

# Back to the Future: an Interaction-oriented Framework for Social Computing

Matteo Baldoni, Cristina Baroglio, Elisa Marengo, Viviana Patti  
Università degli Studi di Torino

Alessandro Ricci  
Università degli Studi di Bologna

## Abstract

*We propose an interaction-oriented framework and the related support infrastructure, that reifies commitment-based interaction protocols into programmable environments and artifacts. The use of commitments gives a normative characterization to coordination artifacts, while the use of artifacts enables the application of software engineering methodologies to protocols.*

## 1 Vision and motivation

The growing pervasiveness of computer networks and of internet is an important catalyst pushing towards the realization of business-to-business, cross-business or, more generally, of *open environment* systems, that are made of heterogeneous and antecedently existing entities, which integrate their capabilities, interact according to some agreed contracts, and cooperatively exploit resources. The integration based on the classical notion of control flow of their softwares, even by means of orchestration languages, does not accomplish the objective of flexibility, of software reuse and of modular development that such systems require. This reality demands abstractions and models where the involved entities are fully *autonomous*. From a software engineering perspective, we think it is necessary to move from the notion of “integration” to a notion of “coordination”, obtained by introducing a social dimension which realizes a new *interaction-centric* approach.

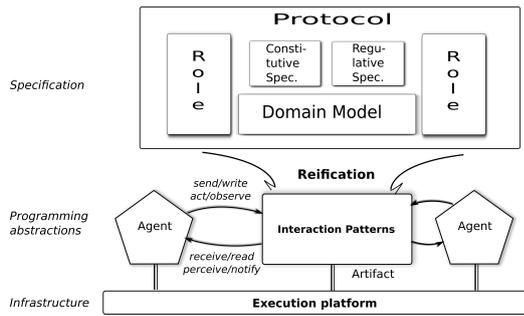
Even though Multi-Agent Systems (MAS) could be the most proper reference paradigm for developing open environment systems, because they adopt an interaction-oriented paradigm, current platforms fail, in our opinion, in that they are too much *content-centric*, and do not account for the social aspects of communication (recognized as fundamental by the research community [3]).

Interaction creates *social expectations* and *bindings/dependencies* in the involved partners [2, 15]. These

should be explicitly accounted for by the agent platform to allow the coordination of autonomous entities. In order to create social expectations on the agents’ behavior, it is necessary to introduce a *normative* characterization of coordination and give a social meaning to their actions. An agent that understands such a specification and that publicly accepts the regulation (i.e. that declares it will behave according to it) allows reasoning about its behavior [6]. This is the key to the development of open environment systems, made of autonomous and heterogeneous components. By not supplying such abstractions, current platforms do not supply agents the means for *observing* or *reasoning* about such meanings of interaction, and do not supply the designers the means to explicitly *express* and *characterize* them when developing an interaction model.

In the literature there are proposals that in principle could explicitly provide the abstractions that we need. The most prominent is the Agents & Artifacts meta-model (A&A) [17, 12], which provides abstractions for *environments* and *artifacts*, that can be acted upon, observed, perceived, notified, and so on. Our proposal is to use them to *reify regulations aimed at coordination*, in a way that agents can examine them (e.g. to decide whether playing one of the foreseen roles), use them (which entails that they explicitly accept the corresponding regulation), construct them e.g. by negotiation, specialize them, compose them, and so forth. In particular, to supply a normative characterization of coordination, we propose to use *commitment-based* approaches, which feature a *social* and *observational* semantics [15, 16, 18]. The use of commitment-based approaches will also provide the semantics of communication that the abstractions of artifact and environment still lack of, thus allowing the realization of both direct and indirect forms of communication (the latter are not supported by most of current platforms). Last but not least, a communication artifact could implement monitoring functionalities for the verification that the on-going interactions respect the regulation, for detecting violations and violators, and so forth.

The distributed nature of artifacts and their compositionality makes our proposal suitable to open environments,



**Figure 1. The schema of our proposal.**

while the regulatory, commitment-based approach to interaction and coordination exactly captures the social aspects on which the coordination of a group of autonomous partners should be based.

## 2 An interaction-centric framework for Social Computing

In our vision, coordination patterns, interaction protocols and, more generally, behavioral conventions or agreements are represented as commitment-based protocols and encoded into artifacts, conceived so as to implement programmable communication media, having a normative characterization. This conjugation helps providing protocol specifications some good and important features, from a Software Engineering point of view:

**abstraction** - the main concepts used to define environments, e.g. artifacts, are first-class entities; the interaction with agents is built around the agent-based concepts of action and perception (use and observation). As a consequence protocols become first-class abstractions as well, that agents can observe, use, handle.

