Towards a Hierarchical Taxonomy of Autonomous Agents *

Predrag T. Tosic and Gul A. Agha
Open Systems Laboratory, Department of Computer Science
University of Illinois at Urbana-Champaign, USA
E-mail: {p-tosic, agha}@cs.uiuc.edu

Abstract – *Autonomous agents have become an influential and powerful paradigm in a great variety of disciplines, from sociology and economics to distributed artificial intelligence and software engineering to philosophy. Given that the paradigm has been around for awhile, one would expect a broadly agreed-upon, solid understanding of what autonomous agents are and what they are not. This, however, is not the case. We therefore join the ongoing debate on what are the appropriate notions of autonomous agency. We approach agents and agent ontology from a cybernetics and general systems perspective, in contrast to the much more common in the agent literature sociology, anthropology and/or cognitive psychology based approaches. We attempt to identify the most fundamental attributes of autonomous agents, and propose a tentative hierarchy of autonomous agents based on those attributes.*

Keywords: Smart Agents, Intelligent Systems, Intelligent and Soft Computing - Systems and Applications

1 Introduction and Motivation

Autonomous agents have become a powerful paradigm in a great variety of disciplines, from sociology and economics to distributed artificial intelligence and software engineering to cognitive sciences to philosophy. While different disciplines have different needs and may have different notions of *agents*, the agents in economics and those in distributed artificial intelligence (DAI), for example, nonetheless tend to share most of the fundamental properties. Given that the paradigm has been around for awhile, one would expect a relatively solid understanding of what autonomous agents are, and what they are not. This, however, is not at all the case.

Is there, then, at least a reasonably unified and broadly agreed upon notion of *autonomous agency* among the computer scientists? The answer is still *'No.'* In particular, the notion of autonomous agency in open distributed computing environments (e.g., [2, 6]), while certainly sharing some of the properties, does not coincide with the corresponding standard notion of agency in artificial intelligence (e.g., [18]). One would hope that understanding this gap can assist in bridging it, thereby enhancing the ability of the software system designers to meet the requirements of various AI and other applications more effectively and easily by readily identifying and efficiently building the required additional functionality ("agent capabilities") on the top of the existing open distributed agent-based (or even merely object-based) software infrastructures.

Herewith, we attempt to contribute to the ongoing debate on what are the appropriate notions of *autonomous agency.* Instead of proposing a single such prescriptive (and therefore necessarily also restrictive), "one size fits all" definition of autonomous agents, we propose an entire hierarchy of agents, from simpler (reactive situated agents) towards quite complicated and capable of human-like complex cognitive tasks (deliberative, intelligent agents). The proposed hierarchy, rather than being based on any particular school of thought in artificial intelligence and cognitive sciences, is chiefly based on ideas and paradigms from other scientific disciplines - mainly *cybernetics* [25] and *systems science* [14, 15]. We argue that learning from other, non-AI and non-cognitive disciplines such as cybernetics or biology can provide some critical, yet thus far for the most part missing, ingredients in building successful and complete theories of artificial autonomous agents and multi-agent systems. This work is intended to be a modest step in that direction.

2 What Are Autonomous Agents?

It has become common to define an appropriate notion of agency by specifying the *necessary attributes* that all agents of the particular kind one has in mind are required to share (e.g., [10, 16, 18]). There has been much of debate, however, what set of properties exactly qualifies an entity, such as a single human decision maker, a firm in the market, a computer program, a robot or an unmanned autonomous vehicle, for an *autonomous* or an *intelligent* agent. Influential position papers, such as [26] for intelligent agents or [10] for autonomous agents, while trying to clarify and unify the terminology, and propose agent taxonomies, also illustrate the heterogeneity and lack of agreement on the definition and the required (as opposed to optional) properties even in the case of autonomous agents that are restricted to computer programs alone (which disallows, say, humans or social insects).

It has been observed that the main division line is the one that separates the (purely) *reactive agents* [17, 18] from the more complex, capable of cognitive-like behaviors *deliberative agents* [16, 18, 26]. A reactive agent is one that is coupled to the environment and is capable of being affected by, and perhaps in turn also affecting, the environment. It need not be capable of cognitive tasks such as learning, planning or reasoning. It need not have any complicated internal structure, or any capability of complex correlations between its internal states and the states of the outer world ("symbolic representations"); it uses a little or no memory, etc.

