

The Role of Roles in Designing Effective Agent Organizations

James J. Odell¹, H. Van Dyke Parunak², and Mitchell Fleischer²

¹ James Odell Associates, 3646 West Huron River Drive, Ann Arbor
MI 48103-9489 USA
email@jamesodell.com
http://www.jamesodell.com

² Altarum, 3520 Green Court, Suite 300, Ann Arbor,
MI 48105 USA
{van.parunak,mitch.fleischer}@altarum.org
http://www.erim.org/{~vparunak,~mfleischer}

Abstract. Agent-based systems are no longer contained within the boundaries of a single, small agent organization. To meet the demands of large-scale system implementations, agent organizations must deal with environmental forces, interact with other agent organizations, and know how they affect individual agents. In this paper, we look to social and organizational systems theory as a source of inspiration. Many of these techniques have been successful for a hundreds and thousands of years. We believe that the designers of agent-based systems can learn a great deal from organization designers. In the first of a series, this paper examines the notion of *role* and its implications on how agents might behave in group settings.

1 Introduction

We simply have hardly any real experience building truly heterogeneous, realistically coordinated multiagent systems that work together, and ... almost no basis for systematic reflection and analysis of that experience [1].

Societies need to employ patterned behavior to exist. The behavior of each individual is determined to a great extent by the requirements of these patterns [2]. However, the current practice of MAS design tends to be limited to individual agents and small face-to-face groups of agents that operate as closed systems. We have little principled understanding for:

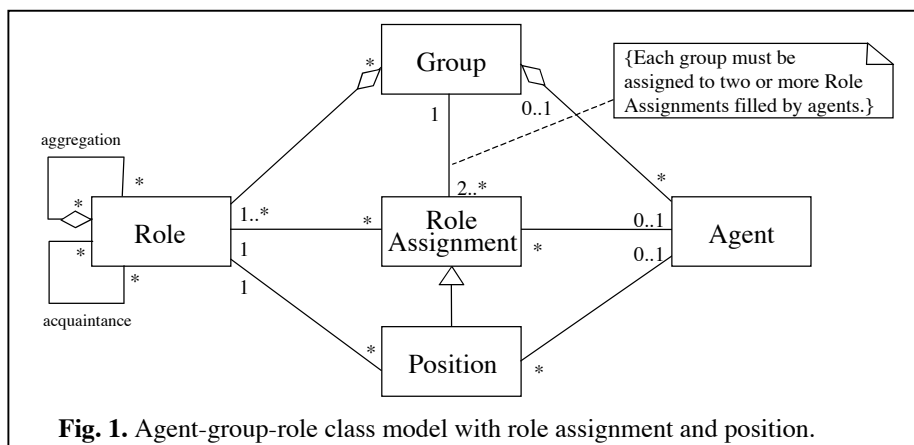
- organizing sophisticated, interactive, heterogeneous agent-based systems.
- grouping the agents in such systems into very large-scale aggregates that exhibit predictable, stable, and reliable behavior.
- achieving economies of scale and scope within large MAS.
- building and operating such systems *in situ*.

From a scientific standpoint, the foundations for constructing large multiagent systems have a long history. Although researchers have been explicitly thinking about MAS/DAI organizations and attempting to link organization theory with MAS/DAI models for decades, the idea of organization, *per se*, has been only a pe-

ripheral theme. MAS/DAI researchers have focused on specific coordination techniques, rather than the central issues involved in MAS organization.

In previous papers, we introduced the notion of groups and roles [9], initially explored the effect of the environment on agents [10], and discussed frequently characterized forms of agent interaction [11]. We are now beginning a series of papers that explore ways of designing effective MAS organizations. This paper addresses *roles*. The usefulness of the concept is apparent by the degree to which it is referenced (often intuitively) in newer agent development methodologies, such as Gaia [24] and MESSAGE [25]. Our concern here is to examine in more detail what the notion means and how roles might be employed in agent-based systems. Forthcoming papers will address related issues, such as:

