An Architecture Proposal for Human-Agent Societies

Holger Billhardt[1], Vicente Julián[2], Juan Manuel Corchado[3], and Alberto Fernández[1]

[1] CETINIA, Universidad Rey Juan Carlos, Spain
{holger.billhardt,alberto.fernandez}@urjc.es
[2] DSIC, Universidad Politécnica de Valencia, Valencia, Spain
vinglada@dsic.upv.es
[3] BISITE, Universidad de Salamanca, Salamanca, Spain
corchado@usal.es

Abstract. Agreement technologies have settled the basis for creating systems that operate on the basis of agreements in societies of independent, autonomous computational entities (agents). However, nowadays more and more systems of such kind rely on a seamless interaction of software agents with humans. Humans work in partnership (directly or indirectly) or closely related with agents that are able to act autonomously and intelligently. Specifically, humans and agents have the ability to establish a series of relationships/collaborative interactions with each other, forming what might be called human-agent teams to meet their individual or collective goals within an organisation or social structure. Systems in which people and agents operate on a large scale offer an enormous potential but also require the consideration of additional issues. In this paper we analyse the open issues that may be addressed for researches in order to develop open human-agent systems. We present a real-world case study and an abstract architecture proposal for such systems.

Keywords: Multi-agent systems, Human-agent societies, Service-oriented multi-agent systems.

1 Introduction

Nowadays more and more humans work in partnership (directly or indirectly) or closely related with computational entities (agents) that are able to act autonomously and intelligently. Considering systems of people and agents operating on a large scale offers an enormous potential and, if performed properly, it will help tackle complex social applications.

As intelligent computational systems pervade the human space, they dramatically change the ways in which human users interact with technological systems. Rather than performing tasks in an isolated and subordinate fashion, computational agents are involved in an increasing range of collaborative relationships with humans, where negotiation and task delegation occur along the agent-agent axis as well as along the human-agent one. Due to its social ground, agreement technologies can be the basic driver to smooth the boundary between the human space and the computational space so as to create cooperative social ensembles of humans and agents that work in strict

J.M. Corchado et al. (Eds.): PAAMS 2014 Workshops, CCIS 430, pp. 344–357, 2014.

partnership to accomplish their individual or collective goals. However, they have to be enhanced in order to take into account the particular requirements that arise when considering interactions among computational agents and humans. Adaptation of the normative context, adjustable and limited autonomy and recognition of user's intentions are some of the characteristics that have to be covered. Semantics, norms, negotiation, learning, context awareness, behavioural modelling and human-agent collaboration will be the building blocks for the specification of such systems.

We call such systems human-agent societies and the research questions to be solved in order to create such systems are: What type of environment or technological platform is appropriate for enabling a computing ecosystem of that kind? What type of services, standards, methods and tools are needed to generate such environment in a way that both, computational agents and humans can seamlessly interact with each other, each of which solving its individual objectives?

The view of human-agent societies also reflects the technological evolution in the areas of Computer Technology and Communications (Internet, WWW, e-commerce, wireless connectivity, etc.). Computation has become an inherently social activity that occurs by means of and through communication among entities (agents or humans) and takes place in our daily lives when and wherever it is needed. Applications are no longer monolithic or distributed applications managed by only one organisation, but rather societies of actors that provide and consume services. Even the notion of application is losing its importance. Instead, computing is something that is enabled by an environment – a virtual ecosystem of intelligent components – in which we are immersed and where complex behaviour is created on the fly. The concepts of "service oriented computing", "computing by interaction and agreement", "cloud computing", the "internet of things" or "social computing" are closely related for such systems, but they have to be enhanced to make computing a truly seamless experience of immersion in such an environment. Possible scenarios of such systems are virtual marketplaces where either agents or humans interact, simulation and training environments (possibly in 3D), the area of health and medical applications, home automation, etc.

In this context, we need to define platforms, as well as development methods and tools that support the creation, operation and evolution of such an ecosystem and allow for a seamless and integrated interoperability of agents, services and humans. This paper presents an analysis of the main open issues related with such systems and which we belief must be taken into account in the research in this area. The analysis is supported by the study of a real-world case: a tourism scenario. Based on this analysis a proposal for an abstract architecture for human-agent societies is presented.