In contrast, a deliberative agent is much more complex in terms of its internal structure, is typically capable of creating and working with abstract representations of a complex outer world (e.g., by performing planning, reasoning and/or learning tasks), has some sense of its purpose (tasks, goals, utilities), usually is pro-active and adaptable, etc. Much of research in the main-stream artificial intelligence (AI) over the past twenty or more years has been focused on the design problem of such artificial deliberative agents, capable of acting in complex environments and autonomously pursuing their complex goals or tasks in such environments (see, e.g., [6, 16, 18, 24, 26] and references therein).

Herein, we attempt to hierarchically classify agents based on their complexity in terms of their capabilities and functionalities, not on (models of) agents' internal structure. An agent is more sophisticated than another, if it is capable of more complex behaviors observable by an outside observer. This natural functionalist, behaviorist and systems theory oriented approach, however, does not seem very common in the mainstream agent literature.

Some of the most frequently encountered general properties of agents found in the literature include reactiveness, pro-activeness, ability to execute autonomously, goal-orientedness or goal-drivenness, a capability of sensing the environment and being affected by the environment, a capability of affecting the environment, sociability, ability to communicate, persistence, purposefulness, and ability to learn and/or reason about the world.

Not all the agents have to possess all of the above mentioned properties, of course. We shall make an attempt, however, to identify those properties that are *necessary* for autonomous agents of a desired level of complexity.

In case of the computer programs, being capable of autonomous execution, that is, an execution that is not (entirely) controlled from the outside, seems to be the most natural requirement for any notion of autonomous agency. However, a question then arises, is this enough? For instance, a *finite state machine (FSM)* executes autonomously (and reactively, inasmuch as the ability of an agent to be affected by the environment suffices for reactivity), but we find it hard to consider *individual FSMs* an appropriate abstraction of autonomous agents. On the other hand, a *coupled* finite automata model has been proposed as an abstraction of reactive situated agents (e.g., [17]). We shall discuss in some detail what we consider to be the necessary attributes of autonomous agency, as well as propose a hierarchy of agents in terms of the attributes they possess, in *Section 4.*

3 A Systems Approach To Agents

Most approaches to classifying various types of (natural as well as artificial) agents are based on specifying the necessary *attributes* of a particular kind of agents, as in, e.g., [10]. We adopt this general approach, as well. However, we also try to be more specific as to *what kinds of attributes* we allow. Tools from other cognitive disciplines, such as psychology, anthropology and sociology, have been liberally applied to characterize the fundamental properties, and therefore the very nature, of various artificial agent systems. In particular, software and robotic agents have been generously ascribed properties that characterize anthropomorphic cognition, such as beliefs, desires, intentions, emotions, etc. One of the most successful examples of such approach are the BDI agent paradigm and architectures [16].

However, we see some potential conceptual and practical problems with assigning too liberally human (cognitive or other) attributes to a piece of software or a robot. In scientific and engineering modeling, the very purpose of a *model* is to be intrinsically simpler, and therefore more amenable to analysis, than the phenomenon being modeled. But when the attributes of beliefs, intentions, emotions, and the like are ascribed to, for instance, a software agent system with individual agents of a fairly modest complexity, it seems that exactly the opposite is the case. While there is some justification in correlating, for instance, how artificial agents represent and interact with complex, partially observable environments and tasks to how humans act (reason, learn, represent knowledge, etc.) with respect to their tasks and environments, there are also certain dangers in this approach. For, after all, robots and software agents are not human, and (unless one believes in the Strong AI hypothesis [19]) perhaps cannot ever be made very human-like in terms of their cognitive capabilities. Furthermore, representing and reasoning about relatively simple software agents encountered in many software engineering applications in terms of highly complex capabilities of human-like cognition seems to be an "overkill", in that the complexity of the model may end up considerably exceeding the sophistication of the modeled.

