of FA/C Distributed Problem Solving. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1347–1363, November 1991.

V. Lesser. Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture. *Autonomous Agents and Multi-Agent Systems*, 1: 89–111, January 1998.

V. Lesser and D. Corkill. Functionally Accurate, Cooperative Distributed Systems. *IEEE Trans. on Systems, Man, and Cybernetics*, SMC-11(1):81–96, January 1981.

V. Lesser and D. Corkill. The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks. *AI Magazine*, 4(3):15–33, 1983.

V. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. Neiman, R. Podorozhny, M. NagendraPrasad, A. Raja, R. Vincent, P. Xuan, and X. Zhang. Evolution of the GPGP/TAEMS Domain-Independent Coordination Framework. *Autonomous Agents and Multi-Agent Systems*, 9(1):87–143, July 2004.

V. Lesser and L. Erman. Distributed Interpretation: A Model and an Experiment. *IEEE Transactions on Computers Special Issue on Distributed Processing*, C-29 (12):1144–1163, December 1980.

H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On acting together. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 94–99, July 1990.

Z. Lin and K. Carley. DYCORP: A computational framework for examining organizational performance under dynamic conditions. *Journal of Mathematical Sociology*, 20(2-3):193–218, 1995.

D. Lybäck. Transient diversity in multi-agent systems. Master's thesis, Department of Computer and Systems Sciences, Stockholm University and the Royal Institute of Technology, September 1999.

R. Mailler and V. Lesser. Solving Distributed Constraint Optimization Problems Using Cooperative Mediation. In *Proceedings of Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 438–445. IEEE Computer Society, 2004.

R. Mailler, V. Lesser, and B. Horling. Cooperative Negotiation for Soft Real-Time Distributed Resource Allocation. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, pages 576–583, Melbourne, July 2003. ACM Press.

T. W. Malone and S. A. Smith. Modeling the performance of organizational structures. *Operations Research*, 36 (3):421–436, 1988.

J. G. March and H. A. Simon. *Organizations*. Wiley, New York, 1958.

S. Marsella, M. Tambe, J. Adibi, Y. Al-Onaizan, G. A. Kaminka, and I. Muslea. Experiences acquired in the design of robocup teams: A comparison of two fielded teams. *Autonomous Agents and Multi-Agent Systems*, 4 (1/2):115–129, 2001.

P. Mathieu, J. C. Routier, and Y. Secq. Dynamic organization of multi-agent systems. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 451–452. ACM Press, 2002.

E. Matson and S. DeLoach. Using dynamic capability evaluation to organize a team of cooperative, autonomous robots. In *Proceedings of The 2003 International Conference on Artificial Intelligence (IC-AI'03)*, pages 744–749, 2003.

F. Maturana, W. Shen, and D. Norrie. Metamorph: An adaptive agent-based architecture for intelligent manufacturing. *International Journal of Production Research*, 37(10):2159–2174, 1999.

C. Mérida-Campos and S. Willmott. Modelling coalition formation over time for iterative coalition games. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 572–579. IEEE Computer Society, August 2004.

T. Montgomery and E. Durfee. Search reduction in hierarchical distributed problem solving. *Group Decision and Negotiation*, 2:301–317, 1993.

Y. Moses and M. Tennenholtz. Artificial social systems. *Computers and AI*, 14(6):533–562, 1995.

A. Mowshowitz. On the theory of virtual organization. *Systems Research and Behavior Science*, 14(6):373–384, 1997.

L. Mui, M. Mohtashemi, and A. Halberstadt. Notions of reputation in multi-agent systems: A review. In *Proceedings of the First International Conference on Autonomous Agents and MAS*, pages 280–287, Bologna, Italy, July 2002. ACM.

R. Nair, T. Ito, M. Tambe, and S. Marsella. The role of emotions in multiagent teamwork: A preliminary investigation. In *Who needs emotions: the brain meets the robot*. Oxford University Press, 2003a.

R. Nair, M. Tambe, and S. Marsella. Role allocation and reallocation in multiagent teams: Towards a practical analysis. In *Proceedings of Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-03)*, pages 552–559, 2003b.

T. J. Norman, A. Preece, S. Chalmers, N. R. Jennings, M. Luck, V. D. Dang, T. D. Nguyen, V. Deora, J. Shao, W. A. Gray, and N. J. Fiddian. Conoise: Agent-based formation of virtual organisations. In *Proceedings of the 23rd SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 353–366. Springer-Verlag, 2003.

L. E. Parker. Designing control laws for cooperative agent teams. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 582–587, 1993.

H. V. D. Parunak, R. Savit, and R. L. Riolo. Agent-based modeling vs. equation-based modeling: A case study and users' guide. In *Proceedings of the First International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pages 10–25, London, UK, 1998. Springer-Verlag.

H. E. Pattison, D. D. Corkill, and V. R. Lesser. Instantiating Descriptions of Organizational Structures. *Distributed Artificial Intelligence, Research Notes in Artificial Intelligence*, I:59–96, 1987.