The rest of the paper is structured as follows: section 2 defines human-agent societies; section 3 presents the analysis of a case study: a tourism scenario; section 4 enumerates open issues for research in the field; section 5 proposes an abstract architecture for human-agent societies and, finally, section 6 gives some conclusions.

2 Human-Agent Societies

With the term human-agent society we refer to a computing paradigm in which the traditional notion of application disappears. Rather than developing software applications that accomplish computational tasks for specific purposes, this paradigm is based on an immersion of the users in an environment that enables computation. The environment itself is populated by computational entities with different capacities and intellectual properties, ranging from simple devices that offer specific capacities or rudimentary information, like screens or sensors, to autonomous artificial agents that provide high level services and are able to engage in complex interaction protocols.

Humans participate in this society either through some personal assistance agent or just through some graphical interface on different computing devices. From the point of view of the inside of the society, humans participate just like any other agent, trying to fulfil their individual goals and objectives. From the point of view of the humans, the interactions follow the principles of ubiquitous computing. That is, humans have the perception to interact with objects they are used to and that recognize their contexts and personal needs and preferences. In this way, the immersion of humans in the system is maximized, minimizing at the same time the difficulties of interaction.

Human-agent societies are intrinsically open; participant can enter and leave the system at any time and there is no direct control over the behaviour of the participating components. Computation in such societies is not predefined. Instead it takes place as a consequence of the fact that autonomous agents try to meet their objectives. In order to do this, they interact with others, establishing relations of different kind and different temporal duration; from the simple provisioning of a piece of information, to long lasting collaborations in some business scenario, or connections based on friendship.

In our vision, human-agent societies are an evolution and integration of different paradigms and technologies that have appeared in the last decades.

In the first place, they are an evolution of multi-agent systems (MAS). In the last years, the research on MAS has been centred primarily on open system [1]; systems that are populated by heterogeneous agents that can dynamically enter and leave the system at any time. They are characterized by their limited confidence, possibly conflicting individual goals and a great probability of non-accordance with specifications [2]. All these characteristics are also present in our vision of human-agent societies. In particular, research in open MAS has concentrated on methods for regulating the interactions among the agents. In order to assure that a system with an a priori unknown population of agents will behave according to some global objectives and preferences, researchers have studied the regulation of artificial agent societies upon the contractual and normative relations found in human social interactions and the concept of organisation has become very important [3]. This concept changes the focus in the design of MAS from an agent-centred approach to an organisation-oriented approach where the problem consists in designing the rules of the game rather than the individual components. A wide variety of more or less flexible meta models for specifying virtual organisations have been proposed. Some examples are AGR [3], MESSAGE [4], GAIA and Roadmap [1,5], Electronic Institutions [6,7], OMNI [8], MOISE+ [9], GORMAS [10] and THOMAS[11] or the model proposed in [12]. Other approaches propose to regulate open MAS through more lightweight coordination mechanisms or artifacts that can be embedded in the environment [13-16].

In recent years, the development of agreement technologies has settled the basis for a paradigm of "computing by agreement", where intelligent computational agents autonomously operate and interact among them in a sophisticated way, and whose higher level of intelligence permits to delegate more and more complex tasks [17]. Semantics, norms, organisations, negotiation, argumentation, coordination and trust have been identified as the enabling technologies for such systems. We belief that these are also some of the building blocks for human-agent societies, but they have to be enhanced beyond the computational space for the development of socially acceptable systems in which humans represent just another component.