**modularity and encapsulation** - it provides an explicit way to modularize the protocol, since artifacts are components representing units of functionality, encapsulating a partially-observable state and operations;

**extensibility and adaptation** - it provides a direct support for protocol extensibility and adaptation, since artifacts can be dynamically constructed (instantiated), disposed, replaced, and adapted by agents;

**reusability** - it promotes the definition of types of protocol that can be reused, such as in the case of coordination artifacts empowering agent interaction and coordination, blackboards and synchronizers.

The reification of protocols and coordination patterns means that practically all interactions among agents will be indirect, even in the case agents exchange messages. This

is actually what happens in the real world (words need air, paper, letterboxes or other media to be delivered). The advantage is that indirect communication allows more variegated ways of interacting, not hindering message exchange nor to rely on the message exchange metaphor when necessary. As Keil and Goldin observed, *indirect communication fosters the collaboration and the coordination inside open systems*, in that “it allows anonymous, non-specialized interfaces that invite participation of new knowledge sources” [10]. Indirect communication can be realized by means of persistent observable state changes, as it is for instance done by stigmergic approaches. Think, for instance, to a business interaction where a client pays for some item into the paypal account of the merchant: it should not be necessary that the client also informs the merchant about the payment, because the merchant can either check the paypal account or it could be notified by the paypal account itself. Even though the latter solution involves sending a piece of information, it would be unnatural to model a paypal account as an agent. A paypal account is an object in the environment of the merchant and of the client that is used by both agents, which is characterized by a programmed behavior, and that, to some extent, acts as a communication and coordination channel between the two (once the merchant verifies the payment was done, he/she can ship the sold items). This kind of components is to be captured by proper abstractions in the interaction model. Figure 1 draws the overall picture of the proposal, identifying three levels.

**Specification Level.** The specification level allows the designer to shape the interactions that will characterize the system by supplying adequate high-level abstractions. To this aim we propose to rely upon commitment-based protocols [15, 16, 18]. By relying on an *observational semantics*, commitment-based approaches can cope both with direct forms of communication (by supporting the implementation of communicative acts), and with forms of interaction, that are mediated by the environment (by supporting the implementation of non-communicative acts, having a social meaning). Agents can not only send and receive messages but they can also *act upon* or *perceive* the social state.

In order to support the specification of regulations, capturing conventions, laws, and the like, we rely on the proposal in [1, 11], that allows the representation of legal patterns of interaction by enriching commitment protocols with temporal regulations. In particular, [1] proposes a decoupled approach that separates a constitutive and a regulative specification. This modularization brings about many advantages: easier re-use of actions, easier customization, greater compositionality.

In [1] the constitutive specification defines the meaning of actions based on their effects on the social state, the regulative specification reinforces the regulative nature of

commitment by adding a set of behavioral rules, given in terms of *temporal constraints* among commitments (and literals). These constraints define schemes on how commitments must be satisfied, defining the legal evolution of the social state, independently from the executed actions. Since interactions are observable and their semantics is shared, each agent is able to draw conclusions concerning the behavior of the partners or concerning the system as a whole.

**Programming Abstractions Level.** This layer realizes at a programming language level the abstractions defined above. This is done by incorporating interaction protocols based on commitments, patterns of interaction, forms of direct and indirect communication and coordination between agents (such as stigmergic coordination) inside the programmable environments envisaged by the A&A meta-model [17, 12]. To this aim, it is necessary to integrate, inside such environments the tools that are necessary to handle commitments (e.g. commitment machines [18]) and to reason about the evolution of the social state. We mean to use the public interface of artifacts to allow agents to examine the encoded coordination rules. As a consequence, the act of using an artifact can be interpreted as a declaration of acceptance of the coordination rules. This will generate social expectations about the agent's behavior and this agrees with the characterization of norms in [6]. Moreover, the fact that the behavior of agents on artifacts is observable and that interactions only occur through artifacts, agrees with the view that regulations can only concern observable behavior [7]. The resulting programmable environments will provide flexible *communication channels* that are specifically suitable for open systems. Notice that the use of a programming environment does not mean that the social state will necessarily be centralized in the physical system: let us recall that an artifact can be composed by a distributed network of artifacts. The notion of environment has always played a key role in the context of MAS; recently, it started to be considered as a first-class abstraction useful for the design and the engineering of MAS [17]. A&A follows this perspective, being a meta-model rooted upon Activity Theory and Computer Support Cooperative Work that defines the main abstractions for modeling a MAS, and in particular for modeling the environment in which a MAS is situated. A&A promotes a vision of an endogenous environment, that is a sort of software/computational environment, part of the MAS, that encapsulates the set of tools and resources useful/required by agents during the execution of their activities. A&A introduces the notion of artifact as the fundamental abstraction used for modeling the resources and the tools that populate the MAS environment.