Another problem with attributing various anthropomorphic features to artificial agents emerges once different types of such agents are compared and contrasted with one another. Software agents, robots and other types of artificial agents are man-designed engineering systems. They should be characterized, studied, compared and contrasted to one another in terms of how they as systems *behave,* not what "mental states" or "beliefs" or "desires" or "emotions" their designer attributes to them. Whether an agent is reactive or adaptable can be, in general, verified by an outside observer that is independent of the agent. What are the belief or desire or emotional states of an agent, on the other hand, cannot.

We shall propose in the sequel a less cognition-oriented, and less anthropomorphic, approach to modeling, classifying and understanding various types of (artificial) autonomous agents and multi-agent systems (MAS). In particular, our approach, instead of cognitive psychology, draws more analogies and paradigms from cybernetics [25] and systems science [14, 15] on one, and biology and natural evolution [9], on the other hand. We argue that this approach fairly naturally leads to various possible hierarchical classifications of autonomous agents, and we propose one such general and broad agent hierarchy.

In particular, instead of comparing various agents in terms of their sophistication by chiefly comparing the complexities of agents' internal representations or "logics", we adopt a cybernetics-inspired approach based on the "black box" abstraction, and consider what kind of properties an agent needs in order to be able to do certain things, or function a certain way. We view an agent system "not a thing, but a list of variables" [5] and relations among those variables. Moreover, to understand an autonomous agent, one has to also understand this agent's environment, as well as various loops (e.g., feed-forward or feedback) that determine how this agent interacts with its environments. Thus our emphasis is on a functionalist, behavioral aspects of agency, and an agent is viewed as a black box whose inner structure (such as beliefs, desires, emotions, etc.) may or may not be accessible or understood, but it is *the interaction of this black box system with the outside world,* mechanisms for that interaction, and observable behavioral consequences of that interaction that are given the "first class" status (see, e.g., [4, 5]).

4 An Agent Hierarchy: From Reactive Towards Deliberative

We now discuss in some detail what are the critical, *necessary* (as opposed to optional) attributes that characterize most known autonomous agents, biological and computational alike. The most elementary attributes of such agents can be expected to be those properties that characterize any *autonomous system* in general. Once a couple of additional attributes that characterize virtually all agents are added, we arrive at a *weak* notion of autonomous agency. Subsequently, some additional properties will be identified that, we argue, characterize nearly all autonomous agents found in AI and DAI. An agent that possesses each of these attributes, as well as those of weakly autonomous agents, we shall call *strongly autonomous.* Finally, one more property will be identified that is absolutely necessary for any (however weak) notion of intelligence. Thus this list of system properties, each to at least some degree observable or testable by an observer external to the system, will implicitly define a tentative *natural hierarchy* of autonomous agents. In addition to similar attempts at classifying various types of agents (e.g., [10, 26]), our approach is also motivated by the general systems theory, and, in particular, by epistemological hierarchies of *(general) systems,* as in, e.g., [15].

The minimal notion of autonomy is the requirement that an entity (at least partially) controls its own internal state. Some degree of control[1] of one's internal state indeed appears necessary for autonomous agency, as well - but it is by no means sufficient. In addition to control over its *internal state,* an autonomous system ought to have at least some degree of control over its *behavior.* In case of a computer program (that is, a software agent), this means autonomous execution. If some autonomous control of a software system's state and execution were all it takes for such a system to be an autonomous agent, then the distinction between software agents and arbitrary computer programs would be rather blurred, and *(almost) all* programs would "qualify" for autonomous agents (see, e.g., discussions in [10, 17]). This is clearly undesirable. The question arises, what is missing - what additional requirements need to be imposed on an arbitrary computer program so that such a program can be considered a legitimate software agent?

Agents cannot be understood in isolation from the environment in which they are embedded [10]. This implies that, in order to develop a meaningful model of an agent, we need (a) an appropriate model of the environment, and (b) a model of the agent's *interaction* with the environment.