- groups – defining sets of related agent roles that support certain kinds of organizations
- group structures – specifying hierarchic and network configurations among groups
- policies and contracts – expressing rights, norms, rules, agreements, and enforcement mechanisms for agent organizations
- communication and interaction protocols – discussing patterns of communications among roles and their uses
- activity structures – expressing work plans and business processes for agent-based systems
- ontology - building and expressing sets of concepts and their terms adopted by agents and their groups
- goals – defining the state a group is trying to attain (e.g., explicit missions, objectives, critical success factors)
- beliefs, desires, intention – imputing cognition to an organization as a way for external entities to reason about the organization and its members.
- emotions and morale – formalizing the notion and application of “feelings” for agent-based organizations as an approach for its human users to reason and understand agent social behavior.



Using a UML class diagram, Fig. 1 illustrates the primary concepts involved in expressing roles. Below are definitions for some of the key concepts we are using in

this paper. In particular, we will be discussing the first four notions more fully (the last three are for background, only).

A *role* is a class that defines a normative behavioral repertoire of an agent. While roles are defined independently from groups, they must be played within groups. Roles can be composed of other roles. Furthermore, roles can have acquaintance associations with other roles, denoting that interaction may occur among the instances of the related roles.

A *position* is a formally recognized role assignment that may occur within a group structure. A position may or may not have an agent assigned to it. Only one agent may be associated with it at any point in time. However, many different agents may be associated with the position over time.

A *role assignment* is the association of a group with a role that may be occupied by an agent. Each mapping of an agent to a role assignment can be a result of an exogenous or an endogenous process.

A *group* is a set of two or more agents that are related via their role assignments, where these relationships must form a connected graph within the group. Agents and Roles are associated with Groups to provide context.

A *social system* consists of a set of interdependent role behaviors, providing a collective pattern in which agents play their parts, or roles.

An *organization* is a group whose roles and interactions are typically expected to be relatively stable and change slowly over time.

An *environment* provides the conditions under which an entity (agent or object) exists. It includes the groups to which an agent may belong.

In the balance of this paper we develop the notion of role intuitively (Section 2), explain how roles can be defined (Section 3) and assigned to agents (Section 4), and then look ahead to how roles enable us to define groups (Section 5).

2 Roles

The notion of *role* is fundamentally a thespian concept, and attention to how it functions in the theater can reinforce our intuitions and provide useful metaphors for application to multiagent systems.

All the world's a stage,
And all the men and women merely players:
They all have their exits and entrances;
And one man in his time plays many parts.
—W. Shakespeare, *As You Like It*, Act II, Scene 7.

Can it be possible that Shakespeare's lines have anything to do with the field of behavioral science known as *role theory* [12]? The similarities between the Shakespeare's characterization and our present-day usage of *role* are noteworthy. Both express a perspective on human behavior. For Shakespeare, social life was similar to stage acting with its scenes, masks, and airs; in role theory, life *is* the stage. Both also employ a special language. Shakespeare wrote of players and parts that are played; role theory discusses individuals and the roles they play.

As humans, we find the perspective and language of the theater a useful analogy for describing and understanding many of the same complex aspects of individual behavior. The *role* perspective consists of those factors presumed to be influential in governing human behavior. Since we commonly employ this device in real life for conceptualizing human behavior, it may also serve as a useful device for other kinds of individuals in a MAS—be they life forms, active software constructs, or hardware devices. The notion of role permeates other aspects of the theater analogy, such as:

