MULTIAGENT SYSTEMS

ANDREA OMICINI

Dipartimento di Elettronica, Informatica e Sistemistica – Università degli Studi di Bologna

AGOSTINO POGGI

Dipartimento di Ingegneria dell'Informazione – Università degli Studi di Parma

1 Introduction

While research on multiagent systems (MAS) mostly develops from the DAI (Distributed ArtiÞcial Intelligence) Þeld, it also liberally draws from the results of a number of different research areasÑ from classical AI to robotics, from distributed & mobile systems to programming languages. Resulting from this multifaceted sources of inspiration and knowledge, an essentially and widely incoherent acceptation of the notion of agent has not precluded an impetuous development of MAS research and applications in the last 15 years. As a weak notion, an agent is mainly to be conceived as an abstraction encapsulating control, along with a criterion (a goal, a task) to govern control in autonomy: so, an agent might either be mobile or not; it might either be task-oriented or goal-orientedÑ but it is always *autonomous*. As a strong notion, an agent is a cognitive entity capable of representing, understanding, and possibly changing the world in which it lives, of rationally planning its actions and deliberating, and of intelligently pursuing its objectives.

Recognition that computational systems are mostly distributed systems pushed research on agent models and technologies from a single agent to a multiagent setting. This emphasised two of the main sources of inspiration for contemporary MAS research: *sociality* of agents, and their situatedness in a given *environment*. So, agents never live aloneÑ instead, they coexist in MAS as complex and articulated organisations. And, agents of a MAS are not conceived, designed, built in the vacuumÑ instead, their features, abilities and behaviour are strictly related to the environment where they live. Historically, MAS were initially conceived as mostly homogeneous systems, made of cooperative agents; but the expressive power of agent-oriented abstractions has soon promoted a more general vision of MAS as open, heterogeneous systems, where also self-interested, competitive agents could play a relevant role.

As agents are social entities, MAS research devotes particular emphasis to those aspects that concern interaction between agents in a MASÑ such as *communication* between agents, MAS *coordination* and *cooperation*Ñ and other social aspectsÑ such as *trust* [23], reputation, coalition formation, and the like. Activity on MAS communicationÑ which has also produced the only meaningful standardisation effort in the MAS Þeld, FIPA [38]Ñ focuses on agents as speaking entities, interacting with other agents through speech acts. Activity on MAS coordination, cooperation, and the like, focuses more on the *global view* of MAS as super-organisms with their own structure, behaviour and goalsÑ which need to be suitably related to the individual structure, behaviour and goals of the agents composing MAS [22, 59].

As agents are situated entities, research concentrates on MAS environment [72]: how to model and build it, and how to model and regulate agent interaction within the environment. Environment representation (of both its status and dynamics) is then a twofold issueÑ for the agent that has to act on it, and for the engineer that has to construct it, by deÞning its structure and behaviour. Models for agent actions within a MAS, and abstractions to structure MAS environment and govern agent-environment interactions are today subject of paramount interest for MAS research.

While agent research is quite pervasive, both in terms of scientiÞc disciplines and in terms of world-wide geographic distribution, in the last ten years the European lead has emerged quite clearly, in particular thanks to many EC funded projects, from AgentCities to AgentLink [50]. In these contexts, contributions from Italian researchers have always be relevant, and have produced many advances and results. The liveliness and the scientiÞc level of the Italian community working on MAS is also witnessed by the activity of MAS-AIIAÑ the AI*IA working group on Agents and Multiagent Systems–which was founded by Cristiano Castelfranchi several years ago (and led by the authors of this paper in the last eight years), and counts some hundreds of participants from both Italian and European universities, research institutes and industries [30].

2 Relation with other research areas

Today, results from AI and MAS research are no longer so easily distinguishable. Agents and MAS have become the introductory metaphors to most of the AI results, both for researcher and for industryÑ as exempliÞed by one of the most commonly used AI textbooks [66]. So, classic AI results on planning, practical reasoning, knowledge representation, machine learning, and the like, have become the most obvious and fruitful starting points for MAS research and technologies. On the other hand, it is quite rare nowadays that new Þndings or lines of research in AI might ignore the agent abstractions at all. Altogether, rather than a mere subÞeld of AI, agents and MAS could be seen as promoting a new paradigm, providing a new and original perspective about computational intelligence and intelligent systems [76, 75].