In the second place, human-agent societies are based on the developments in the field of service oriented computing (SOC) and semantic technologies. Service oriented computing is a field that has evolved in the past in parallel to MAS, but both have a principal aspect in common: the idea of computing by delegation; of creating complex behaviour by composing capacities (services) or information provided by different independent and possibly heterogeneous entities. However, whereas MAS has more concentrated on the autonomy and intelligent behaviour of the components, research in SOC has been rather concentrated on appropriate descriptions that facilitate and enable interoperability, service discovery and service composition. All these issues are relevant for human-agent societies. Thus, the languages and methods that have been developed in this field can provide a starting point for the description of capacities of the different components. Especially of interest are languages for semantic descriptions of web services like OWL-S [18] or WSMO [19], or the more lightweight approaches like SAWSDL [20], SA-REST [21] and WSMO-Lite [22]. However, such languages have to be extended in two ways. First, it is not only necessary to describe the services any component is able to provide to others, but also their intellectual and social capacities – their abilities to negotiate, argue and establish agreements, their social constrains and relationships and their preferences and contextual situations. Some approaches in this sense have already been proposed in the context of Service Oriented Multi-Agent Systems (e.g., [23, 24]. The second necessary extension of current semantic service description languages is related to interoperability. It is necessary to find mechanisms for a transparent access of humans and artificial agents to other components. This implies the need for bridging the gap between ontology based descriptions and easily understandable representations for humans. Some initial ideas in this sense have been proposed for service search and discovery based on key words or tag clouds (e.g., [25, 26]).

Human-agent societies do not only rely on the use of services or capacities other provide, but also on the transparent and integrated access to the available information regardless the location of the actual information sources or the representational structure of information pieces. Semantic technologies have been used in recent years to link data and information sources, both at the instance and the meta level. From these efforts emerges the Linked Data[1] initiative whose objective is to share and interconnect structural data in Internet. The integration of data and information from different possibly heterogeneous sources require the alignment of different languages and

[1] http://linkeddata.org

vocabularies used to describe the data schemata and the fusion of data at the instance level. RDF Schema and OWL provide mechanisms to define relationships among different vocabularies and that allow transformations between schemata in an easy way. The openness of the envisioned systems however, requires methods that allow obtaining such alignments in an automatic manner (e.g., [27, 28]).

Finally, the third pillar for constructing human-agent societies arises from the field of ambient intelligence and context aware systems. Ambient intelligence has the same perspective of an environment of computational devices that support humans and in which humans are immersed in their daily lives [29]. Such an environment should be able to perceive and respond to the presence of individuals in a simple, non-intrusive and often invisible manner [30]. Ambient intelligence is closely related to ubiquitous computing, a concept that has been introduced in the 90th [31], and that describes a world where different computational objects communicate and interact with the aim to support the daily activities of humans. Research in this field has concentrated on the following issues [32], all of which are also relevant for human-agent societies: i) the technologies should be transparent for users, ii) service provision should adapt to the context and preferences of users, and iii) the systems have to provide user friendly interfaces. Intelligent systems and Multi-agent technologies are considered to play a fundamental role in meeting these requirements [33].

The concept of context based systems, introduced by Want et al. [34], deals with the issue of adapting the provision of services to the context of users. Context aware applications been created in many different domains (e.g., [35-39]). In most of those applications, however, the definition of context is restricted to some particular attributes. Context should be something that is not predefined. What should be considered as context depends on the interaction a component is carrying out in a particular moment. We belief that the use of agents with higher intellectual capacities (e.g., adaptation, learning, etc.) as core components in human-agents societies will facilitate the creation of systems with a higher degree of context awareness. Moreover, this will also make possible to provide and obtain additional contextual information, like social relationships, or social and institutional roles, for example.

3 Case Study Analysis

A case study of a tourism market is used to illustrate the research requirements of human-agent societies. However, it is important to point out that the requirements identified are independent of the domain of the example. A tourism market can be viewed as a large dynamic problem that interconnects clients (individuals) and suppliers – a huge number of public and private institutions (companies, travel agencies, hotel chains, individual hotels, airlines, …) – of tourism services. Nowadays, many of such services include additional context-based facilities to adapt to the age, genre, acquisition capacity and interest of users and, thus, to fulfil the final users needs.

In a concrete situation, let's suppose a group of friends wants to organize a holiday trip together. As a group, they have to negotiate with several service providers (e.g., for accommodation, travel arrangements, leisure activities) to get a travel package for the group. But at the same time they need to deal with possible conflicts of interest since each member of the group will have his/her own desires and preferences. Thus, they have to agree upon a package that would somehow satisfy all members of the group. Considering human negotiation, such a task cannot be done without a great deal of effort. It would be time consuming and, context related issues, like geographical distance, different time schedules or idiomatic barriers would complicate the required negotiation and argumentation processes.