- *Actor* - The individual assigned to play a part, or role. In a MAS, an actor is an agent.
- *Part* - In ancient Greece and Rome, the parts of theatrical characters were literally read from rolls of parchment. When translated into French, the term “roll” becomes “role” [13]. Since the sixteenth century, *role* has been an English synonym for a part that one plays or assumes in society or life [14]. Employing the notion of role for multiagent systems is in keeping with centuries of accepted usage.
- *Script* – In a multiagent system, the social script may be as constraining as that of a traditional play, or it may permit the agent actor some degree of improvisation that is consistent with the role. In other words, role behavior can be strictly or loosely defined. However, even in rote-based scripts a degree of latitude always exists. Here, actors can at least employ their own style and interpretation as well as exercise some degree of variance in terms of appearance, movement, and so on.
- *Performance* – The player acts the part. Actors can at least employ their own style and interpretation of a script (even in a verbatim performance). They can also exercise some degree of variance in terms of appearance, movement, and so on. In society, role performance is determined by social norms; in agent-based systems, performance can be determined by sets of constraints that govern both the behavior of each role and the behavior of all agents in a society, independent of the role.
- *Play* – In the theater, a play is a literary composition in the form of dialogue assigned to specific roles adapted for staged performance. In society, this is present in certain situations or contexts where scripts and role performances are required or expected. For example, a marriage ceremony has well defined roles and a verbatim-style script. In contrast, driving a car involves roles that have no predefined verbal dialogue. Instead, the dialogue of the Driver role involves signals and hand gestures along with driving behavior that is expected to conform to certain social and legal norms. In agent-based systems, the analog of a *play* would be a *system*, *application*, or *scenario*. For example, a scenario for a British auction has a well defined interaction protocol involving agents playing the roles of bidder, seller, and auctioneer.
- *Director* – In a play, the director supervises and controls its production; he or she determines how the roles are to be filled and how the parts should be interpreted. Supervisors, assembly line coordinators, parents, coaches, and police are examples

of those who govern individuals and roles in human societies. Agent societies, too, can require supervision and control by agents that have been given authority known as *distinguished agents*. Distinguished agents can come in either human form (e.g., an agent designer, system architect, or operator) or automated form (a central controller or another agent with the appropriate power). Directors are not always a required element. The directors could be free to determine for themselves how the script is played out. The end result may or may not be as good as having a knowledgeable director to refine the movement of the actors. Yet, improvisational plays are probably best performed without central control.

- *Audience* – The viewers of a play both observe the performance and react to the actors' role playing. An agent system could be similarly observed for several reasons. The primary mechanism would be by monitoring the system and recording its progress and status.

In essence, the role perspective in theater, life, human-resource departments, and agent-based societies assumes that performance results from social proscriptions and individual behavior and that the individual variations in performance are expressed within the framework created by these factors. Here, an individual's behavior is shaped by the demands and rules of others, the individual's own understanding of appropriate behavior, and the individual's competence in the performance. In human-based organizations, defining normative forms of behavior is quite common. For example, human resource departments define and specify roles to ensure some degree of standardization within an organization. Such roles can be defined generally as Staff Member and Manager or, more specialized, as Intermediate Programmer III and Principal Analyst. This perspective does not deny individual differences. Instead, it highlights the social determinants that may be involved in creating such differences. Furthermore, it can help *role designers* focus on those conditions where social determinants will be more influential.

3 Designing Agent Roles

Roles define normative behavioral repertoires for agents. They provide both the building blocks for agent social systems and the requirements by which agents interact. Each agent is linked to other agents by the roles it plays by virtue of the system's functional requirements—which are based on the expectations that the system has of the agent. This section first describes how roles are formed, and second, how roles can be configured both in terms of the number and complexity of actions supported by a given role and the degree of control the agent has over various actions.

3.1 Role Formation

Roles can be assigned to agents in a multiagent system in at least two ways: endogenously (by emergent self-organization as the system runs), and exogenously (by the system designer when the system is constructed or modified).

Endogenous self-organization is a widespread phenomenon in natural systems [3][4][5]. As agents interact with each other, system-level structures and patterns emerge

that can adapt robustly to changes in the system's environment. Role differentiation in social insects is a prominent example of such self-organization [7]. For instance, individual *Polistes* wasps in a nest spontaneously take on different roles (nurse vs. forager vs. "chief") in response to changing food resources and brood sizes [8]. Synthetic systems can exploit similar dynamics. Experiments with information brokers in a web-based information system [6] show that the brokers can emergently specialize into complementary niches, thus adopting distinct information roles.