Apart from this straightforward connection with AI, relationships between MAS and other research areas are wide and numerous. Among the computational disciplines, software engineering (SE) is likely to be mostly impacted by contributions from the MAS Þeld. In fact, agent-based abstractions induce different meta-models for computational systems, so have the potential to bring about new ways to conceive and design software systems. Starting from the pioneering work on Gaia [74], a number of different AOSE (agent-oriented software engineering) methodologies have appeared [10] that promise to deeply change the theory and practise of SEÑ a signiÞcant deal of which resulting from the work of Italian research groups [31, 41, 57].

Agent-based systems are also characterised by the fading boundaries between human and software agents. In the context of Þelds like CSCW [67] and WfMS [70, 16], for instance, distinctions between artiÞcial and human participants to shared activities are mostly blurred. This paves the way for contributions from many disciplines of human and natural sciences to profoundly inßuence MAS research. Findings from economics [25], social and organisational sciences [21], biology and ethology [60], among the many others, have been drawn and quite liberally re-used in the MAS Þeld. Also, it should be noted that this trend is not going to be just a fashion phenomenon, since it appears to be instead the result of the deep co-relation among the many disciplines dealing with complex systems.

A novel point of convergence between MAS and complex systems is represented by research on emergence, self-organisation, and self-* properties in general. In fact, the demanding requirements of scenarios like ubiquitous computing, dynamic supply chains and health care call for a level of ßexibility, adaptability and robustness that only MAS coordinated by self-organisation and emergence mechanisms apparently feature [36]. As a result, MAS are today the emerging computational paradigm for self-organising systems [15].

On the other hand, if today Òall science is computer scienceÓpossibly holds for computational disciplines in general [46], this seems particularly true for MAS research and technologies. In particular, a new deal of research is proposing MAS as a natural paradigm for simulation of complex systems of the most various sortsÑ where agent-based models and technologies are used to reproduce the laws, structure and behaviour of systems ranging from natural to artiÞcial, from biological to social ones. Apart from agent-based simulation infrastructures like RePast [56], many research efforts have been (successfully) devoted to system simulation: MAS have then been used for simulating the behaviour of crowds [7], the dynamics of biological systems [20], social systems & organisations [29], and many other sorts of complex systems. Also, it should also be noticed that the new deal of scientiÞc events in the MAS simulation Þeld has been promoted and organised by researchers coming from the Italian AI community, like Bandini, Petta & Vizzari [8], and Fortino, Cossentino & Russo [39].

3 State of the art

Initial efforts in the Þeld of MAS focused on the development of intelligent agent architectures, and the early years established several lasting styles of architecture. These range from purely reactive agents, which operate in a simple stimulus-response fashionÑ see, for example, the subsumption architecture [14]Ñ , to deliberative agents that reason about their actionsÑ see, for example, the belief-desire-intention (BDI) architecture [65]. Nowadays, BDI works as a sort of reference architecture for building agents of a MAS; however, many researchers are working on agent architectures to improve and extend the BDI model, and also to deÞne new models. In particular, one of the main goals is to deÞne agent architectures able both to make the construction of MAS simpler, and to implement efÞcient MAS to be used in real-world settings [2, 6, 53].

One of the key components in a MAS is communication. In fact, agents need to be able to communicate with users, with system resources, and with each other if they are to cooperate, collaborate, negotiate and so on. Therefore, a number of researchers focused on communication components for MAS and, in particular, on the deÞnition of a language for the communication between agents. Agent Communication Languages (ACL) rely on speech act theory [69] and are based on a separation between the communicative acts and the content language. Currently the most used and studied agent communication language is the FIPA ACL [38], whose main features are the possibility of using different content languages and the management of conversations through predeÞned interaction protocols. However, some researchers proved the limits of this languages and are working on the improvement of this language to provide alternative semantics, new ontological supports, and also new content languages [37, 71, 5].

MAS infrastructures and development tools are an important component that can affect the diffusion of the use of agent technologies in the different application domains. The work towards standards for agents interoperability made from the middle of nineties by FIPA boosted the study and development of such tools: now, as a result, a large number of both open source and commercial agent development environments and toolkits are availableÑ see, for example, JADE [9] and JACK [73]. In particular, it is here worth noting that JADEÑ today the most used agent-oriented platform worldwideÑ is a remarkable product of the activity of many Italian researchers. In the last years, research mainly focused on enhancing the most widely used development tools with new features, aimed at simplifying software development, as well as to extend their use in other application domains. In particular, a number of researchers are working in: *(i)* the development of tools for bridging agent technologies with both Web services and Semantic Web technologies [54], *(ii)* the deÞnition of agent programming layer on the top of the most known peer-to-peer middleware [12], and *(iii)* the introduction of the most sophisticated security techniques in the MAS architectures [64].