Now let's suppose that this problem is solved in a human-agent society, in which each friend of the group would be represented by his/her electronic assistant agent (EAA) that will act in the society on behalf of its user. Moreover, bigger tourist companies will have software agents that offer their products and smaller providers (e.g., small hotels) may also be represented by assistant agents. The assistant agents should be aware of the personal preferences of their users regarding booking conditions. Furthermore, they should also be able to recognize certain context parameters of their users that might be of interest (for instance, one of the friend may have his birthday during the trip). In a first stage, the EAAs would need to create a team whose aim is to collaboratively solve the problem at hand. The system requires mechanisms for searching for different tourist services, hiding away any barriers regarding possible heterogeneous descriptions of such services. In an open system as the one described here, there exist an intrinsic uncertainty about the reliability of service providers, and trust and reputation mechanisms may play an important role in order to assure successful interactions. In our case, a particular hotels may be discarded by the group because one of the members has had a bad experience in this hotel in a previous trip.

The society should provide facilities that allow the EAAs to negotiate, to argue and to establish agreements/contracts with agents that offer tourist services. It should also facilitate the set up of some type or organisational structure that defines the role of each EAA in the group, as well as certain norms or protocols that regulate the negotiation and argumentation process among the EAA within the group in an efficient way. Such norms and protocols should assure the best possible outcome, e.g., a tourist package with which all group members will be happy. Re-organisation should be possible, that is, the system should be adaptable, for instance, if another friend wants to join the trip or if a new offers is available. The society should allow or even may require the human intervention in different moments. Some of the friend may relay totally on the autonomy of their EAAs, but others may want to participate in certain decisions, or may want to be informed on the decision progress. The society should be able to deal with such situations, adapting its operation to the different intervention requirements of the human participants. It should allow the participation of humans through different devices and interfaces and provide the required information to human agents in a human-readable manner.

4 Open Issues

Extending current technologies in order to develop a computational ecosystem that allows for a seamless integration of humans, artificial agents and other computation components poses several challenging problems. As we have tried to demonstrate in section 3, the complexity of such systems is high and there are several open issues for research that should be tackled. In this section we try to identify such issues.

Human-agent collaboration, adjustable and limited autonomy.
In order to obtain a true benefit of the collaboration of humans and computational agents, it is necessary to understand the capabilities and capacities of both components and how these capacities can complement each other in computational systems. Collaborative problem solving and decision-making through collective intelligence and crowdsourcing is an issue to be studied. We need to focus on how humans and computational agents define the agreement space that permits them to collaborate to solve complex problems. Human-agent planning, negotiation, argumentation, mixed-initiatives, teamwork and team monitoring are key technologies that are required.

Within this issue, the concept of adjustable autonomy is of importance, as a way for agents to decide whether to take a decision autonomously or to delegate it to a human user [40,41]. But this concept needs to be extended to what could be called limited autonomy, integrating the fact that humans might want to have the final say on certain operations and decisions made by intelligent agents. Computational agents must be provided with effective protocols that allow them to interact in the system without surpassing the limits of autonomy and authorization that have been imposed on them. This, on the other hand, will require the adaptation of the (normative or organizational) contexts under which interactions take place.

Semantics and information.
The basis of effective and smooth partnership between human and computational agents depends on the transmission of information and knowledge across the boundary between the computational space and the human space. Semantics, knowledge representation and argumentation are the basis of the mutual understanding and knowledge exchange. The development of mechanisms to reliably blend information generated by users and computational agents running on dispersed digital devices and sensors is a topic that has to be analysed and studied. Means are required to transparently provide information and knowledge in a way understandable for each participating entity in the system. Semantic alignment techniques, as proposed for agent-agent interactions [42, 43] and approaches of explaining information that is passed to other entities [44, 45] can provide a basis but have to be extended to provide support for human-agent interactions.

Context awareness and user preferences.
Conveying the information related to human preferences, contexts and objectives, and understanding these pieces of knowledge by computational agents is essential for creating socially acceptable systems in which humans participate in a transparent way. In human-agent societies, where the interactions take place between components that are unknown “a priori” there is a need for creating an infrastructure that allows to obtain any contextual information regarding a given interaction “on the fly”. Such an

infrastructure should allow any component to declare/publish contextual information and also to search for contextual information in a transparent manner.