Self-organizing role assignment is robust to system change and replaces expensive and time-consuming human design effort with autonomous configuration. Thus it is a particularly attractive way to design systems for domains that are subject to unexpected change. Even the human social structures with which we are intuitively familiar are ultimately generated and maintained by emergent self-organization. Many social psychologists suggest that roles are formed when a recurring set of behaviors is found useful in a group context [15, 16, 17, 18]. Initially, the behavior could have been the actions of a single individual. Over time, however, the individual's actions might be identified as a useful set of behaviors that can be employed by other individuals to produce similar results. This reusable set of actions becomes standardized in the form of a role. These roles can be extracted analytically by techniques such as social network analysis [23][24] or Dooley graphs [21].

However, there are several reasons that one might want to begin with existing human structural patterns and impose them exogenously on a multiagent system at design time. These reasons fall into two main categories: cognitive fit and dynamic properties.

Imposed roles may be preferable to emergent roles in a system that must interface closely with human agents. There are usually several different ways that roles can be assigned to a set of agents with a given function. Those roles that emerge in a self-organizing digital system may not align naturally with those that a human organization would adopt to address the same function. This difference can make it difficult for people to understand the behavior of the digital system. If human stakeholders cannot understand a system, they will not trust it, severely compromising its usefulness.

Dynamically imposed roles may be useful in initiating a system to which one wishes to apply emergent self-organization. One initializes the multiagent system with a set of roles known to provide the desired functionality in human organizations, which greatly reduces the burden of subsequent self-organization. Instead of searching the entire space of possible role assignments, self-organization can be used to fine-tune the initial assignment, reducing the operating time necessary to achieve a desired level of function. By starting with known human roles, one also encourages self-organization to find assignments that are "close" to those roles and thus more tractable to human cognition.

Readers will recognize our interest in and research on self-organizing mechanisms in other publications [3][10][20]. In this paper, we focus our attention on defining patterns of human-like roles that can be imposed at design time. The synthesis of these two approaches is a major area of current research for us.

3.2 Role Configuration

Configuring a role can be thought of along two dimensions: breadth and depth. Role breadth, or *horizontal specialization*, addresses the number and complexity of actions supported by a given role. Depth, or *vertical specialization*, separates the performance of actions from the administration of them. That is, it relates to the degree of control an agent can have over its actions and the actions of other agents [19].

Horizontal specialization of roles

At one extreme, a role could require an agent to be a jack-of-all-trades, handling any kind of request. In this scenario, there is no horizontal specialization—only “fat” homogeneous agents. In small, simple organizations, such versatility can be useful. However, in large, diversified organizations (whether human or MAS systems), agents cannot be expected to do everything. Instead, diversity and limited capability are the guidelines for role design.

In the other extreme, a role could simply require the repertoire of a single, highly specialized action. Such an approach is common in the traditional assembly line where each worker has a single, specific task to perform. Simple roles have two primary characteristics. First, each incumbent of a role assignment can perform its action separately and independently from other role assignments. Second, the incumbent can perform its task without knowing how other assembly-line actions fit together to produce the end product. Simple roles are commonly used because they provide building blocks that are both easier to understand and simpler to develop than more broadly defined roles. However, the resulting population of narrowly defined roles will be larger than more elaborate ones for two reasons. The simpler the individual role, the more roles will be required by a complex activity to accomplish the same goal. Because of their narrow design, simple roles can sometimes require coordinator roles to organize them. Coordination is not always required for simple roles, particularly when it can occur among the roles through mechanisms such as mutual adjustment, standards, and stigmergy. Coordinator roles are just one of many mechanisms for obtaining coordination.

A high level of horizontal specialization is particularly useful when two situations exist: independent action is useful and an understanding exists of how the other actions fit together to deliver the end product. Some degree of horizontal specialization can increase productivity for several reasons:

- Specialized roles can be identified and specified that use repetitive actions which are expected to produce successful results.
- Incremental improvement is encouraged due to its modular nature. As actions become more specialized and standardized they tend to become smaller and simpler units of behavior. Incrementally modifying behavior at a fine-grain level tends to be easier than at a coarse-grain level.
- Role assignment could be based on particular criteria, enabling the system to match and assign an individual agent to a role.

