# Coordination and Agreement in Multi-Agent Systems

Sascha Ossowski

Centre for Intelligent Information Technologies (CETINIA),
Universidad Rey Juan Carlos,
Calle Tulipán s/n,
28933 Móstoles (Madrid), Spain
sascha.ossowski@urjc.es

**Abstract.** It is commonly accepted that coordination is a key characteristic of multi-agent systems and that, in turn, the capability of coordinating with others constitutes a centrepiece of agenthood. However, the key elements of coordination models, mechanisms, and languages for multi-agent systems are still subject to considerable debate. This paper provides a brief overview of different approaches to coordination in multi-agent systems. It will then show how these approaches relate to current efforts working towards a paradigm for smart, next-generation distributed systems, where coordination is based on the concept of agreement between computational agents.

## 1 Introduction

Most current transactions and interactions at business level, but also at leisure level, are mediated by computers and computer networks. From email, over social networks, to virtual worlds, the way people work and enjoy their free time has changed dramatically in less than a generation time. This change has made that IT research and development focuses on aspects like new Human-Computer Interfaces or enhanced routing and network management tools. However, the biggest impact has been on the way applications are thought and developed. These applications require components to which more and more complex tasks can be delegated, components that show higher levels of intelligence, components that are capable of sophisticated ways of interacting, as they are massively distributed, sometimes embedded in all sort of appliances and sensors. In order to allow for an efficient design and implementation of systems of these characteristics, it is necessary to effectively enable, structure, and regulate their communications in different contexts.

Such an enterprise raises a number of technological challenges. Firstly, the open distributed nature of such systems adds to the *heterogeneity* of its components. The system structure may evolve at runtime, as new nodes may appear or disappear at will. There is also a need for on-the-fly alignment of certain concepts that interactions relate to, as the basic ontological conventions in such systems will be very limited. The *dynamicity* of the environment calls for a continuous *adaptation* of the structures that regulate the components' interactions, so as to achieve and sustain desired functional properties. But also non-functional issues related to *scalability*, *security*, and *usability* need to be taken into account. When designing mechanisms that address these challenges, the notion of *autonomy* becomes central: components may show

M. Klusch, M. Pechoucek, and A. Polleres (Eds.): CIA 2008, LNAI 5180, pp. 16–23, 2008.

complex patterns of activity aligned with the different goals of their designers, while it is usually impossible to directly influence their behaviour from the outside.

Coordination in multi-agent system (MAS) aims at harmonising the interactions of multiple autonomous components or agents. Therefore, it appears promising to review different conceptual frameworks for MAS coordination, and to analyse the potential and limitations of the work done in that field with regard to some of the aforementioned challenges.

This paper is organised as follows. Section 2 provides brief overview of coordination in MAS. Section 3 proposes the notion of *agreement* as a centrepiece of an integrated approach to coordination in open distributed systems, and outlines some research topics related to the vision of a technology of agreement. Some conclusions are drawn in Section 4.

## 2 Coordination in Multi-Agent Systems

Maybe the most widely accepted conceptualisation of coordination in the MAS field originates from Organisational Science. It defines coordination the *management of dependencies* between organisational activities [21]. One of the many workflows in an organisation, for instance, may involve a secretary writing a letter, an official signing it, and another employee sending it to its final destination. The interrelation among these activities is modelled as a *producer/consumer* dependency, which can be managed by inserting additional *notification* and *transportation* actions into the workflow.

It is straightforward to generalise this approach to coordination problems in multi-agent systems. The subjects whose activities need to be coordinated are the agents, while the entities between which dependencies are usually goals, actions or plans. Depending on the characteristics of the MAS environment, a taxonomy of dependencies can be established, and a set of potential coordination actions assigned to each of them (e.g.[36], [26]). Within this model, the *process* of coordination is to accomplish two major tasks: first, a *detection* of dependencies needs to be performed, and second, a *decision* respecting which coordination action to apply must be taken. A coordination *mechanism* shapes the way that agents perform these tasks [24].

The *result* of coordination, and its *quality*, is conceived differently at different levels of granularity. Understanding coordination as *a way of adapting to the environment* [36] is quite well suited to address this question from a *micro-level* (agent-centric) perspective. This is particularly true for multi-agent settings. If new acquaintances enter an agent's environment, coordination amounts to re-assessing its former goals, plans and actions, so as to account for the new (potential) dependencies between itself and other agents. If a STRIPS-like planning agent, for instance, is put into a multi-agent environment, it will definitely have to accommodate its individual plans to the new dependencies between its own prospective actions and potential actions of others, trying to exploit possible synergies (others may free certain relevant blocks for it), and avoiding harmful dependencies (making sure that others do not unstack intentionally constructed stacks etc). At this level, the result of coordination, the agent's adapted individual plan, is the better the closer it takes the agent to the achievement of its goals in the multi-agent environment.