Infrastructures also play a key role in the area of MAS coordination, while research on cooperation mostly deals with individual agents sharing goals, and acting toward their fulÞlmentÑ and often neglects infrastructural issues. On the one hand, cooperation in MAS largely draws from the results of Þelds like CSCW, CIS (cooperative information systems) and the work on cooperative information agents, and mostly promotes a view over MAS where social behaviour result from the composition of individual cooperative behaviours (and attitudes). On the other hand, literature on MAS coordination freely reuses Þndings from a wide spectrum of research Þelds of many sorts, from DAI to parallel & distributed computing, from robotics to mobile computing. While most of the work on agent coordination adopts the *subjective* viewpoint on coordinationÑ where agents are the coordinating entitiesÑ , a good deal of the efforts in the last years concentrate on the *objective* aspectsÑ where agents are the coordinated entities [68]. A large number of contributions from Italian researchers insist exactly on the latter aspects [17, 28], and have also produced several well-known coordination technologies, ranging from Þeld-based [51] to tuple-based [55, 61] coordination infrastructures.

Among the wide spectrum of MAS technologies that are today under study and development, agent-oriented programming languages are today at the core of many research efforts. After some initial efforts devoted to the design of agent programming languagesÑ embodying the agent abstraction at the language level, and focussing on individual agent programmingÑ activity has then extended to cover all the many aspects of MAS that require speciÞc language technologiesÑ thus gearing research toward the deÞnition of a whole range of new classes of programming languages for MAS [33].

Finally, the impact of agent-oriented approaches on the engineering of complex computational systems is not limited to technologies, or methodologies. A important portion of AOSE research in the last years has in fact been devoted to Þnd out suitable meta-models, which could allow researchers to build up a coherent conceptual framework, for the numerous, sparse and loosely-related results on MAS technologies, infrastructures, methodologies and tools [11].

4 Applications

MAS are being used in an increasingly wide variety of applications, ranging from comparatively small systems for personal assistance to open, complex, mission-critical systems for industrial applications [44].

Industrial applications are very important for MAS because they represent the Þeld where the MAS techniques were Þrst experimented, and where they Þrst showed their huge potential. Today, MAS are used for a number of different industrial applications: in particular, they are employed in application scenarios like process control [45], system diagnostics [1], manufacturing [62] and network management [13], whose distributed nature easily falls within the reach of MAS techniques.

One of the Þrst and most important application Þelds for MAS is information management [35]. In particular, the Internet has been described as an ideal domain for MAS, given its distributed nature and the sheer volume of information available that make the use of agents of great interest for searching and Þltering the information [47]. Internet has also pushed the use of MAS technologies in the Þelds of commerce and business process management. In fact, before the spreading of the Internet, commerce and business process management are almost entirely driven by human interactions; humans decide when to buy goods, how much they are willing to pay, and so on. Today, electronic commerce and automated business processes have increasingly assumed a pivotal role in many organisations because they offers opportunities to signiÞcantly improve the way in which the many entities involved in the business process interact. In this scenario, MAS have been shown both to be suitable for the modelling and the design of business process management systems [19], and to be amenable to work as key components for the automation of some or all the steps of these processes [43].

The distributed nature of trafÞc and transport processes, along with the strong independence among the entities involved in such processes, have made MAS a key tool for the engineering of effective, real-world applications for both trafÞc management and transport logistics [34]. Different applications have been already realised; in particular, one of themÑ OASIS [49]Ñ can be considered as the proof that MAS are the ideal means for building open, complex, mission-critical systems. OASIS is in fact a sophisticated agent-based air-trafÞc control system based on the BDI agent model, which was used with success at the Sydney airport in Australia.

Other interesting MAS applications can be found in the health care domain [52]. There, in fact, MAS have already been proposed to deal with many different kinds of problems such as patient scheduling and management, senior and community care, medical information access and management, and decision support[26]. Also, the applications implemented till now already shown that MAS are likely to be the right solution to build up medical decision support systems [42], and to improve the coordination between the different professionals involved in the health care processes [48].