Regardless of the nature and mechanisms of this interaction between an agent and its environment (where the environment may also include other agents), there would be no point to any such interaction if it were not able to *affect* either the agent, or the environment outside of the agent, or, most often in practice, *both.*

Consequently, we consider *reactivity* (or what is called *"responsiveness"* in [26]) to be another necessary attribute of any notion of autonomous agency, as the agent has to be able to (1) notice changes in the environment, (2) appropriately respond to those changes, and (3) affect what input or stimuli it will receive from the environment in the future. Hence, the necessary attributes for any reasonable notion of autonomous agency identified thus far are (i) some control of one's internal state and execution, and (ii) reactivity as a prerequisite for the agent-environment interactions that, in general, may affect both the agent and the environment.

Any "proper" computational or biological autonomous agent can also be expected to be at least somewhat *persistent,* that is, to "live on" beyond completing a single task on a single occasion. In case of software agents, persistence makes an agent different from say a *subroutine* of a computer program whose "turning on and off" is controlled from outside of that subroutine (see, e.g., [10]). This necessity of some form of persistence is evidently strongly related to the most basic requirement of *(weakly) autonomous agency,* namely, that an agent ought to have some degree of control of its internal state and behavior.

[1] Full and exclusive control of one's internal state, if understood in the sense of that "nothing from the outside" can affect the entity's state, is clearly not desirable in case of agents, as one would like the agent to be able to be effected by its environment.

We summarize below our notion of *weakly autonomous agency (WAA)*[2] in terms of the necessary agent attributes:

$$\textit{weak autonomous agency} \approx \textit{control of own state} + \textit{reactivity} + \textit{persistence}$$

Hence, at the bottom level of the emerging hierarchy of autonomous agents, we find purely reactive embedded (or situated) agents [17]. Such agents can be appropriately abstracted via finite state machines (deterministic case) or discrete Markov chains (probabilistic case). A combination of reactivity and persistence characterizes many of both the simplest life forms and simple artificial agents. When some degree of control of the agent's internal state and behavior is also present, one arrives at what we shall call herewith *weakly autonomous agency (WAA).* We suggest the actor model of distributed computing [1, 2] to be a canonical example of *WAA* among the software agents.

One common feature found in all or nearly all interesting autonomous agents, biological and computational alike, is some form of *goal-orientedness* or *goal-drivenness.* In case of the living organisms, the highest level driving mechanisms are the instincts of *survival* and *reproduction.* The single most fundamental instinct in all of known life forms (to which an appropriate notion of an instinct can be ascribed at all) is that of survival. Indeed, the instinct of reproduction is related to the survival of the species or, perhaps, of the particular genes and gene patterns, as opposed to the "mere" survival of the individual organisms [9]. At lower levels, the driving mechanisms - finding food or a sexual partner - are those that are expected to provide, promote and enhance the two highest-level goals, survival and reproduction.

In the case of artificial computational agents such as a web crawler or a robot or an autonomous unmanned vehicle, these agents are designed and programmed with a particular goal or a set of goals in designer's mind. Thus, the ability to act autonomously is typically related to an agent having some goal(s) to accomplish, and therefore being goal-driven.

From a systems perspective, in order for an agent to be reactive, it has to be *coupled* to its environment via some appropriate sensors (or "input channels") and effectors ("output channels"). Due to agent's sensors, the environment can affect the agent; due to agent's effectors, the agent can affect the outside environment. For a stronger notion of agency than *WAA,* in addition to some sort of sensors and effectors, necessary to ensure that the agent can interact with, affect and be affected by the outside world, it seems natural that an appropriate feedback, or control, loop exists between the agent and the outside world, so that this feedback loop affects how the agent responds to the environmental changes. A feedback loop provides the agent with knowledge of "how well it is doing". In particular, an agent will have use of a feedback loop only if it has an appropriate notion of its goals or tasks, and an evaluation function (task value, utility, resource consumption, or the like) associated with it. That is, an agent needs some sort of a *control loop* in order to be capable of goal-oriented or utility-oriented behavior.

Finally, in addition to responsiveness, persistence and goal-drivenness or goal-orientedness, one more characteristic found in nearly all interesting autonomous agents, not altogether unrelated to goal-orientedness, is that of *pro-activeness* [10, 18, 26]. While some literature on autonomous agents treats pro-activeness and goal-drivenness as synonyms, we briefly discuss why, in general, the two attributes ought to be distinguished.