From a *macro-level* (MAS-centric) perspective, the outcome of coordination can be conceived a “global” plan (or decision, action etc.). This may be a “joint plan” [29] if the agents reach an explicit agreement on it during the coordination process, or just the sum of the agents' individual plans (or decisions, actions etc. – sometimes called “multi-plan” [27]) as perceived by an external observer. Roughly speaking, the quality of the outcome of coordination at the macro-level can be evaluated with respect to the agents’ joint goals or the desired functionality of the MAS as a whole. If no such notion can be ascribed to the MAS, other, more basic features can be used instead. A good result of coordination, for instance, often relates to efficiency, which frequently comes down to the notion of Pareto-optimality. The amount of resources necessary for coordination (e.g. the number of messages necessary) is also sometimes used as a measure of efficiency.

The dependency model of coordination appears to be particularly well suited to *represent* relevant features of a coordination problem in MAS. The TAEMS framework [11], for instance, has been used to model coordination requirements in a variety of interesting MAS domains. It is also useful to rationalise observed coordination behaviour in line with a knowledge-level perspective [22]. Still, dependency detection may come to be a rather knowledge intensive task, which is further complicated by incomplete and potentially inconsistent local views of the agents. Moreover, making timely decisions that lead to efficient coordination actions is also everything but trivial. The problem becomes even more difficult when agents pursuing partially conflicting goals come into play [26]. In all but the simplest MAS, the instrumentation of these tasks gives rise to complex patterns of interactions among agents.

From a design perspective, coordination is probably best conceived as the effort of *governing the space of interaction* [6] of a MAS, as the basic challenge amounts to how to make agents converge on interaction patterns that adequately (i.e. instrumentally with respect to desired MAS features) solve the dependency detection and decision tasks. A variety of approaches that tackle this problem can be found in the literature, shaping the interaction space either directly, by making assumptions on agent behaviours and/or knowledge, or indirectly, by modifying the *context* of the agents in the MAS environment. The applicability of these mechanisms depends largely on the number and type of assumptions that one may make regarding the possibility of manipulating agent programs, agent populations, or the agents’ environment. This, in turn, is dependent on the characteristics of the coordination problem at hand.

The RICA-J framework [31], for instance, provides an ontology of interaction types, together with their associated protocols. Agents can freely choose to play or abandon certain roles within an interaction but, when using the framework, an agent programmer is limited to using protocol compliant actions.

Governing coordination infrastructures make a clear separation between the *enabling* services that they provide (e.g. communication channel or blackboard-based communication primitives) and the *governing* aspects of interaction, which are usually described within a declarative language (e.g. programmable tuple spaces) [25]. The access regulations for the elements of the MAS environment (resources, services, etc) expressed in such a language are sometimes called *environment laws* [30].

Electronic Institutions (EI) [23] use organisational abstractions to shape the interactions of the agents participating in them. Agents play different roles in the (sub-) protocols that, together with additional rules of behaviour, determine the legal sequences of illocutions that may arise within a particular instance of a scene. Scenes, in turn, are interconnected and synchronised by means of transitions within a performative structure. Norms, as additional institutional abstractions, express further behaviour restrictions for agents. In the EI framework, agents can only interact with each other through specific institutional agents, called governors [13], which assure that all behaviour complies with the norms and that it obeys the performative structure. So, different from the aforementioned approaches, the governing or regulating responsibility is transferred from the infrastructure to specialized middle agents.

From the point of view of an individual agent, the problem of coordination essentially boils down to finding the sequence of actions that, given the regulations within the system (or, if possible in a certain environment, the expected cost of transgressing them), best achieves its goals. In practice, this implies a series of non-trivial problems. Models of coalition formation determine when and with whom to form a team for the achievement of some common (sub-) goal, and how to distribute the benefits of synergies that arise from this cooperation [32]. Distributed planning approaches [12] may determine how to (re-)distribute tasks among team members and how to integrate results. From an individual agent's perspective, the level of trustworthiness of others is central to almost every stage of these processes, so as to determine whether other agents are likely to honour the commitments that have been generated [33].