The area of computers & law is another hot Þeld for MAS [24]. There, on the one hand, MAS technologies are becoming essential tools for the application and integration of traditional AI techniques (such as expert systems, automated deduction, deontic logic) within computational systems supporting techno-legal activities [27]. On the other hand, the complexity of agent-based systems proposes novel issues to the analysis of the legislative framework, such as the possible law-abidingness of software agents.

Italian research groups already gave, and are still giving, an important contribution to prove MAS methodologies and technologies as beneÞcial for many complex application scenarios. In fact, many systems could be mentioned in all the application domains where MAS are up to now applied, where the contribution of Italian groups and researchers is at least relevant. Even more, some of them are today considered among the most relevant systems in such application domains: for example, in trafÞc control [32], information management [18, 40], health care [48, 3], decision support systems [63] and bioinformatics [4].

5 Research perspectives and potential applications

At present, MAS are widely studied in both academic and industry laboratoriesÑ but they are not yet widely available in real-world applications. However, many studies foresee an evolution of MAS technologies and the extension of their use to the large part of computing based applications [75]. For example, AgentLink researchers [50] expect an evolution of MAS in four phases that moves from systems designed by one design team for one corporate environment and based on agents sharing common high-level goals in a single domain (phase 1), to systems designed by the same team to cross corporate boundaries, based on agents with fewer goals in common, that act in a common domain interacting through standard agent communications languages, but by using non-standard interaction protocols (phase 2), to systems based on heterogeneous agents designed by different teams on the basis of publicly-stated requirements and standards (phase 3), to, Þnally, open systems spanning over multiple application domains, and involving heterogeneous agents developed by diverse design teams (phase 4).

Therefore, one of the main goals of MAS research in the next ten years will be oriented to realise truly scalable open and distributed MAS. This will be possible only in case some important advances will be also achieved in some other research areas as, for instance, knowledge representation, machine learning, planning and reasoning. Another important factor will likely be the integration with Web technologies, which will promote the advent of MAS as a fundamental means for both information management and service composition applications. Also, the availability of advanced tools for the development of agent-based software and the strengthening of agent interoperability standards will be one of the key features to encourage the adoption of MAS by an increasing number of software developers. Finally, the possibility of MAS techniques to keep on working as well-founded integrators, and also as vehicles toward real-world application scenarios, for the most advanced results from such a huge number of heterogeneous research areas (from AI to Grid computing, from wireless sensor networks to pervasive computing) will probably be another key factor in establishing MAS as the fundamental paradigm for the engineering of complex artiÞcial systems [75].

Italian researchers in the MAS Þeld are already working along that direction. In particular, we can easily devise out at least three different areas where their contributions could result essential: i) the development of new AI techniques, allowing agents to act and deliberate intelligently and effectively in real-world open environments, and their integration in agent architectures and software technologies; ii) the integration of MAS with Web technologies and service-oriented architectures, and iii) the realisation of advanced methodologies and tools for the design, development, deployment, monitoring and maintenance of complex agent-based software systems.

6 Conclusion

Research on MAS has already a long history behind, but also several stimulating perspectives for its future. While MAS research is well-rooted in the AI tradition, it has developed along a number of different and heterogeneous lines, often crossing the (often vague) boundaries of computational sciences and AI. Today, there is a lot of excitement about MAS technologies, and their prospective applications; and Italian researchers in the Þeld are expected to keep on providing their precious contributions.

REFERENCES

[1] M. Albert, T. Laengle, H. Woern, M. Capobianco, and A. Brighenti. Multi-agent systems for industrial diagnostics. In *5th IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes*

(*SAFEPROCESS 2003*, pages 483–488, Washington, WA, USA, 9–11 June 2003.

[2] D. Ancona, V. Mascardi, J. F. Hubner, and R. H. Bordini. Coo-AgentSpeak: Cooperation in AgentSpeak through plan exchange. In N. R. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)* , volume 2, pages 696–705, New York, NY, USA, 19–23 July 2004. ACM.

[3] L. Ardissono, A. D. Leva, P. G, S. M, and M. Sonnessa. Adaptive medical workßow management for a context-dependent home healthcare assistance service. *Electronic Notes in Theoretical Computer Science*, 146(1):59–68, 2006.

