

1000 Years of Coo-BDI*

Viviana Mascardi and Davide Ancona

DISI - Università di Genova
Via Dodecaneso 35, 16146, Genova, Italy
{viviana.mascardi,davide.ancona}@unige.it

Abstract. The idea of extending the BDI architecture with cooperativity started shaping in 2003 when two independent proposals to support cooperation in a BDI setting were presented at DALT. One proposal, Coo-BDI, extended the BDI architecture by allowing agents to cooperate by exchanging and sharing plans in a quite flexible way; the other extended the BDI operational semantics for introducing speech-act based communication, including primitives for plan exchange. Besides allowing a natural and seamless integration with speech-act based communication for BDI languages, the intuitions behind Coo-BDI have proved to be promising and attractive enough to give rise to new investigations. In this retrospective review we discuss papers that were influenced by Coo-BDI and we outline other potential developments for future research.

1 Life after Coo-BDI

The paper introducing Coo-BDI [3] ended with the following statement:

We are currently working with the authors of [27] to realize this extension..

The planned extension has consisted in the design and implementation of a unified architecture for highly cooperative BDI agents meeting the following requirements:

- messages adhere to the form proposed in [27], including a `<tellHow, SenderId, Plan>` performative allowing the receiver to add `Plan` to its plan library if `SenderId` is trusted, and
- plans are associated with access specifiers as in Coo-BDI so that agents can decide when a plan should be shared with others by means of a `tellHow` message.

Together with J. F. Hübner and R. H. Bordini we worked one year to finish the design and implementation of our planned extension, and finally the Coo-BDI approach was successfully and smoothly integrated with AgentSpeak [9, 29] in the context of Jason [10]. Jason implemented the operational semantics

* There are only 10 types of people in the world: those who understand binary, and those who don't.

given in [9] as well as the extensions in [27], giving the necessary formal and practical basis for plan exchange among BDI agent in the way required by Coo-BDI. The obtained language was named Coo-AgentSpeak and was presented at AAMAS 2004 [4]. The extensions purposely made to Jason for supporting Coo-AgentSpeak features are part of the standard Jason release.

Encouraged by the promising results, we explored the applicability of the Coo-BDI principles to other concrete scenarios. In particular, we were interested in investigating if and how Web Services (WSs) technologies could support a component, described in terms of beliefs, desires and intentions, that dynamically adapts its behavior to new environments (namely, a Coo-BDI agent). A positive answer came from CooWS [11] which implements the ideas behind the Coo-BDI by means of WS technologies. In CooWS plan bodies are expressed in BPEL [1], a high-level scripting language for Web Services built on top of WSDL [12]. Agents able to execute a BPEL specification can execute the body of any plan, making the exchange of plans among agents a fruitful extension of the basic BDI architecture.

In parallel with this practical research activity, theoretical work was carried out for finding a BDI logic suitable for modeling the behavior of Coo-BDI agents. BDI^{ATL} [26] was the result of that effort. By replacing ATL^* (Alternating-Time Temporal Logic [2]) with CTL^* (an extension of Computation Tree Logic and Linear Temporal Logic [20]) in Rao and Georgeff's BDI logic [32, 30, 31], BDI^{ATL} allows us to express new commitment strategies that could not be defined there. In particular, we can express three variants of Rao and Georgeff's "open minded" commitment: "independent open minded", "optimistic open minded", and "pessimistic open minded". In these commitment strategies the new features that ATL^* adds to CTL^* , namely *cooperation modalities*, are exploited for expressing the way of thinking of rational Coo-BDI agents.

After the intense activity of the beginning, research on Coo-BDI slackened for a few years during which we pursued other scientific goals, including that of deepening our knowledge on semantic web issues. When the competencies acquired on these themes in general, and on ontology matching [21] in particular, were mature enough, an inspired intuition of A. Ricci gave us the chance to resume Coo-BDI and apply to it the techniques we were experimenting in other domains. The result was CooL-AgentSpeak [24], the "Cooperative Description-Logic AgentSpeak" language integrating Coo-BDI and AgentSpeak-DL [28] and enhancing them with *ontology matching capabilities*. In CooL-AgentSpeak, search for a plan takes place as in Coo-BDI. However, handling an event is more flexible as it is not based solely on unification and on the subsumption relation between concepts as in AgentSpeak-DL, but also on ontology matching. Belief querying and updating take advantage of ontological matching as well. The syntax of the language and motivating scenarios for its adoption are given. A sketch of the operational semantics and of how CooL-AgentSpeak can be implemented on top of JASDL [23] are also provided.

2 The Lives of the Others

Many research activities carried out under the agents and MASs umbrella share with Coo-BDI the idea of exchanging knowledge among peers that, otherwise, could not properly cope with some situations.