**Infrastructure Level.** In the state of the art numerous applications of the endogenous environments, i.e. environ-

ments used as a computational support for the agents' activities, have been explored, including coordination artifacts, artifacts used for realizing argumentation by means of proper coordination mechanisms, artifacts used for realizing stigmergic coordination mechanisms, organizational artifacts. Our starting point is the CArtAgO framework [14]. CArtAgO provides the basis for the engineering of MAS environments on the base of: (i) a proper computational model and (ii) a programming model for the design and the development of the environments on the base of the A&A meta-model.

### 3 Discussion and conclusions

This paper proposes a framework for open environment systems that: (i) is *interaction-centric*; (ii) captures the *social meaning of interactions* and not only the meaning related to the content of communication; (iii) foresees a notion of *programmable communication channel*, realized by means of artifacts, which can also be empowered with *monitoring* functionalities; (iv) the semantics of communication is supplied declaratively by a *commitment-based* approach, and this is an improvement w.r.t. what is done by platforms and frameworks, like JADE, Jadex, 2APL, which do not integrate the semantics of the adopted ACL but only guarantee a syntactic interoperability; (v) the use of commitments gives a *normative value* to the encoded protocol; (vi) the act of using a communication artifact amounts to the *explicit acceptance* of the rules of the protocol; the proposal allows the interaction of any kind of agent. The proposal conjugates the flexibility and the openness that are typical of MAS with the needs of modularity and compositionality that are typical of design and development methodologies. In particular, relying on artifacts allows the application of object-oriented software engineering methodologies to the specification and composition of interaction protocols. For instance, it makes it simple to structure the protocol into a constitutive part and a regulative part as well as to build a new protocol by assembling previously existing ones. As such, our proposal shows the properties hoped for in recent works on patterns like [5]. Notice that the compositionality we think of is not procedural due to the declarative nature of commitments.

The realization of commitment protocols as artifacts/environments (first-class objects of the model) which are capable to monitor the interaction, taking into account also the regulative specification, represents an advancement of research on commitment-based approaches. These features are advantageous w.r.t. approaches like [4], where these elements reside in the middleware, and are therefore shielded from the agents and from the designer.

Although our proposal is more general than e-institutions, it is interesting to comment some work car-

ried on in this context. The current proposals for electronic institutions do not supply yet all of the solutions that we need: either they do not account for indirect forms of communication or they lack mechanisms for allowing the a priori verification of global properties of the interaction. As [8] witness, there is an emerging need of defining a more abstract notion of action, which is not limited to direct speech acts, whose use is not always natural. For what concerns organizations, instead, there are some attempts to integrate them with artifacts, e.g. ORA4MAS [9] and JaCaMo <http://jacamo.sourceforge.net>, which also accounts for BDI agents. Following the A&A perspective, artifacts are concrete bricks used to structure the agents' world: part of which is the organizational infrastructure, part amounts to artifacts introduced by specific MAS applications, including entities/services belonging to the external environment. In [9] the organizational infrastructure is based on *Moise*<sup>+</sup>, which allows both for the enforcement and the regimentation of the rules of the organization. This is done by defining a set of conditions to be achieved and the roles that are permitted or obliged to perform them. The limit of this approach is that it cannot capture contexts in which regulations are, more generally, norms because norms cannot be restricted to achievement goals.

For what concerns the infrastructure, the key idea is the definition of environments based on A&A and on CArTAgo, that reifies the proposed model of interaction protocols. Among the specific results related to this, we foresee an advancement of the state of the art with respect to the definition and the exploitation of forms of stigmergic coordination [13] in the context of intelligent agent systems.