[4] G. Armano, G. Mancosu, and A. Orro. Using multiple experts for predicting protein secondary structure. In *Artificial Intelligence and Applications*, pages 451–456, 2005.

[5] M. Baldoni, C. Baroglio, A. Martelli, and V. Patti. VeriÞcation of protocol conformance and agent interoperability. In F. Toni and P. Torroni, editors, *Computational Logic in Multi-Agent Systems* , volume 3900 of *LNAI*, pages 265–283. Springer, 2006. 6th International Workshop (CLIMA VI) London, UK, June 27-29, 2005, Revised Selected and Invited Papers.

[6] M. Baldoni, L. Giordano, A. Martelli, and V. Patti. Programming rational agents in a modal action logic. *Annals of Mathematics and Artificial Intelligence*, 41(2-4):207–257, 2004.

[7] S. Bandini, M. L. Federici, and G. Vizzari. A methodology for crowd modelling with situated cellular agents. In F. Corradini, F. De Paoli, E. Merelli, and A. Omicini, editors, *WOA 2005: Dagli Oggetti agli Agenti. 6th AI*IA/TABOO Joint Workshop "From Objects to Agents": Simulation and Formal Analysis of Complex Systems* , pages 91–98, Camerino, MC, Italy, 14–16 Nov. 2005. Pitagora Editrice Bologna.

[8] S. Bandini, P. Petta, and G. Vizzari. *International Symposium on Agent Based Modeling and Simulation (AbModSim 2006)*. University of Vienna, 18th European Meeting on Cybernetic Science and Systems Research (EMCSR 2006), Vienna, Austria, Apr. 18 2006.

[9] F. Bellifemine, A. Poggi, and G. Rimassa. Developing multi-agent systems with a FIPA-compliant agent framework. *Software, Practice and Experience*, 31(2):103–128, 2001.

[10] F. Bergenti, M.-P. Gleizes, and F. Zambonelli, editors. *Methodologies and Software Engineering for Agent Systems: The Agent-Oriented Software Engineering Handbook*. Kluwer Academic Publishers, June 2004.

[11] C. Bernon, M. Cossentino, and J. Pavon. An overview of current trends in european AOSE research. *Informatica*, 29(4):379–390, Nov. 2006.

[12] D. Bertolini, P. Busetta, A. Molani, M. Nori, and A. Perini. Designing peer-to-peer applications: An agent-oriented approach. In R. Kowalczyk, J. P. Muller, H. TianÞeld, and R. Unland, editors, *Agent Technologies, Infrastructures, Tools, and Applications for E-Services*, volume 2592 of *LNCS*, pages 92–106. Springer, 2003.

[13] A. Bieszczad, T. White, and B. Pagurek. Mobile agents for network management. *IEEE Communications Surveys*, 1998.

[14] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[15] S. A. Brueckner, G. Di Marzo Serugendo, D. Hales, and F. Zambonelli, editors. *Engineering Self-Organising Systems*, volume 3910 of *LNAI*. Springer, 2006. 3rd International Workshop (ESOA 2005), Utrecht, The Netherlands, 26 July 2005, Revised Selected Papers.

[16] P. A. Buhler and J. M. Vidal. Towards adaptive workßow enactment using multiagent systems. *Information Technology and Management*, 6:61–87, 2005.

[17] N. Busi, P. Ciancarini, R. Gorrieri, and G. Zavattaro. Coordination models: A guided tour. In A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors, *Coordination of Internet Agents: Models, Technologies, and Applications*, chapter 1, pages 6–24. Springer-Verlag, Mar. 2001.

[18] G. Cabri, F. Guerra, M. Vincini, S. Bergamaschi, L. Leonardi, and F. Zambonelli. MOMIS: Exploiting agents to support information integration. *International Journal of Cooperative Information Systems* , 11(3):293–314, 2002.

[19] L. Camarinha-Matos and H. Afsarmanesh. Virtual enterprise modeling and support infrastructures: applying multi-agent system approaches. In J. Carbonell and J. Siekmann, editors, *Multi-agents systems and applications*, pages 335–364. Springer-Verlag, New York, NY, USA, 2001.

[20] N. Cannata, F. Corradini, E. Merelli, A. Omicini, and A. Ricci. An agent-oriented conceptual framework for systems biology. *Transactions on Computational Systems Biology*, III:105–122, Dec. 2005. 4th International Workshop on NETwork Tools and Applications in Biology (NETTAB 2004), Camerino, MC,

Italy, September 5–7, 2004. Revised Selected and Invited Papers.