Specialization is a matter of degree, so that a role can include more than one specialized action and actions can be broad rather than narrow. The richer and more complex a role is, the greater the requirements for the agent filling the role assign-

ment. Coarse-grained roles carry with them the problems and benefits of coarse-grained components. For example, driving a hook-and-ladder fire engine requires a front driver, rear driver, and a navigator/radioman. Each performs a separate specialized action that depends on close coordination with the others. Here, actions are predominantly concurrent and interdependent; whereas the assembly-line specialization (described above) involves actions that are predominantly serially dependent in nature. In other words, these roles are less specialized than the traditional assembly line where each worker role is not concerned about close coordination with other roles.

Vertical specialization of roles

While horizontal specialization addresses the performance of work, vertical specialization addresses the management of work. More precisely, vertical specialization addresses the degree of control that a role has over its actions as well as the actions of other roles. Management roles implement those tasks associated with planning, coordination, and monitoring of the agents within a group. The primary purpose of specializing roles vertically is to insure that horizontally specialized behavior results in the identified goals. In essence, then, these roles define management functions that administer the linkage among agents playing planned horizontal roles. Manager roles differ both in terms of how, why, where, and who they control, as well as the breadth and management level of their control. In all cases, they manage behavior that leads to specific goals.

At one extreme, the role of an agent that simply carries out its activities only when directed by another agent is very narrow vertically. In contrast, the role of an agent that assumes control of much of the decision making for other agents is vertically broad. The shift from passive responder to active manager is another way of saying that a role's vertical dimension shifts from narrow (or specialized) to broader (or more general). At the far extreme of vertical generalization, a fat jack-of-all-managers role would control the actions of every agent including itself.

In between these extremes exists a range of possibilities. For example, a manager role could still require management in a system where every agent requires checks and balances. And a very simple agent could follow simple rules involving agent-based pheromones. For instance, if it detects a pheromone, it will follow the "scent" gradient towards the "food." No manager role is necessary. In fact, many effective solutions for agent-based organizational designs employ swarm intelligence techniques that require little or no vertical specialization. Some forms of human organizations do not require management roles, either. Without vertical specialization, however, many human-based organizations, such as most armed forces and legal systems, would fail utterly.

4 Role Assignments and Positions

Roles are established so that actors may play them. When an agent is actually assigned to play a role, this is known as *role assignment*. Assignments to roles can be made by a distinguished agent or established upon request by the agent itself. In

either case, a set of prerequisite conditions must be met for the assignment to be made.

A *role assignment* is the association of a group with a role that may be occupied by an agent. Each mapping of an agent to a role assignment can be a result of an exogenous or an endogenous process.

Formally, a role assignment is a relation between an agent and a role. However, having a role-filling mechanism that does not always have a particular agent assigned is also useful. These are called *positions*.

A *position* is an officially recognized potential for role assignment within a group structure. A position may or may not have an agent assigned to it. Only one agent may be associated with it at any point in time. However, many different agents may be associated with the position over time.

The notion of “position” is useful to organization designers, because it encourages the designer to identify and specify those jobs that need to be performed. This is accomplished before any agents physically exist to fill the positions. A position can be established to last for long periods of time or merely created for a brief, ad hoc, usage. For example, long-term positions could include positions such as facilitators and librarians. Short-term positions could be established for each instance of a workflow activity. For instance, the workflow for Order #12345 would include activities such as Accept Order, Fill Order, and Ship Order. Here, a temporary position could be requisitioned for each activity defined for Order #12345 and filled by a qualified agent. The manager agent in charge of the Order #12345 workflow would be responsible for creating and filling each position. When each activity completes its job, the position is “vacated” and terminated. In fact, even the manager job could be requisitioned as a short-term position in an Order workflow, to be filled as needed by an agent designed as a general order-processing manager.