Namely, a situated, reactive agent can be goal-oriented without being pro-active: given an input from the "world", the goal-oriented agent acts so as to ensure, e.g., avoiding being in certain of its internal states that it views incompatible with its limited knowledge of "what is going on out there". If there are no changes in the environment, the agent simply keeps "sitting" in whatever its current state happens to be. Thus, this reactive agent has a goal (although admittedly a very simplistic one), but is not pro-active. Similarly, an agent can be pro-active without being goal-oriented, as long as we require of agent's goal(s) to be non-trivial, and, in particular, to possibly entail some deliberate effect that the agent's actions may be required, under appropriate circumstances, to have on the environment. Under this assumption, an agent may "pro-actively" perform a more or less random walk among its internal states, without any observable effects on the outside world, and therefore without accomplishing - or, indeed, having - any specific goals insofar as the agent's deliberate influence on the environment.

Thus, while pro-activeness and goal-orientedness are usually closely related, they are not synonymous, and, moreover, neither subsumes the other.

Once a *WAA* agent is additionally equipped with some form of goal-drivenness and pro-activeness, we arrive at what we define as *strongly autonomous agency (SAA).* Most agents encountered in AI, whether they are software agents, robots, unmanned vehicles, or of any other kind, are of this, strongly autonomous type (see, e.g., [18, 16, 24]).

Therefore, we find that it is precisely the properties of (i) some degree of control of one's own internal state and behavior, (ii) reactiveness or responsiveness, (iii) persistence, (iv) pro-activeness, and (v) goal-drivenness or goal-orientedness that, together, and in synergy with each other, make an agent *truly (or strongly) autonomous* in an AI sense:

$$\textit{strong autonomous agency} \approx \textit{weak autonomous agency} + \textit{goal-orientedness} + \textit{pro-activeness}$$

Granted, much of the agent literature has identified properties (i) - (v) as common to autonomous agents (see, e.g., [26, 10] and references therein). We claim, however, that these five agent capabilities are *the necessary* properties that are all found in nearly every reasonable model of autonomous agency, whereas other characteristics, including sociability, mobility, "mental states", beliefs-desires-intentions, etc., are not as essential, and are found in (or can be reasonably attributed to) only *some,* but by no means

[2] This notion of weak agent autonomy, mainly based on the dominant notion of autonomous agents in the area of software design for open distributed systems, is obviously not called *weak* by those who consider it sufficient for their purposes.

(nearly) all of the known autonomous agents, whether biological or artificial.

However, even those living organisms that one would never consider intelligent have one more fundamental property, absent from our notion of *SAA,* and that is the ability to *adapt* (e.g., through metabolism). Adaptability is a necessary prerequisite for biological survival, as well as for any reasonable notion of intelligence. The "control loop" between the agent and the world serves no purpose, if the agent has no goals or notions of "goodness" with respect to which it tries to optimize its behavior. But such goal- or utility-drivenness is useless, if the agent cannot dynamically adjust its behavior based on the feedback, i.e., if it cannot adapt.

To summarize, based on how is an agent coupled to its environment, how complex properties of that environment the particular type of coupling (e.g., type of sensors, "control loop", effectors) can capture, and how complex behaviors or actions the agent is capable of, we have proposed a tentative general hierarchical classification of autonomous agents embedded in, and acting as a part of, their environments.

Whether one would consider an agent that (i) has at least some control over its internal state and behavior, and is (ii) reactive, (iii) persistent, (iv) pro-active, (v) goal- or utility-driven, and (vi) adaptable, to automatically be *intelligent,* depends on one's definition of intelligence and is subject to debate. What seems clear, however, is that no proper subset of the properties (i) - (vi) satisfies even the weakest notion of intelligence. Moreover, we argue that, as one keeps adding on properties from (i) towards (vi), one can recognize many well-known examples of agents found in the literature, yet not as a part of what we argue is a reasonable and natural hierarchy of agents[3]. For instance, artificial agents that possess only (i), (ii) and possibly (iii) are studied in detail in [17]. Some examples of *WAA* agents possessing (i) - (iii) and *SAA* agents having attributes (i) - (v) are discussed next.