An appealing way to tackle both the system-level and the agent-level requirements is to take an organisation-oriented tack towards the problem of MAS coordination. Organisational models underlying approaches such as Agent-Group-Role [14], MOISE [18], EI [23], or RICA [31] provide a rich set of concepts to specify and structure mechanisms that govern agent interactions through the corresponding infrastructures or middleware. But they can also facilitate the agents' local decision-making tasks. For instance, role and interaction taxonomies can be used to find suitable interactions partners, by providing additional information regarding the usability of services in a certain interaction context [15]. Structural information about roles can also be used for the bootstrapping of reputation mechanism, when only very limited information about past interactions is available in the system [5]. Role hierarchies, and other types of structural information, can also be extended on-the-fly to improve system performance [17]. In general, the fact that organisational structures may dynamically evolve, shifts the attention from their traditional use as a *design-time* coordination mechanism for mainly closed distributed problem-solving systems, to an adaptive *run-time* coordination mechanism also applicable to open MAS [24].

## 3 Towards a Technology of Agreement

The previous section has given a brief overview of work on coordination mechanisms that has been carried in the MAS field. Even though an attempt has been made to structure and present it in some coherent manner, the reader will have noticed that several quite different approaches and mechanisms coexist under the "umbrella" of the term coordination. Not all of them are relevant to the challenges for the design of open distributed systems outlined in the introduction. For instance, the whole set of *coupled* coordination mechanisms [35] are effectively useless for the purpose of this paper, as they require having a direct influence on the agent programs. On the other hand, the problem of semantic interoperability is usually outside the scope of MAS coordination models and languages.

The notion of *agreement* among computational agents appears to be better suited as the fundamental notion for the proposal outlined in this paper. Until recently, the concept of agreement was a domain of study mainly for philosophers, sociologists and was only applicable to human societies. In recent years, the growth of disciplines such as social psychology, socio-biology, social neuroscience, together with the spectacular emergence of the information society technologies, have changed this situation. Presently, agreement and all the processes and mechanisms implicated in reaching agreements between different kinds of agents are a subject of research and analysis also from technology-oriented perspectives.

The process of agreement-based coordination can be designed based on two main elements:

(1) a normative context, that determines the rules of the game, i.e. interaction patterns and additional restrictions on agent behaviour; and
(2) a call-by-agreement interaction method, where an agreement for action between the agents that respects the normative context is established first; then actual enactment of the action is requested.

The techniques based on organizational structures discussed in the previous section will be useful to specify and design such systems. In addition, semantic alignment, norms, argumentation and negotiation, as well as trust and reputation mechanisms will be in the "agreement technology sandbox".

*Semantic* technologies constitute a centrepiece of the approach as semantic problems pervade all the others. Solutions to semantic mismatches and alignment of ontologies [4] are needed to have a common understanding of norms or of deals, just to put two examples. The use of semantics-based approaches to service discovery and composition will allow exploring the space of possible interactions and, consequently, shaping the set of possible agreements [15].

At system-level, *norms* are needed to determine constraints that the agreements, and the processes to reach them, have to satisfy. Reasoning about a system's norms is necessary at design-time to assure that the system has adequate properties, but it may also be necessary at run-time, as complex systems usually need dynamic regulations [16]. *Organisational* structures further restrict the way agreements are reached by fixing the social structure of the agents: the capabilities of their roles and the relationships among them (e.g. power, authority) [3].

Moving further towards the agent-level, *negotiation* methods are essential to make agents reach agreements that respect the constraints imposed by norms and organisations. These methods need to be complemented by an argumentation-based approach: by exchanging arguments, the agents' mental states may evolve and, consequently, the status of offers may change [2] [7]. Finally, agents will need to use *trust* mechanisms that summarise the history of agreements and subsequent agreement executions in order to build long-term relationships between the agents. Trust is the technology that complements traditional security mechanisms by relying on social mechanisms that interpret the behaviour of agents [34].

One may conceive the aforementioned topics in a "tower structure", with semantic technologies at the bottom layer and trust mechanisms at the top, where each level provides functionality to the levels above [1]. Notice, however, that there is also a certain feedback from higher to lower layers as, for instance, reputation mechanisms may influence organisational structures such as role and interaction hierarchies [17]; and this information can as well be used for semantic alignment [4] and discovery [15].

## 4 Discussion

This paper has presented an overview of different approaches to coordination in the MAS field. It has been argued that the notion of agreement is essential to instil coordination in open distributed systems. Some existing technologies from the field of MAS coordination can be applied to this respect, but others – and in particular semantic technologies – need to be added. Several research efforts are currently ongoing that may contribute to the development of a "technology of agreement" in one or another way. The attempt to harmonise these efforts, which is currently being carried out at European level, promotes the emergence of a new paradigm for next generation distributed systems based on the notion of *agreement* between computational agents [9].