Context-aware system uses contextual information to modify their behaviour to meet user needs better. Although current techniques provide a solid foundation to model and interpret the context, there are aspects that have not been covered completely. In particular, due to the huge quantity of information that is generated continuously, new data fusion algorithms and mechanism are required that are able to filter and process this information, and to generate knowledge.

Normative contexts, regulation and trust.

Systems where different autonomous entities, with possibly different individual objects coexist, require normative contexts or other regulation mechanisms that rule the interactions between the entities. Such context may specify the norms and protocols that should be followed in a particular interaction scenario of a limited duration. But they may also be used to set up long-term relationships between components in the sense of virtual organisations or electronic institutions. The enabling infrastructure for human-agent societies should facilitate the creation of normative contexts for both, short-term interactions and long-term organisations. Due to the intrinsic openness of the envisioned systems (e.g., participants in the society can‘t be controlled directly), especial importance should be given to the development of "soft" regulation mechanisms (incentives, penalties, social acceptance, etc.). In many situations, such mechanisms may be the way to drive the evolution of the system behaviour towards a desirable global outcome, coping at the same time with individual objectives of humans and computational agents.

Also the concepts of trust and reputation as a means for self-regulation and security in open systems are important. Different trust mechanisms have been proposed in the agent community but it is necessary to adapt such methods to situations where both, humans and agents are involved. Humans are not always rational and often establish trust based on their social relationships. In order to employ reliable trust and reputation mechanisms, the underlying computational infrastructure has to provide means to identify individuals in the system regardless the device or interface they are using.

Adaptation and reorganisation.

The world that surrounds us always changes and evolves. Computational human-agent societies, in order to be socially accepted and to persist in time, have to adapt their behaviour to such changes in a natural way.

At the micro level, computational agents should be able to anticipate and adapt to the context of the world outside the computational scope, as well as to the goals and intentions of human user. Learning by human feedback, intention recognition and behavioural pattern classification are some of the topics that should be studied.

At the macro level, reorganisation at the group or society level, allows to cope with environmental changes. That is, the normative contexts and the organisational structures that govern the interactions among entities should evolve in order to adapt the behaviour of a system to a dynamic environment. Key issues to be studied here are the factors that should trigger reorganization (when), and which changes in the organisational structures will have the desired effect on the adaptation of the system's behaviour (how). Furthermore, reorganisation should take place dynamically, with minimal interference and smooth transitions, and without sacrificing efficiency.

Discovery and composition.
Participants in human-agent societies use the services and capacities that others provide. In this regard, the envisioned ecosystem has to provide means to adequately describe, locate, use and compose services or capacities. Much work in this regard has been done in the Web services community. However, these approaches are not directly transferable to human-agent societies. In service-oriented computing the core component, in general, is the service, and not the service provider. In human-agent societies, individuals should be the core component for description and location. Apart from providing certain services, each individual is characterized by other factors, like its social relationships, roles, reliability, objectives, preferences, etc. and all these characteristics will have to be taken into account when interacting with others. Research should be carried out to define new methods for describing, locating individuals and the services they provide as well as new methods for composing the services provided by several agents or humans. Furthermore, the infrastructure for searching and locating individuals should be decentralized to cope with scalability issues and it should be accessible from different devices with different capacities.

Computational platform.
Human-agent societies rely on the seamless interactions of humans and computational entities based on network connections through physical devices with different computational capacities. A computational platform is needed to cope with this situation. Such a platform should be scalable, facilitating the connection to the society through the network and regardless the devices or interfaces used. Furthermore, it should facilitate the participation of entities (humans or software agents) in a transparent manner, making the development, publication, provision and usage of services an easy and transparent task. From the outside, the platform should provide "universal" computational power, hiding away any problems related to physical restrictions of devices. We believe that such a platform has to be decentralized. It should rely on possibly many computational resources that distribute computation tasks among themselves in a decentralized way and based on parameters like workload, importance, etc.

5 Abstract Architecture for Human-Agent Societies

According to the previous analysis, we propose an abstract architecture for human-agent societies, called iHAS (intelligent Human-Agent Societies). The architecture intents to tackle all the above-mentioned open issues. A general view is shown in Figure 1.