5 Discussion and Some Applications

To illustrate the usefulness of the proposed hierarchy of agents, we consider some software engineering developments in the context of open distributed systems. Agent-oriented programming [20] can be viewed both as a novel paradigm and the natural successor to object-oriented paradigm [13]. The transition from object-oriented towards agent-oriented programming was motivated by the design of open distributed platforms, so that concurrency and resource sharing can be exploited in heterogeneous distributed environments [3].

To place the development of a general paradigm of autonomous agency into a broader computer science perspective, we briefly make a comparison to the development of the object-oriented paradigm. The primary motivation for moving away from function evaluation based classical imperative programming towards the object-oriented programming paradigm was primarily motivated by the nature of a great number of emerging applications, where it was more natural to think in terms of objects and their classes and hierarchies, their capabilities ("methods"), etc., then in terms of functions being evaluated on variables. One particular domain that gave a huge impetus to the growth and success of object-oriented programming was that of computer simulation of various complex and distributed infrastructures [8]. As computing started becoming increasingly distributed both physically and logically, and these distributed systems getting increasingly heterogeneous, complex and open, the individual components, whether hardware or software, were moving away from non-autonomous components of a single, tightly coupled system, towards being increasingly sophisticated, autonomous and complex (sub)systems themselves, that were only loosely coupled into an overarching larger system. Hence, a novel paradigm capturing the increasing requirements in terms of autonomy, flexibility and complexity of the individual components in such distributed systems was sought - and, in case of the software, the *agent-based programming* paradigm was born [20].

We thus see the relationship of agent-oriented programming to object-oriented programming, in essence, similar to the relationship of object-oriented to classical imperative programming: each is a novel metaphor and a radical departure from its predecessor - yet a novel metaphor that clearly builds on the top of its predecessor, and adds more desirable properties that brings thus enhanced model considerably closer to the target applications.

Actors [1, 2] are a powerful model for specifying coordination in open distributed systems. In addition to its internal state, an actor also encapsulates its behavior (both data and procedure). An actor can communicate via asynchronous message passing with other actors; this asynchronous communication is, therefore, central to how actors interact with their environment. An actor is responsive or reactive; it also may (but need not) be persistent. Actors thus fit well into our concept of *weakly autonomous agency.*

Actors can also be used as a building block towards implementing more complex systems and, in particular, for software design of autonomous agents with stronger autonomous capabilities, via appropriate extensions (added functionality) of the basic actor model. For instance, ref. [3] addresses an important problem of how to extend actors into a powerful concurrent programming paradigm for *distributed artificial intelligence (DAI)* [6, 24]. There are other examples of designing strongly autonomous applications on the top of weakly autonomous infrastructures. For instance, an actor-based (hence, "weakly autonomous") software infrastructure is used in [11, 12] to build a simulator of a particular kind of strongly autonomous agents, namely, autonomous unmanned vehicles [22, 23]. The basic agent capabilities are provided by the actor infrastructure, whereas the higher-order autonomous abilities, such as the pro-active pursuit of an agent's goals or the agents' coordination strategies, are built on the top of the basic actor architecture, i.e., at the "application software" level [11, 22, 23].

[3] However, see [10] for another proposal of a hierarchical taxonomy of various types of agents.

6 Conclusions

The subject of this paper are *autonomous agents* from a systems perspective. First, we survey some relevant literature and offer some general thoughts about various properties and notions of autonomous agents. We then propose a hierarchy of autonomous agents based on the complexity of their behaviors, and the necessary attributes that can yield particular behaviors. Instead of marking various types of agents as more or less complex in terms of the sophistication of their (supposed) mental or emotional states, we make distinction in terms of the basic agent capabilities whose presence - or lack thereof - is readily observable and measurable by an observer outside of the agent itself. Thus, instead of a "cognitive" or symbolic AI approach, we propose classifying autonomous agents in more behaviorist, functionalist and systems theory terms. In particular, we identify the three absolutely necessary properties for even the weak(est) notion of autonomous agency, and three additional, more advanced properties that are necessary for an agent, whether biological or artificial, to be reasonably considered deliberative or intelligent. We also show how some well-known, existing agent models fit into the appropriate layers of our proposed agent hierarchy. Finally, we point out some examples of how the lower-level agents can be used as "building blocks" in design of more complex, higher-level autonomous agents and *MAS*.