M. Baldoni, C. Baroglio, A. Martelli, V. Patti and C. Schifanella [5–8] face the issue of allowing an entity to play a role in an interaction ruled by a choreography, even when it owns no policy conforming to that role. The scenario of interest is Service-oriented Computing. As the authors recognize, in an agent framework the solution might easily come from a Coo-BDI-like approach: one might think of dynamically enriching the set of behaviors of the agent, which failed the conformance test, by asking other agents to supply a correct interaction policy. In Service-oriented Computing, however, a Coo-BDI-like approach can not be applied since in that scenario it is fundamental that knowledge is available before the interaction among the peers takes place.

The work by S. Costantini, P. Dell’Acqua and L. M. Pereira [15] discusses issues related to learning rules from other agents. The origins of that work date back to 2005, with the prototype implementation presented in [19], developed in DALI [18]. In 2008 that implementation has been enriched with temporal-logic-like operators [14, 17], and experiments in Ambient Intelligent applications have been carried out [16]. In the more recent paper [15], the authors further enrich the approach with a meta-evaluation component that prevents agents to blindly accept and incorporate new knowledge by allowing them to evaluate (and thus possibly discard) it according to its usefulness. The proposal adds to Coo-BDI the very relevant aspect of meta-reasoning for evaluating, activating and deactivating the new knowledge, where evaluation may in principle affect the level of trust of source agents.

The work by Meneguzzi and Luck [25] describes how a procedural agent model can be modified to allow an agent to build new plans at runtime by chaining existing fine-grained plans from a plan library into high-level plans. The applicability of the approach is demonstrated through a modification to the AgentSpeak architecture, where declarative and procedural aspects are combined together. Meneguzzi and Luck propose an integration with the Coo-BDI approach as a possible future extension to their architecture to partially overcome efficiency issues, since getting plans from other planning-capable agents may significantly reduce the amount of time spent to create plans from scratch.

The Coo-BDI approach to plan failure has been easily incorporated into the AgentSpeak meta-interpreter designed and implemented by M. Winikoff [34] and into the guidelines on how to create multi-agent systems using Erlang provided by C. Varela, C. Abalde, L. M. Castro and J. Gulías [33].

Finally, the framework Agent Coordination and Cooperation Cognitive Model, AC³M [13], exhibits connections with Coo-BDI as well: the relationships between coordination, cooperation, BDI and OODA (Observe-Orient-Decide-Act cycle) are analyzed, with a particular focus on uncertain environments.

3 The Future

Research on Coo-BDI has not been financed within a specific project, but has been mainly driven by the willingness of several researchers to collaborate together, by exploiting cross-fertilization fostered by their rather different research backgrounds. Nevertheless, in these eight years the main ideas coming out from this collaboration have proved to have a certain influence on the research community of multi-agent systems, and we believe that there are still many interesting opportunities for improving and extending them in the near future.

Cooperative multi-agent systems find their natural applications in mobile code, context-aware and self-adaptive systems, but also semantic web applications. The most recent and interesting extensions to AgentSpeak discussed in [24] open new interesting scenarios in the intersection of multi-agent systems and advanced semantic web applications, including the Linked Open Data and the Federated Social Web. However, as we highlighted in the IAT 2011 paper, “cross-ontological” knowledge and reasoning may lead to unwanted behavior. Precision and recall of the best performing ontology matching algorithms seldom reach 100% on real ontologies (see <http://oaei.ontologymatching.org/2010/results/benchmarks/index.html>), and this means that using real ontologies and real ontology matchers, wrong matches might be returned, with possibly destructive consequences.

In order to cope with the intrinsic limitations of ontology matching techniques available today, we would greatly benefit from meta-reasoning capabilities similar to those discussed in [15]. Such capabilities might in fact allow agents to reason on the consequences of adopting new plans involving cross-ontological knowledge for ensuring a better control on which plans might be safely incorporated into the plan base, thus limiting risks.

More in general, safety and security are properties of paramount importance for cooperative multi-agent systems, especially when exploited in the context of mobile code, and much work still have to be done to make Coo-BDI usable in practice in contexts where safety and security are serious concerns. Interesting research directions include static and dynamic typechecking and verification of Coo-BDI agents, exploiting for instance session types [22].

The investigation of safety issues and the implementation of CooL-AgentSpeak in Jason is another short term research goal. The far future is too far to be predicted (especially when projects are not funded!), but we are confident that we will be able to talk about Coo-BDI in the next 10000 years¹!

References

1. Alves, A., et al.: Web Services Business Process Execution Language version 2.0, public review draft, August 23 (2006), <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-specification-draft.html> (accessed on August 30, 2011)

¹ Check the footnote in the first page...

2. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *Journal of the ACM* 49(5), 672–713 (2002)
3. Ancona, D., Mascardi, V.: Coo-BDI: Extending the BDI Model with Cooperativity. In: Leite, J., Omicini, A., Sterling, L., Torroni, P. (eds.) *DALT 2003*. LNCS (LNAI), vol. 2990, pp. 109–134. Springer, Heidelberg (2004)
4. Ancona, D., Mascardi, V., Hübner, J.F., Bordini, R.H.: Coo-AgentSpeak: Cooperation in AgentSpeak through plan exchange. In: *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2004*, pp. 696–705. IEEE Computer Society (2004)
5. Baldoni, M., Baroglio, C., Martelli, A., Patti, V., Schifanella, C.: Interaction Protocols and Capabilities: A Preliminary Report. In: Alferes, J.J., Bailey, J., May, W., Schwertel, U. (eds.) *PPSWR 2006*. LNCS, vol. 4187, pp. 63–77. Springer, Heidelberg (2006)
6. Baldoni, M., Baroglio, C., Martelli, A., Patti, V., Schifanella, C.: The need of capability requirements inside choreographies and interaction protocols. In: Yan, Y., Zhang, L. (eds.) *Proceedings of the International Workshop on Service Oriented Techniques, SOT 2006*, pp. 17–24 (2006)
7. Baldoni, M., Baroglio, C., Martelli, A., Patti, V., Schifanella, C.: Preserving players goals: a choreography-driven matchmaking approach. In: Baldoni, M., Boccalatte, A., De Paoli, F., Martelli, M., Mascardi, V. (eds.) *Proceedings of the 8th AI*IA/TABOO Joint Workshop “From Objects to Agents”, WOA 2007*, pp. 132–139. Seneca Edizioni Torino (2007)
8. Baldoni, M., Baroglio, C., Martelli, A., Patti, V., Schifanella, C.: Reasoning on choreographies and capability requirements. *International Journal of Business Process Integration and Management* 2(4), 247–261 (2007)
9. Bordini, R.H., Moreira, Á.F.: Proving BDI properties of agent-oriented programming languages. *Annals of Mathematics and Artificial Intelligence* 42, 197–226 (2004)
10. Bordini, R.H., Wooldridge, M., Hübner, J.F.: *Programming Multi-Agent Systems in AgentSpeak using Jason*. John Wiley & Sons (2007)
11. Bozzo, L., Mascardi, V., Ancona, D., Busetta, P.: COOWS: Adaptive BDI agents meet service-oriented computing. In: Gleizes, M.P., Kaminka, G.A., Nowé, A., Ossowski, S., Tuyls, K., Verbeeck, K. (eds.) *Proceedings of the 3rd European Workshop on Multi-Agent Systems, EUMAS 2005*, p. 473. Koninklijke Vlaamse Academie van Belie voor Wetenschappen en Kunsten (2005); longer version of this paper also appeared in the *Proceedings of the WWW/Internet 2005 Conference*, edited by P. Isaias and M. B. Nunes, vol. II, pp. 205–209 (2005)
12. Christensen, E., Curbera, F., Meredith, G., Weerawarana, S.: *Web Services Description Language (WSDL) 1.1*, W3C note, March 15 (2001), <http://www.w3.org/TR/wsd1> (accessed on August 30, 2011)
13. Consoli, A., Tweedale, J., Jain, L.C.: An Architecture for Agent Coordination and Cooperation. In: Apolloni, B., Howlett, R.J., Jain, L.C. (eds.) *KES 2007, Part III*. LNCS (LNAI), vol. 4694, pp. 934–940. Springer, Heidelberg (2007)
14. Costantini, S., Dell’Acqua, P., Pereira, L.M.: A multi-layer framework for evolving and learning agents. In: *Proceedings of the AAAI 2008 Workshop on Metareasoning: Thinking about Thinking*. Stanford University, AAAI Press (2008)
15. Costantini, S., Dell’Acqua, P., Pereira, L.M.: Conditional Learning of Rules and Plans by Knowledge Exchange in Logical Agents. In: Pasche, A. (ed.) *RuleML 2011 - Europe*. LNCS, vol. 6826, pp. 250–265. Springer, Heidelberg (2011)