In other words, defining positions enables system designers and agent managers:

- to anticipate organizational needs by specifying positions first and populating them later.
- to create positions on-the-fly and fill them with agents they deem qualified.
- to remove an agent from its position when it cannot function effectively (e.g., its platform is down, or the agent is overloaded with other activities at a particular moment) and be replaced with another agent.
- to fill the same position serially with different agents over time, allowing agents to be purchased from different vendors, constructed for specific platforms using particular languages, or adapted over time to fill a given position.

In short, a position is a specification for an agent to fulfill *potentially*. It specifies the role to be assigned, criteria for occupying the position, and the supervisory relationship that it might have with other positions.

By their nature, positions partition the organization into discrete processing units, thereby placing barriers across the organization. Organizations that were thought of as seamless entities instead become aggregations of positions that must communicate

and coordinate. The more positions there are, the more complexity there is within the organization—whether the positions are vertical or horizontal. The fewer positions, the more complexity there is inside of the *agent*. Also, the more complex the environment, the more positions there can be.

Position specialization also creates a balancing problem. When positions are partitioned into more specialized positions, how many of each kind should be created? For example, if all the Sales positions are specialized as either Merchandise Sales or Service Sales positions, how many of each should be created? If there were ten Sales positions before, should there now be five of each new specialization? What is the optimal number? What happens when the number required at any time swings widely from needing many Service Sales positions to needing many Merchandise Sales positions? Having a single type of Sales position handles all situations. Yet, if merchandise and service sales situations are vastly different, a single type of Sales position could require too many duties, skills, responsibilities, and supervisory relationships for any one agent to adequately satisfy.

5 Groups – The Next Frontier

The only class in Fig. 1 that we have not discussed is the Group. As defined earlier:

A *group* is a set of agents that are related via their roles, where these relationships must form a connected graph within the group.

This definition implies not only that a group is a function of the roles contained within it, but also that roles have no meaning without their group referent. Hence, our ability to understand roles is limited by our ability to understand the groups of which they are a part. As with roles, groups may be deliberately established (i.e., by a system designer) or they may be emergent. In human organization terms, a deliberately established group could be a department or other workgroup that has been defined by some organizational authority. In contrast, an emergent group might be a social group that forms when several individuals decide to go out for a beer after work. Over time, they define themselves as a group (“My Friday Afternoon Drinking Buddies”).

In a subsequent paper we will consider the properties of agent groups and how they can either be deliberately established or system parameters set in order to foster the emergence of more effective groups. Some questions that we expect to address in that subsequent paper include:

- What system parameters affect emergent group formation and group effectiveness?
- What (if anything) differentiates deliberate groups from emergent groups?
- How should positions be grouped? To form groups that support individual, assemblage, team, or workflow activities—or some other organization unit?
- How should groups be aggregated into larger groups?
- To what extent should group structure and content be standardized?
- What is the purpose and function of the manager roles in each group?
- How much administrative power should be delegated to the manager agents and how much to the other agents in each group?

6 Conclusion

Societies need to employ patterned behaviors to exist. The behavior of each individual is determined to a great extent by the requirements of these patterns. In this paper, we have begun examining human-based organizational techniques as a source of inspiration for designing effective agent-based organizations. In particular, this paper has addressed the notion of *roles*: what the notion means and how roles might be employed in agent-based systems. Forthcoming work will address groups, group structures, policies and contracts, communication and interaction protocols, activity structures, ontology, goals, BDI's, and emotions and morale for agent organizations.