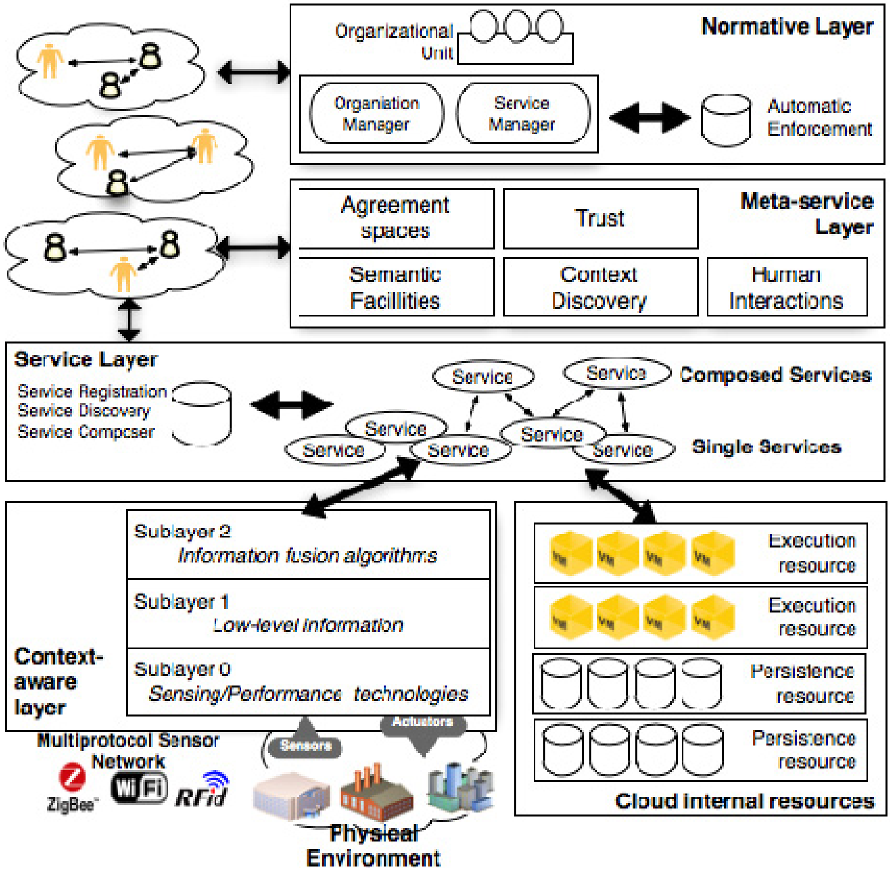

Fig. 1. iHAS abstract architecture proposal

The iHAS framework is organised as a set of layers, which offers Agents or Human access to the iHAS infrastructure. We envision the service as the basic computational unit for interaction. Artificial and human agents may provide services for others and interact with each other through the execution of services. On the lower level, a cloud-like computational platform, enhanced with context aware facilities, enables a transparent deployment of services. On top, a variety of high-level facilities are offered through the Normative and Meta-service layers which might be used by agents in order to set up more complex societies possibly stable in time. In the following we give a detailed description of each layer:

1. *Low-level layers*: it is formed by the *Cloud internal resources*, where services are deployed providing a virtual hosting service with automatic scaling according with reached service level agreements (SLA) and functions for balancing workload. This layer is internal to the system and is formed by the physical resources, which allows the abstraction of these resources shaped as virtual machines (VM).

In order to allow a more realistic representation of the physical world, a Context-aware layer is added that facilitates the integration of pervasive components in the system. Such components are connected trough hybrid Wireless Sensor Network (WSN) and allow the immersion of human actors in the system through the employment of mobile or embedded devices. The layer is divided in sub-layers following the JDL fusion model [46] in order to refine raw data by means of information fusion techniques.