We think that our proposal will give significant contributions to industrial applicative contexts. Among the most interesting examples are the integration and the cooperation of e-Government applications (services) spread over the nation. In this context, the aim is to verify the adherence of bureaucratic procedures, of the public administration, to the current laws (e.g. <http://www.ict4law.org>). Another interesting application regards (Web) services. Some of the fundamental aspects promoted by the SOA model, such as autonomy and decoupling, are addressed in a natural way by the agent-oriented paradigm. The development and analysis of service-oriented systems can benefit from the increased level of abstraction offered by agents, by reducing the gap between the modeling, design, development, and implementation phases. In this context, it is necessary to deploy complex interactions having those characteristics of flexibility that agents are able to guarantee.

## References

[1] M. Baldoni, C. Baroglio, E. Marengo, and V. Patti. Constitutive and Regulative Specifications of Commitment Proto-

cols: a Decoupled Approach. *ACM Trans. on Int. Sys. and Tech., Spec. Iss. on Agent Communication*, 2011. To appear.

[2] C. Castelfranchi. Principles of Individual Social Action. In G. Holmstrom-Hintikka and R. Tuomela, editors, *Contemporary action theory: Social action*, volume 2, pages 163–192. Dordrecht, 1997. Kluwer.

[3] A. K. Chopra, A. Artikis, J. Bentahar, M. Colombetti, F. Dignum, N. Fornara, A. J. I. Jones, M. P. Singh, and P. Yolum. Research Directions in Agent Communication. *ACM Trans. on Int. Sys. and Tech., Spec. Iss. on Agent Communication*, 2011. To appear.

[4] A. K. Chopra and M. P. Singh. An Architecture for Multi-agent Systems: An Approach Based on Commitments. In *Proc. of ProMAS*, 2009.

[5] A. K. Chopra and M. P. Singh. Specifying and applying commitment-based business patterns. In *Proc. of AAMAS. IFAAMAS*, 2011.

[6] R. Conte, C. Castelfranchi, and F. Dignum. Autonomous Norm Acceptance. In *Proc. of ATAL*, volume 1555 of *LNCS*, pages 99–112. Springer, 1998.

[7] M. Dastani, D. Grossi, M. J.-J. Ch., and N. A. M. Tinemeyer. Normative Multi-agent Programs and Their Logics. In *KRAMAS*, pages 16–31, 2008.

[8] N. Fornara, F. Viganò, and M. Colombetti. Agent communication and artificial institutions. *Autonomous Agents and Multi-Agent Systems*, 14(2):121–142, 2007.

[9] J. F. Hubner, O. Boissier, R. Kitio, and A. Ricci. Instrumenting multi-agent organisations with organisational artifacts and agents: “Giving the organisational power back to the agents”. *Autonomous Agents and Multi-Agent Systems*, 20, 2009.

[10] D. Keil and D. Goldin. Modeling indirect interaction in open computational systems. In *Proc. of TAPOCS*, pages 355–360. IEEE Press, 2003.

[11] E. Marengo, M. Baldoni, C. Baroglio, A. K. Chopra, V. Patti, and M. P. Singh. Commitments with Regulations: Reasoning about Safety and Control in REGULA. In *Proc. of AAMAS*, 2011.

[12] A. Omicini, A. Ricci, and M. Viroli. Artifacts in the A&A meta-model for multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 17(3):432–456, 2008.

[13] A. Ricci, A. Omicini, M. Viroli, L. Gardelli, and E. Oliva. Cognitive stigmergy: Towards a framework based on agents and artifacts. In *Proc. of EAMAS*, volume 4389 of *LNCS*, pages 124–140. Springer, 2007.

[14] A. Ricci, M. Piunti, M. Viroli, and A. Omicini. Environment Programming in CArTAgo. In *Multi-Agent Programming II: Languages, Platforms and Applications, Multiagent Systems, Artificial Societies, and Simulated Organizations*. 2009.

[15] M. P. Singh. An Ontology for Commitments in Multiagent Systems. *Artif. Intell. Law*, 7(1):97–113, 1999.

[16] M. P. Singh. A social semantics for agent communication languages. In *Issues in Agent Communication*, volume 1916 of *LNCS*, pages 31–45. Springer, 2000.

[17] D. Weyns, A. Omicini, and J. Odell. Environment as a first class abstraction in multiagent systems. *Autonomous Agents and Multi-Agent Systems*, 14(1):5–30, 2007.

[18] P. Yolum and M. P. Singh. Commitment machines. In *Proc. of ATAL*, pages 235–247, 2001.