[21] C. Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103(1–2):157–182, Aug. 1998. Special Issue: ArtiÞcial Intelligence 40 Years Later.

[22] C. Castelfranchi. Engineering social order. In A. Omicini, R. Tolksdorf, and F. Zambonelli, editors, *Engineering Societies in the Agents World*, volume 1972 of *LNAI*, pages 1–18. Springer-Verlag, Dec. 2000.

[23] C. Castelfranchi and R. Falcone. Principles of trust for MAS: Cognitive anatomy, social importance, and quantiÞcation. In *3rd International Conference on Multi Agent Systems (ICMAS'98)*, pages 72–79. IEEE Press, 3-7 July 1998.

[24] C. Cevenini, editor. *The Law of Electronic Agents – Proceedings of the LEA'04 Workshop*, Roma, Italy, June 2004. Gedit Edizioni.

[25] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodrıguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, Jan. 2006.

[26] A. Ciampolini, P. Mello, and S. Storari. Distributed medical diagnosis with abductive logic agents. *AI*IA Notizie*, Anno XV(3):12–21, Sept. 2002.

[27] A. Ciampolini and P. Torroni. Using abductive logic agents for modelling judicial evaluation of criminal evidence. *Applied Artificial Intelligence*, 18(3-4):251–275, Feb. 2004. Special Issue ÒThe Construction of Judicial Proof: A Challenge for ArtiÞcial Intelligence ModellingÓ.

[28] P. Ciancarini, A. Omicini, and F. Zambonelli. Multiagent system engineering: The coordination viewpoint. In N. R. Jennings and Y. Lesperance, editors, *Intelligent Agents VI. Agent Theories, Architectures, and Languages*, volume 1757 of *LNAI*, pages 250—259. Springer-Verlag, 2000. 6th International Workshop (ATALÕ99), Orlando, FL, USA, 15–17July1999. Proceedings.

[29] R. Conte, B. Edmonds, S. Moss, and R. K. Sawyer. Sociology and social theory in agent based social simulation: A symposium. *Computational and Mathematical Organization Theory*, 7(3):183–205, Oct. 2001.

[30] A. Corradi, A. Omicini, and A. Poggi, editors. *WOA 2000 – Dagli oggetti agli agenti: tendenze evolutive dei sistemi software*, volume 1195 of *Atti di Congressi*, Parma, Italy, 29–30 May 2000. Pitagora Editrice Bologna. Proceedings.

[31] M. Cossentino, L. Sabatucci, and A. Chella. Patterns reuse in the PASSI methodology. In A. Omicini, P. Petta, and J. Pitt, editors, *Engineering Societies in the Agents World IV*, volume 3071 of *LNAI*. Springer, 2004. 4th International Workshop (ESAW 2003) London, UK, October 29-31, 2003, Revised Selected and Invited Papers.

[32] A. Cuppari, P. L. Guida, M. Martelli, V. Mascardi, and F. Zini. An agent based prototype for freight trains trafÞc management. In P. Larsen, editor, *FMERail'99 Workshop, satellite event of FM'99*, Toulouse, Trance, 1999.

[33] M. Dastani and J. J. Gomez-Sanz. Programming multi-agent systems. In *The Knowledge Engineering Review* [58], pages 151–164. Special Issue: Revised Reports from Selected Technical Forum Groups at the 1st & 2nd AgentLink III Technical Forums.

[34] P. Davidsson, L. Henesey, L. Ramstedt, J. Tornquist, and F. Wernstedt. *Agent-Based Approaches to Transport Logistics*. Whitestein, 2005.

[35] K. Decker and K.Sycara. Intelligent adaptive information agents. *Journal of Intelligent Information Systems*, 9(3):239–260, 1997.

[36] G. Di Marzo Serugendo, M.-P. Gleizes, and A. Karageorgos. Self-organisation and emergence in mas: An overview. *Informatica*, 30(1):45–54, Jan. 2006.

[37] A. Di Stefano, C. Santoro, G. Pappalardo, and E. Tramontana. Enforcing agent communication laws by means of a reßective framework. In H. Haddad, A. Omicini, R. L. Wainwright, and L. M. Liebrock, editors, *2004 ACM Symposium on Applied Computing (SAC)*, pages 462–468, Nicosia, Cyprus, 14—17 Mar. 2004. ACM, ACM.