16. Costantini, S., Dell'Acqua, P., Pereira, L.M., Toni, F.: Learning and evolving agents in user monitoring and training. In: Proceedings of the 48th National Conference of the Italian Association for Computer Science and Automatic Computation, AICA Conference 2010 (2010)
17. Costantini, S., Dell'Acqua, P., Pereira, L.M., Tsintza, P.: Runtime verification of agent properties. In: Proceedings of the International Conference on Applications of Declarative Programming and Knowledge Management, INAP 2009 (2009)
18. Costantini, S., Tocchio, A.: The DALI Logic Programming Agent-Oriented Language. In: Alferes, J.J., Leite, J. (eds.) JELIA 2004. LNCS (LNAI), vol. 3229, pp. 685–688. Springer, Heidelberg (2004)
19. Costantini, S., Tocchio, A.: Learning by knowledge exchange in logical agents. In: Corradini, F., De Paoli, F., Merelli, E., Omicini, A. (eds.) Proceedings of the 6th AI*IA/TABOO Joint Workshop “From Objects to Agents”, WOA 2005, pp. 1–8. Pitagora Editrice Bologna (2005)
20. Emerson, E.A., Halpern, J.Y.: “Sometimes” and “not never” revisited: on branching versus linear time temporal logic. *Journal of the ACM* 33(1), 151–178 (1986)
21. Euzenat, J., Shvaiko, P.: *Ontology Matching*. Springer, Heidelberg (2007)
22. Honda, K.: Types for Dynamic Interaction. In: Best, E. (ed.) *CONCUR 1993*. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993)
23. Klapiscak, T., Bordini, R.H.: JASDL: A Practical Programming Approach Combining Agent and Semantic Web Technologies. In: Baldoni, M., Son, T.C., van Riemsdijk, M.B., Winikoff, M. (eds.) *DALT 2008*. LNCS (LNAI), vol. 5397, pp. 91–110. Springer, Heidelberg (2009)
24. Mascardi, V., Ancona, D., Bordini, R.H., Ricci, A.: Cool-AgentSpeak: Enhancing AgentSpeak-DL agents with plan exchange and ontology services. In: Boissier, O., Bradshaw, J., Cao, L., Fischer, K., Hacid, M.-S. (eds.) *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT 2011*, pp. 109–116. IEEE Computer Society (2011)
25. Meneguzzi, F., Luck, M.: Composing High-Level Plans for Declarative Agent Programming. In: Baldoni, M., Son, T.C., van Riemsdijk, M.B., Winikoff, M. (eds.) *DALT 2007*. LNCS (LNAI), vol. 4897, pp. 69–85. Springer, Heidelberg (2008)
26. Montagna, R., Delzanno, G., Martelli, M., Mascardi, V.: BDI^{ATL} : An alternating-time BDI logic for multiagent systems. In: Gleizes, M.P., Kaminka, G.A., Nowé, A., Ossowski, S., Tuyls, K., Verbeeck, K. (eds.) *Proceedings of the 3rd European Workshop on Multi-Agent Systems, EUMAS 2005*, pp. 214–223. Koninklijke Vlaamse Academie van Belie voor Wetenschappen en Kunsten (2005)
27. Moreira, Á.F., Vieira, R., Bordini, R.H.: Extending the Operational Semantics of a BDI Agent-Oriented Programming Language for Introducing Speech-Act Based Communication. In: Leite, J., Omicini, A., Sterling, L., Torroni, P. (eds.) *DALT 2003*. LNCS (LNAI), vol. 2990, pp. 135–154. Springer, Heidelberg (2004)
28. Moreira, Á.F., Vieira, R., Bordini, R.H., Hübner, J.F.: Agent-Oriented Programming with Underlying Ontological Reasoning. In: Baldoni, M., Endriss, U., Omicini, A., Torroni, P. (eds.) *DALT 2005*. LNCS (LNAI), vol. 3904, pp. 155–170. Springer, Heidelberg (2006)
29. Rao, A.S.: AgentSpeak(L): BDI Agents Speak Out in a Logical Computable Language. In: Perram, J.W., Van de Velde, W. (eds.) *MAAMAW 1996*. LNCS, vol. 1038, pp. 42–55. Springer, Heidelberg (1996)
30. Rao, A.S., Georgeff, M.P.: Asymmetry thesis and side-effect problems in linear-time and branching-time intention logics. In: Mylopoulos, J., Reiter, R. (eds.) *Proceedings of the 12th International Joint Conference on Artificial Intelligence, IJCAI 1991*, pp. 498–505. Morgan Kaufmann (1991)

31. Rao, A.S., Georgeff, M.P.: Deliberation and intentions. In: D'Ambrosio, B., Smets, P. (eds.) Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence, UAI 1991. Morgan Kaufmann (1991)
32. Rao, A.S., Georgeff, M.P.: Modelling rational agents within a BDI-architecture. In: Proceedings of the 2nd International Conference of Principles of Knowledge Representation and Reasoning, KR 1991. Morgan Kaufmann Publishers (1991)
33. Varela, C., Abalde, C., Castro, L.M., Gulas, J.: On modelling agent systems with Erlang. In: Cesarini, F., Wadler, P. (eds.) 3rd ACM SIGPLAN Workshop on Erlang, Proceedings, pp. 65–70. ACM (2004)
34. Winikoff, M.: An AgentSpeak Meta-interpretter and Its Applications. In: Bordini, R.H., Dastani, M.M., Dix, J., El Fallah Seghrouchni, A. (eds.) PROMAS 2005. LNCS (LNAI), vol. 3862, pp. 123–138. Springer, Heidelberg (2006)