M. N. Prasad and V. Lesser. Learning Situation Specific Coordination in Cooperative Multi-Agent Systems. *Autonomous Agents and Multi-Agent Systems*, 2:173–207, 1999.

D. V. Pynadath and M. Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of AI research*, 16:389–423, 2002.

S. Reed and V. Lesser. Division of Labor in Honey Bees and Distributed Focus of Attention. *University of Massachusetts/Amherst Computer and Information Science Department Technical Report 80-17*, November 1980.

J. Rodriguez, P. Noriega, C. Sierra, and J. Padget. FM96.5 A Java-based Electronic Auction House. In *Proceedings of 2nd Conference on Practical Applications of Intelligent Agents and MultiAgent Technology (PAAM'97)*, pages 207–224, London, UK, Apr. 1997.

P. Romelaer. Organization: A Diagnosis Method. Technical Report 78, University Paris IX Dauphine, Crepa Laboratory, June 2002.

D. H. S. D. Ramchurn and N. R. Jennings. Trust in multiagent systems. *The Knowledge Engineering Review*, 19(1):1–25, 2004.

J. Sabater and C. Sierra. Social regret, a reputation model based on social relations. *SIGecom Exch.*, 3(1):44–56, 2002.

T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

T. Sandholm. Optimal winner determination algorithms. In *Combinatorial Auctions*. MIT Press, 2005.

T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-2):209–238, 1999.

T. Sandholm and V. Lesser. Coalitions Among Computationally Bounded Agents. *Artificial Intelligence, Special Issue on Economic Principles of Multi-Agent Systems*, 94(1):99–137, January 1997.

P. Scerri, D. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17:171–228, 2002.

J. Schmitt and U. Roedig. Sensor Network Calculus - A Framework for Worst Case Analysis. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS05), Marina del Rey, USA*. IEEE Computer Society Press, June 2005.

M. Schwaninger. A theory for optimal organization. Technical Report 38, Institute of Management at the University of St. Gallen, Switzerland, 2000.

S. Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proc. of the Second International Conference on Multiagent Systems*, pages 322–329, 1996.

O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101 (1–2):165–200, 1998.

O. Shehory, K. Sycara, P. Chalasani, and S. Jha. Agent cloning: an approach to agent mobility and resource allocation. *IEEE Communications Magazine*, 36(7):58–67, 1998.

J. Shen, X. Zhang, and V. Lesser. Degree of Local Cooperation and its Implication on Global Utility. In *Proceedings of Third International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2004)*, volume 2, pages 546–553, New York, New York, July 2004. IEEE Computer Society.

W. Shen and D. Norrie. A hybrid agent-oriented infrastructure for modeling manufacturing enterprises. In *Proceedings of the Knowledge Acquisition Workshop (KAW'98)*, pages 1–19, 1998.

Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: Off-line design. *Artificial Intelligence*, 73(1-2):231–252, 1995.

J. Sichman and Y. Demazeau. On social reasoning in multi-agent systems. *Revista Ibero-Americana de Inteligncia Artificial*, 13:68–84, 2001.

C. Sierra, J. Sabater, J. Augusti, and P. Garcia. SADDE: Social agents design driven by equations. In F. Bergenti, M. Gleizes, and F. Zambonelli, editors, *Methodologies and software engineering for agent systems*. Kluwer Academic Publishers, 2004.

H. A. Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, MA, 1968.

M. Sims, D. Corkill, and V. Lesser. Separating Domain and Coordination in Multi-Agent Organizational Design and Instantiation. In *Proceedings of the International Conference on Intelligent Agent Technology (IAT 2004)*, pages 155–161, Beijing, China, September 2004.

M. Sims, C. Goldman, and V. Lesser. Self-Organization through Bottom-up Coalition Formation. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, pages 867–874, Melbourne, AUS, July 2003. ACM Press.

N. Singh, M. R. Genesereth, and M. Syed. A distributed and anonymous knowledge sharing approach to software interoperation. *International Journal of Cooperative Information Systems*, 4(4):339–368, 1995.

R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transctions on Computers*, 29(12):1104–1113, 1980.

Y. So and E. H. Durfee. Designing tree-structured organizations for computational agents. *Computational and Mathematical Organization Theory*, 2(3):219–246, 1996.

Y. So and E. H. Durfee. Designing organizations for computational agents. In *Simulating Organizations*, pages 47–64. AAAI Press/ MIT Press, 1998.

L.-K. Soh, C. Tsatsoulis, and H. Sevay. A Satisficing, Negotiated, and Learning Coalition Formation Architecture. In V. Lesser, C. Ortiz, and M. Tambe, editors, *Distributed Sensor Networks: A multiagent perspective*, pages 109–138. Kluwer Academic Publishers, 2003.

D. Stefanoiu, M. Ulieru, and D. Norrie. Fuzzy modeling of multi-agent systems behavior. vagueness minimization. In *Proceedings of World Multiconference on Systemics, Cybernetics and Informatics (SCI'2000)*, volume III, pages 118–123, July 2000.