## Acknowledgements

Many ideas reported in this paper draw upon joint work with Carles Sierra, Vicent Botti, and others, in the framework of a Spanish national project on "Agreement Technology". This term was first mentioned by Mike Wooldridge in internal discussions at the AAMAS conference in 2004. It has also been used as a title for a conference by Nick Jennings. I am also thankful to Axel Polleres, Cristiano Castelfranchi and Leila Amgoud for their comments regarding the different technological challenges related to a computational notion of agreement, as well as Andrea Omicini for our discussions on the different stances on coordination in multi-agent systems. This work was partially supported by the Spanish Ministry of Science and Innovation, grants TIN2006-14630-C03-02 and CSD2007-00022 (CONSOLIDER-INGENIO 2010).

## References

[1] Agreement Technologies project homepage, `http://www.agreement-technologies.org/`

[2] Amgoud, L., Dimopolous, Y., Moraitis, P.: A unified and general framework for argumentation-based negotiation. In: Proc. 6th Int. Joint Conference on Autonomous Agents and Multi-Agents Systems (AAMAS 2007), pp. 963–970. IFAAMAS (2007)

[3] Argente, E., Julian, V., Botti, V.: Multi-Agent System Development based on Organizations. Electronic Notes in Theoretical Computer Science 150(3), 55–71 (2006)

[4] Atienza, M., Schorlemmer, M.: I-SSA - Interaction-situated Semantic Alignment. In: Proc Int. Conf. on Cooperative Information Systems (CoopIS 2008) (to appear, 2008)

[5] Billhardt, H., Hermoso, R., Ossowski, S., Centeno, R.: Trust-based Service Provider Selection in Open Environments. In: Proc. ACM Symposium on Applied Computing (SAC-2007), pp. 1375–1380. ACM Press, New York (2007)

[6] Busi, N., Ciancarini, P., Gorrieri, R., Zavattaro, G.: Coordination Models - A Guided Tour. In: Omicini, et al. (eds.) Coordination of Internet Agents: Models, Technologies, and Applications, pp. 6–24. Springer, Heidelberg (2001)

[7] Caminada, M., Amgoud, L.: On the evaluation of argumentation formalisms. Artificial Intelligence Journal 171(5-6), 286–310 (2007)

[8] Castelfranchi, C., Dignum, F., Jonker, C., Treur, J.: Deliberative Normative Agents - Principles and Architecture. In: Jennings, Lespérance (eds.) Intelligent Agents VI. LNCS, vol. 1757, pp. 364–378. Springer, Heidelberg (2000)

[9] COST Act. IC0801, http://www.cost.esf.org/index.php?id=110&action_number=IC0801

[10] Debenham, J., Sierra, C.: Merging intelligent agency and the Semantic Web. Knowledge-Based Systems 21(3), 184–191 (2008)

[11] Decker, K.: TAEMS: A Framework for Environment Centered Analysis and Design of Coordination Mechanisms. In: O'Hare, Jennings (eds.) Foundations of Distributed Artificial Intelligence. John Wiley and Sons, Chichester (1996)

[12] Durfee, E.: Distributed Problem Solving and Planning. In: Luck, M., Mařík, V., Štěpánková, O., Trappl, R. (eds.) ACAI 2001 and EASSS 2001. LNCS (LNAI), vol. 2086, pp. 118–149. Springer, Heidelberg (2001)

[13] Esteva, M., Rosell, B., Rodríguez-Aguilar, J.A., Arcos, J.L.: AMELI - An agent-based middleware for electronic institutions. In: Proc. of the Third Int. Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004), pp. 236–243. ACM Press, New York (2004)

[14] Ferber, J., Gutknecht, O., Fabien, M.: From Agents to Organizations - An Organizational View of Multi-agent Systems. In: Giorgini, P., Müller, J.P., Odell, J.J. (eds.) AOSE 2003. LNCS, vol. 2935, pp. 214–230. Springer, Heidelberg (2004)

[15] Fernández, A., Ossowski, S.: Exploiting Organisational Information for Service Coordination in Multiagent Systems. In: Proc. of the Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2008), pp. 257–264. IFAAMAS (2008)

[16] Gaertner, D., García-Camino, A., Noriega, P.,Rodríguez-Aguilar, J.A., Vasconcelos, W.: Distributed norm management in regulated multiagent systems. In: Proc. Int. Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2007), pp. 624–631. IFAAMAS (2007)