[38] FIPA speciÞcations, 2002. http://www.Þpa.org.

[39] G. Fortino, M. Cossentino, and W. Russo. *Workshop on Multi-Agent Systems and Simulation (MAS&S 2006)*. University of Palermo, Industrial Simulation Conference 2006 (ISC 2006), Palermo, Italy, June 5 2006.

[40] F. Gandon, A. Poggi, G. Rimassa, and P. Turci. Multiagent corporate memory management system. *Applied Artificial Intelligence*, 16(9-10):699–720, 2002.

[41] P. Giorgini, M. Kolp, J. Mylopoulos, and M. Pistore. The Tropos methodology: An overview. In Bergenti et al. [10], pages 89–106.

[42] D. Hudson and M. Cohen. Use of intelligent agents in the diagnosis of cardiac disorders. *Computers in Cardiology*, pages 633–636, 2002.

[43] N. Jennings, P. Faratin, M. J. Johnson, T. J. Norman, P. OÕBrien, and M. E. Wiegand. Agent-based business process management. *International Journal of Cooperative Information Systems* , 5(2-3):105–130, 1996.

[44] N. Jennings and M. Wooldridge. Applications of intelligent agents. In *Agent technology: foundations, applications, and markets*, pages 3–28. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.

[45] N. R. Jennings. The ARCHON system and its applications. In *2nd International Working Conference on Cooperating Knowledge Based Systems (CKBS-94)* , pages 13–29, Keele, UK, 1994.

[46] G. Johnson. All science is computer science. *The New York Times*, Mar. 25 2001.

[47] M. Klusch. Information agent technology for the Internet: a survey. *Data Knowledge Engineering*, 36(3):337–372, 2001.

[48] G. Lanzola and H. Boley. *Experience with a functional-logic multi-agent architecture for medical problem solving*, pages 17–37. Idea Group Publishing, Hershey, PA, USA, 2002.

[49] M. Ljungberg and A. Lucas. The OASIS air-trafÞc management system. In *2nd Pacific Rim International Conference on Artificial Intelligence (PRICAI'92)*, Seoul, Korea, 1992.

[50] M. Luck, P. McBurney, O. Shehory, and S. Willmott. Agent technology: Computing as interaction a roadmap for agent-based computing, 2005. http://www.agentlink.org/roadmap.

[51] M. Mamei, F. Zambonelli, and L. Leonardi. Co-Þelds: Towards a unifying approach to the engineering of swarm intelligent systems. In P. Petta, R. Tolksdorf, and F. Zambonelli, editors, *Engineering Societies in the Agents World III* , volume 2577 of *LNCS* , pages 68–81. Springer-Verlag, Apr. 2003.

[52] A. Moreno and J. Nealon. *Applications of software agent technology*. Whitestein, 2003.

[53] V. Morreale, S. Bonura, F. Centineo, A. Rossi, M. Cossentino, and S. Gaglio. PRACTIONIST: Implementing PRACTIcal reasONIng sySTems. In F. Corradini, F. De Paoli, E. Merelli, and A. Omicini, editors, *WOA 2005: Dagli Oggetti agli Agenti. 6th AI*IA/TABOO Joint Workshop "From Objects to Agents": Simulation and Formal Analysis of Complex Systems* , pages 66–74, Camerino, MC, Italy, 14–16 Nov. 2005. Pitagora Editrice Bologna.

[54] E. Motta, J. Domingue, L. Cabral, and M. Gaspari. Irs-ii: A framework and infrastructure for semantic web services. In D. Fensel, K. P. Sycara, and J. Mylopoulos, editors, *2nd International Semantic Web Conference*, volume 2870 of *LNCS*, pages 306–318, Sanibel Island, FL, 2003. Springer.

[55] R. D. Nicola, G. Ferrari, and R. Pugliese. KLAIM: A kernel language for agent interaction and mobility. *IEEE Transaction on Software Engineering*, 24(5), 1998.

[56] M. J. North, N. T. Collier, and J. R. Vos. Experiences creating three implementations of the Repast agent modeling toolkit. *ACM Transactions on Modeling and Computer Simulation* , 16(1):1–25, Jan. 2006.