Acknowledgments: This work was supported by the *DARPA IPTO TASK Program*, contract *F30602-00-2-0586*.

References

[1] G. Agha, "Actors: A Model of Concurrent Computation in Distributed Systems", MIT Press, Cambridge, MA, 1986

[2] G. Agha, "Concurrent Object-Oriented Programming", in *Communications ACM,* vol. 33 (9), 1990

[3] G. Agha, N. Jamali, "Concurrent Programming for DAI", in *G. Weiss (ed.), "Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence", The MIT Press, 1999*

[4] W. Ross Ashby, "Design for a Brain", Wiley, 1960

[5] W. Ross Ashby, "An Introduction to Cybernetics", 4th impression, Chapman & Hall Ltd., London, 1961

[6] N. M. Avouris, L. Gasser (eds.), "Distributed Artificial Intelligence: Theory and Praxis", Euro Courses Comp. & Info. Sci. vol. 5, Kluwer Academic Publ., 1992

[7] M. E. Bratman, "Intentions, Plans and Practical Reason", Harvard Univ. Press, Cambridge, MA, 1987

[8] T. Budd, "An Introduction to Object-Oriented Programming", Addison-Wesley, 1991

[9] R. Dawkins, "The Selfish Gene", Oxford Press, 1990

[10] S. Franklin, A. Graesser, "Is it an Agent, or just a Program? A Taxonomy for Autonomous Agents", Proc. 3rd Int'l Workshop on Agent Theories, Architectures & Languages, Springer-Verlag, 1996

[11] M. Jang, S. Reddy, P. Tosic, L. Chen, G. Agha, "An Actor-based Simulation for Studying UAV Coordination", Proc. 15th Euro. Symp. Simul. (ESS '03), Delft, The Netherlands, October 2003

[12] M. Jang, G. Agha, "On Efficient Communication and Service Agent Discovery in Multi-agent Systems," 3rd Int'l Workshop on Software Engineering for Large-Scale Multi-Agent Systems (SELMAS '04), Edinburgh, Scotland, May 2004

[13] N. R. Jennings, "An agent-based approach for building complex software systems", *Communications of the ACM* vol. 44 (4), 2001

[14] G. J. Klir, "Facets of Systems Science" (2nd ed.), Plenum Press, New York, 2001

[15] G. J. Klir, "Systems Science", Encyclopedia of Information Systems, Vol. 4, Elsevier Sci. (USA), 2003

[16] A. S. Rao, M. P. Georgeff, "BDI Agents: From Theory to Practice", Proc. 1st Int'l Conf. on MAS (ICMAS'95), San Francisco, USA, 1995

[17] S. J. Rosenschein, L. P. Kaelbling, "A Situated View of Presentation and Control", *Artificial Intelligence* vol. 73, 1995

[18] S. Russell, P. Norvig, "Artificial Intelligence: A Modern Approach", 2nd ed., Prentice Hall, 2003

[19] John Searle, "Minds, Brains and Science", Harvard University Press, Cambridge, 1984

[20] Y. Shoham, "Agent-oriented programming", *Artificial Intelligence* vol. 60, 1993

[21] H. A. Simon, "Models of Man", J. Willey & Sons, New York, 1957

[22] P. Tosic, M. Jang, S. Reddy, J. Chia, L. Chen, G. Agha, "Modeling a System of UAVs on a Mission", Proc. SCI'03, Orlando, Florida, July 2003

[23] P. Tosic, G. Agha, "Modeling Agent's Autonomous Decision Making in Multi-Agent, Multi-Task Environments", Proc. 1st Euro. Workshop on *MAS* (EUMAS'03), Oxford, England, December 2003

[24] G. Weiss (ed.), "Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence", The MIT Press, Cambridge, MA, 1999

[25] N. Wiener, "Cybernetics", J. Willey, New York, 1948

[26] M. Wooldridge, N. Jennings, "Intelligent Agents: Theory and Practice", Knowledge Engin. Rev., 1995