[17] Hermoso, R., Centeno, R., Billhardt, H., Ossowski, S.: Extending Virtual Organizations to improve trust mechanisms (Short Paper). In: Proc. of the Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2008), pp. 1489–1492. IFAAMAS (2008)

[18] Hubner, J., Sichman, J., Boissier, O.: Developing organised multiagent systems using the MOISE+ model: programming issues at the system and agent levels. Int. Journal of Agent-Oriented Software Engineering 1(3/4), 370–395 (2006)

[19] Klusch, M., Sycara, K.: Brokering and matchmaking for coordination of agent societies: a survey. In: Coordination of Internet Agents: Models, Technologies, and Applications (Omicini y otros), pp. 197–224. Springer, Heidelberg (2001)

[20] Klusch, M., Fries, B., Sycara, K.: Automated Semantic Web Service Discovery with OWLS-MX. In: Proceedings of 5th International Conference on Autonomous Agents and Multi- Agent Systems (AAMAS-2006), pp. 915–922. ACM Press, New York (2006)

[21] Malone, T., Crowston, K.: The Interdisciplinary Study of Co-ordination. Computing Surveys 26(1), 87–119 (1994)

[22] Newell, A.: Reflections on the Knowledge Level. Artificial Intelligence 59, 31–38 (1993)

[23] Noriega, P., Sierra, C.: Electronic Institutions – Future Trends and Challenges. In: Klusch, M., Ossowski, S., Shehory, O. (eds.) CIA 2002. LNCS (LNAI), vol. 2446, pp. 14–17. Springer, Heidelberg (2002)

[24] Omicini, A., Ossowski, S.: Objective versus Subjective Coordination in the Engineering of Agent Systems. In: Klusch, M., Bergamaschi, S., Edwards, P., Petta, P. (eds.) Intelligent Information Agents. LNCS (LNAI), vol. 2586, pp. 179–202. Springer, Heidelberg (2003)
[25] Omicini, A., Ossowski, S., Ricci, A.: Coordination Infrastructures in the Engineering of Multiagent Systems. In: Bergenti, Gleizes, Zambonelli (eds.) Methodologies and software engineering for agent systems – The Agent-Oriented Software Engineering Handbook, pp. 273–296. Kluwer, Dordrecht (2004)
[26] Ossowski, S.: Co-ordination in Artificial Agent Societies. LNCS (LNAI), vol. 1535. Springer, Heidelberg (1998)
[27] Ossowski, S.: Constraint Based Coordination of Autonomous Agents. Electronic Notes in Theoretical Computer Science 48, 211–226 (2001)
[28] Ossowski, S., Menezes, R.: On Coordination and its Significance to Distributed and Multi-Agent Systems. Journal of Concurrency and Computation - Practice and Experience 18(4), 359–370 (2006)
[29] Rosenschein, J., Zlotkin, G.: Designing Conventions for Automated Negotiation. AI Magazine 15(3), 29–46 (1995)
[30] Schumacher, M., Ossowski, S.: The governing environment. In: Weyns, D., Van Dyke Parunak, H., Michel, F. (eds.) E4MAS 2005. LNCS (LNAI), vol. 3830, pp. 88–104. Springer, Heidelberg (2006)
[31] Serrano, J.M., Ossowski, S.: On the Impact of Agent Communication Languages on the Implementation of Agent Systems. In: Klusch, et al. (eds.) Cooperative Information Agents VIII. LNCS, vol. 2782, pp. 92–106. Springer, Heidelberg (2004)
[32] Shehory, O., Sycara, K., Somesh, J.: Multi-agent Coordination through Coalition Formation. In: Rao, A., Singh, M.P., Wooldridge, M.J. (eds.) ATAL 1997. LNCS, vol. 1365, pp. 143–154. Springer, Heidelberg (1998)
[33] Sabater, J., Sierra, C.: Review on Computational Trust and Reputation Models. Artificial. Intelligence Review 24(1), 33–60 (2005)
[34] Sierra, C., Debenham, J.: Information-Based Agency. In: Proc Int. Joint Conference on AI (IJCAI-2007), pp. 1513–1518. AAAI Press, Menlo Park (2007)
[35] Tolksdorf, R.: Models of Coordination. In: Omicini, A., Tolksdorf, R., Zambonelli, F. (eds.) ESAW 2000. LNCS (LNAI), vol. 1972. Springer, Heidelberg (2000)
[36] von Martial, F.: Coordinating Plans of Autonomous Agents. LNCS (LNAI), vol. 610. Springer, Heidelberg (1992)