References

- [1] Gasser, Les, "Perspectives on Organizations in Multi-Agent Systems," *Multi-Agent Systems and Applications*, Michael Luck *et al.* eds., Springer-Verlag, Berlin, 2001, pp. 1-16.
- [2] Katz, Daniel, and Robert L. Kahn, *The Social Psychology of Organizations*, (2nd ed.), John Wiley and Sons, New York, 1978.
- [3] Parunak, H. V. D., "'Go to the Ant': Engineering Principles from Natural Agent Systems," *Annals of Operations Research* 75: 69-101, 1997.
- [4] Bonabeau, E., M. Dorigo, *et al.*, *Swarm Intelligence: From Natural to Artificial Systems*, New York, Oxford University Press, 1999.
- [5] Camazine, S., J.-L. Deneubourg, *et al.*, *Self-Organization in Biological Systems*. Princeton, NJ, Princeton University Press, 2001.
- [6] Kephart, J. O., J. E. Hanson, *et al.*, "Dynamics of an Information-Filtering Economy," *Cooperative Information Agents II*, M. Klusch and G. Weiss, eds. Berlin, Springer-Verlag, 1435: 160-170, 1998.
- [7] Beshers, S. N. and J. H. Fewell, "Models of Division of Labor in Social Insects," *Annual Review of Entomology*, 46: 413-440, 2001.
- [8] Theraulaz, G., S. Goss, *et al.*, "Task Differentiation in *Polistes* Wasp Colonies: A Model for Self-Organizing Groups of Robots," *First International Conference on Simulation of Adaptive Behavior*, MIT Press, 1991.
- [9] Parunak, H. Van Dyke, and James Odell, "Representing Social Structures in UML," *Agent-Oriented Software Engineering (AOSE) II*, Michael Wooldridge *et al.* eds., Springer-Verlag, Berlin, 2002, pp. 1-16.
- [10] Odell, James, H. Van Dyke Parunak, Mitch Fleischer, and Sven Brueckner, "Modeling Agents and their Environment," *Agent-Oriented Software Engineering (AOSE) III*, Fausto Giunchiglia *et al.* eds., Springer-Verlag, Berlin, 2003 (forthcoming).
- [11] Parunak, H. Van Dyke, S. Brueckner, M. Fleischer, and J. Odell, "Co-X: Defining what Agents Do Together," In *Proceedings of Workshop on Teamwork and Coalition Formation*, AAMAS 2002, 2002.
- [12] Biddle, Bruce J., and Edwin J. Thomas, *Role Theory: Concepts and Research*, John Wiley and Sons, New York, 1966.
- [13] Moreno, J.L. ed., *The Sociometry Reader*, The Free Press, Glencoe, IL, 1960.

- [14] Oxford English Dictionary, (2nd ed.), Oxford University Press, Oxford, 1992.
- [15] Bormann, E.G., and N.C. Bormann, *Effective Small Group Communication*, Burgess Publications, Minneapolis, MN, 1988.
- [16] Burke, Peter J., "The Development of Task and Social-emotional Role Differentiation," *Sociometry*, 30 (December), 1967, pp. 379-92.
- [17] Burke, Peter J., "Role Differentiation and the Legitimation of Task Activity," *Sociometry*, 32 (June), 1968, pp. 159-68. Reprinted by Warner Modular Publications, Inc. (R466), 1973.
- [18] Diamond, M.A. and Seth Allcorn, "Role Formation As Defensive Activity In Bureaucratic Organizations," *Political Psychology*, Vol. 7, No. 4, December, 1986.
- [19] Mintzberg, Henry, *Structure in Fives: Designing Effective Organizations*, Prentice Hall, Englewood Cliffs, NJ, 1993.
- [20] Zambonelli, F. and H. V. D. Parunak, "From Design to Intention: Signs of a Revolution," First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), Bologna, Italy, 2002.
- [21] Parunak, H. V. D., "Visualizing Agent Conversations: Using Enhanced Dooley Graphs for Agent Design and Analysis," Second International Conference on Multi-Agent Systems (ICMAS'96), 1996.
- [22] Krackhardt, D. and J. Hanson, "Informal Networks: the Company behind the Chart," *Harvard Business Review*, 71(4), July/Aug., 1993, pp. 104-111.
- [23] Rogers, E.M. and D.L. Kincaid, *Communication Networks: toward a new paradigm for research*, New York: Free Press, 1981.
- [24] Wooldridge, M., N.R. Jennings, and D. Kinny, The Gaia Methodology for Agent-Oriented Analysis and Design. *International Journal of Autonomous Agents and Multi-Agent Systems*, 3(3), September 2000, pp. 285-312.
- [25] Evans, R., et al., *Methodology for Agent-Oriented Software Engineering*, EURESCOM, Project P907 Deliverable 3, final report, 20 September, 2001.