[57] A. Omicini. *SODA*: Societies and infrastructures in the analysis and design of agent-based systems. In P. Ciancarini and M. J. Wooldridge, editors, *Agent-Oriented Software Engineering*, volume 1957 of *LNCS*, pages 185–193. Springer-Verlag, 2001. 1st International Workshop (AOSE 2000), Limerick, Ireland, 10 June 2000. Revised Papers.

[58] A. Omicini and P. McBurney. The Agentlink III Technical Forums: Introduction to the Special Issue. *The Knowledge Engineering Review*, 20(2):95–97, June 2005. Special Issue: Revised Reports from Selected Technical Forum Groups at the 1st & 2nd AgentLink III Technical Forums.

[59] A. Omicini and S. Ossowski. Coordination and collaboration activities in cooperative information systems. *International Journal of Cooperative Information Systems* , 13(1):1–7, Mar. 2004. Special Issue: Coordination and Collaboration Technologies.

[60] A. Omicini, A. Ricci, and M. Viroli. *Agens Faber*: Toward a theory of artefacts for MAS. *Electronic Notes in Theoretical Computer Sciences*, 150(3):21–36, 29 May 2006. 1st International Workshop ÒCoordination and OrganizationÓ (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 22 Apr. 2005. Proceedings.

[61] A. Omicini and F. Zambonelli. Coordination for Internet application development. *Autonomous Agents and Multi-Agent Systems* , 2(3):251–269, Sept. 1999.

[62] H. V. D. Parunak. Manufacturing experience with the Contract Net. In M. N. Huhns, editor, *Distributed Artificial Intelligence*, pages 285–310, London, UK, 1987. Pitman.

[63] A. Perini and A. Susi. Developing a decision support system for integrated production in agriculture. *Environmental Modelling and Software* , 19(9):821–829, 2004.

[64] A. Poggi, M. Tomaiuolo, and G. Vitaglione. A security infrastructure for trust management in multiagent systems. In R. Falcone, K. S. Barber, J. Sabater-Mir, and M. P. Singh, editors, *Trusting Agents for Trusting Electronic Societies*, volume 3577 of *LNCS*, pages 162–179. Springer, 2005.

[65] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI architecture. In J. F. Allen, R. Fikes, and E. Sandewall, editors, *2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 473–484, San Mateo, CA, 1991. Morgan Kaufmann Publishers.

[66] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall / Pearson Education International, Englewood Cliffs, NJ, USA, 2nd edition, Dec. 2002. 2nd Edition.

[67] K. Schmidt and C. Simone. Mind the gap! towards a unified view of CSCW. In *The Fourth International Conference on the Design of Cooperative Systems COOP 2000*, May 2000.

[68] M. Schumacher. *Objective Coordination in Multi-Agent System Engineering – Design and Implementation*, volume 2039 of *LNAI*. Springer-Verlag, Apr. 2001.

[69] J. R. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, 1969.

[70] W. M. P. van der Aalst and K. van Hee. *Workflow Management: Models, Methods, and Systems*. MIT Press, 2002.

[71] M. Verdicchio and M. Colombetti. A commitment-based communicative act library. In *4rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 755–761, Utrecht, The Netherlands, 2005.

[72] D. Weyns, M. Schumacher, A. Ricci, M. Viroli, and T. Holvoet. Environments in multiagent systems. In *The Knowledge Engineering Review* [58], pages 127—141. Special Issue: Revised Reports from Selected Technical Forum Groups at the 1st & 2nd AgentLink III Technical Forums.

[73] M. Winikoff. JACK intelligent agents: An i ndustrial strength platform. In R. Bordini, M.Dastani, J.Dix, and A. E. Fallah-Seghrouchni, editors, *Multi-Agent Programming*, pages 175–193. Springer, Berlin, Germany, 2005.

[74] F. Zambonelli, N. R. Jennings, and M. Wooldridge. Developing multiagent systems: The Gaia methodology. *ACM Transactions on Software Engineering and Methodology*, 12(3):317–370, 2003.

[75] F. Zambonelli and A. Omicini. Challenges and research directions in agent-oriented software engineering. *Autonomous Agents and Multi-Agent Systems*, 9(3):253–283, Nov. 2004. Special Issue: Challenges for Agent-Based Computing.

[76] F. Zambonelli and H. V. D. Parunak. Towards a paradigm change in computer science and software engineering: A synthesis. *The Knowledge Engineering Review*, 18(4):329–342, Dec. 2003.