M. Stonebraker, P. M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Staelin, and A. Yu. Mariposa: A wide-area distributed database system. *VLDB Journal: Very Large Data Bases*, 5(1):48–63, 1996.

K. Sycara, K. Decker, and M. Williamson. Middle-agents for the internet. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 578–583, January 1997.

M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

M. Tambe, J. Adibi, Y. Alonaizon, A. Erdem, G. A. Kaminka, S. Marsella, and I. Muslea. Building agent teams using an explicit teamwork model and learning. *Artificial Intelligence*, 110(2):215–239, 1999.

G. Tidhar, C. Heinze, and M. Selvestrel. Flying together: Modelling air mission teams. *Journal of Applied Intelligence*, 8(3):195–218, May 1998.

G. Tidhar, A. Rao, and L. Sonenberg. Guided team selection. In *Proceedings of the 2nd International Conference on Multi-agent Systems (ICMAS-96)*, pages 369–376, Kyoto, Japan, 1996.

M. Tsvetovat and K. Sycara. Customer coalitions in the electronic marketplace. In C. Sierra, M. Gini, and J. S. Rosenschein, editors, *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 263–264, Barcelona, Catalonia, Spain, 2000. ACM Press.

M. Tsvetovatyy, M. Gini, B. Mobasher, and Z. Wieckowski. MAGMA: An agent-based virtual market for electronic commerce. *Journal of Applied Artificial Intelligence*, 11(6):501–523, 1997.

R. M. Turner. The tragedy of the commons and distributed AI systems. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, pages 379–390, Hidden Valley, Pennsylvania, 1993.

M. Ulieru. Emergence of holonic enterprises from multi-agent systems: A fuzzy-evolutionary approach. In V. Loia, editor, *Soft Computing Agents: A New Perspective on Dynamic Information Systems*, pages 187–215. IOS Press, 2002.

M. Ulieru, S. Walker, and R. Brennan. Holonic enterprise as a collaborative information ecosystem. In *Proceedings of the Workshop on Holons: Autonomous and Cooperating Agents for Industry 2001*, pages 1–14, May 2001.

M. van Alystyne. The state of network organization: a survey in three frameworks. *Journal of Organizational Computing and Electronic Commerce*, 7(2&3):83–151, 1997.

J. Vázquez-Salceda, V. Dignum, and F. Dignum. Organizing multiagent systems. Technical Report UU-CS-2004-015, Institute of Information & Computing Sciences, March 2004.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

T. Wagner and V. Lesser. Relating Quantified Motivations for Organizationally Situated Agents. In *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages, Lecture Notes in Artificial Intelligence*, pages 334–348. N. R. Jennings and Y. Lesperance (eds.), Springer-Verlag, Berlin, April 1999.

A. Walker and M. J. Wooldridge. Understanding the emergence of conventions in multi-agent systems. In *Proceedings of the First International Conference on Multi–Agent Systems*, pages 384–389, San Francisco, CA, 1995.

M. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

M. Wellman. Online marketplaces. In M. P. Singh, editor, *Practical Handbook of Internet Computing*. Chapman Hall & CRC Press, Baton Rouge, 2004.

M. Wellman, W. Walsh, P. Wurman, and J. MacKie-Mason. Auction protocols for decentralized scheduling. Technical report, University of Michigan, July 1998.

M. Wellman and P. Wurman. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24:115–125, 1998.

M. Wellman and P. Wurman. A trading agent competition for the research community. In *Proceedings of the IJCAI-99 Workshop on Agent-Mediated Electronic Trading*, August 1999.

G. Wiederhold, P. Wegner, and S. Cefi. Toward megaprogramming. *Communications of the ACM*, 33(11):89–99, 1992.

S. Willmott, J. Dale, B. Burg, P. Charlton, and P. O'Brien. Agentcities: A worldwide open agent network. *Agentlink Newsletter*, 8:13–15, November 2001.

P. Wurman, M. Wellman, and W. Walsh. A parameterization of the auction design space. *Games and Economic Behavior*, 35(1-2):304–338, 2001.

F. L. y López, M. Luck, and M. d'Inverno. Normative agent reasoning in dynamic societies. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 732–739. IEEE Computer Society, 2004.

O. Yadgar, S. Kraus, and C. Ortiz. Scaling up distributed sensor networks: cooperative large-scale mobile-agent organizations. In V. Lesser, C. Ortiz, and M. Tambe, editors, *Distributed Sensor Networks: a multiagent perspective*, pages 185–218. Kluwer publisching, 2003.

X. Zhang, V. Lesser, and T. Wagner. Integrative Negotiation in Complex Organizational Agent Systems. In *Proceedings of the 2003 IEEE/WIC International Conference on Intelligent Agent Technology (IAT 2003)*, pages 140–146, Halifax, Canada, 2003. IEEE Computer Society.

X. Zhang and D. Norrie. Holonic control at the production and controller levels. In *Proceedings of IMS 99*, pages 215–224, 1999.