2. *Service layer:* provides the system functionality, offering services that agents provide and can use for interacting with each other. Services can be single or can be formed through the composition of other services. The layer contains a set of (possibly federated) *directory services* that allow for registering and locating services based on given service descriptions. Such directories contain *matchmaking* functionalities that account for semantic search of services. The component contains composition meta-services endowed with an on-line planning functionality, capable of on-the-fly orchestration of services with regard to a given query.
3. *Meta-Service layer:* this layer offers a set of general high-level facilities that are of special interest for human-agent systems and may be used when needed. The main functionalities provided are *agreement* capabilities to deal with conflicts inside the society; *human interaction* facilities which will try to maximize the immersion of humans in the system; service discovery and *context-aware* mechanisms allowing a transparent and integrated access to the required information; *semantic* interoperability facilities; and, *trust and reputation* models, which will employ probabilistic techniques and psychological and sociological theories to estimate the behaviour of an entity (either human or artificial) in a given context.
4. *Normative layer:* might be used by human-agent societies that require a regulatory framework for the coordination, communication, and interaction among different entities (humans or software agents). This layer incorporates mechanisms for specifying, adapting and enforcing organizational constraints (e.g., norms and protocols) that should govern the interactions of entities and coordinate the use of resources. From a technical view, these functionalities can be obtained using the THOMAS platform [11], which consists basically of a set of modular services that enable the development of agent-based organisations in open environments.

6 Conclusions

In this paper an analysis of the main open issues related with human-agent societies is presented. These systems are based on a new paradigm where the key is the immersion of the users in a complex environment that enables computation.

These kinds of systems can be supported by existing platforms, development methods and tools that are appropriated for open systems, which are populated by heterogeneous agents that can dynamically enter and leave the system at any time. But beyond these, it is essential to evolve them and enable the interoperability of agents, services and humans, all the entities together as a unique ecosystem.

In a first place, the essential parts that compose and give sense to these societies must be improved, such as semantics, norms, negotiation, reorganisation, argumentation, delegation and other concepts already mentioned. Moreover, SOC technologies should be easily integrated, creating complex behaviours by composing capacities modelled as services or information provided by different independent and possibly heterogeneous entities. Another challenge is to achieve a complete human-agent society that could be able to perceive and respond to the presence of individuals in a simple, non-intrusive and often invisible manner. Advanced data fusion algorithms and context aware system techniques will be an important area of work to reach such requirements.

It is important to highlight that the final goal is obtaining an intelligent computational system in which, when pervading the human space, computational agents and human users interact and coexist successfully. As presented in this paper, our proposal is a framework called iHAS. This framework offers a variety of facilities to satisfy all the special requirements of human-agent societies and to guide designers of such systems though the integral development of them.

Acknowledgments. Work partially supported by the Spanish Government through the projects OVAMAH (grant TIN2009-13839-C03-02; co-funded by Plan E), AT (grant CSD2007-0022; CONSOLIDER-INGENIO 2010) and iHAS (grant TIN2012-36586-C01/C02/C03).

References

1. Zambonelli, F., Jennings, N., Wooldridge, M.: Developing multiagent systems: The GAIA methodology. ACM Trans. on Soft. Eng. and Methodology 12, 317–370 (2003)
2. Argente, E., Julian, V., Botti, V.: Multi-Agent System Development based on Organizations. Electronic Notes in Theoretical Computer Science 150(3), 55–71 (2006)
3. Ferber, J., Gutknecht, O., Michel, F.: From Agents to Organizations: An Organizational View of Multi-agent Systems. In: Giorgini, P., Müller, J.P., Odell, J. (eds.) AOSE 2003. LNCS, vol. 2935, pp. 214–230. Springer, Heidelberg (2004)
4. Caire, G., et al.: Agent Oriented Analysis Using Message/UML. In: Wooldridge, M.J., Weiß, G., Ciancarini, P. (eds.) AOSE 2001. LNCS, vol. 2222, p. 119. Springer, Heidelberg (2002)
5. Juan, T., Pierce, A., Sterling, L.: Roadmap: Extending the Gaia methodology for complex open systems. In: Proc. AAMAS 2002, pp. 3–10 (2002)
6. Esteva, M., Rodríguez-Aguilar, J.-A., Sierra, C., Garcia, P., Arcos, J.L.: On the Formal Specification of Electronic Institutions. In: Sierra, C., Dignum, F.P.M. (eds.) AgentLink 2000. LNCS (LNAI), vol. 1991, pp. 126–147. Springer, Heidelberg (2001)
7. 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)
8. Dignum, V., Vázquez-Salceda, J., Dignum, F.: OMNI: Introducing Social Structure, Norms and Ontologies into Agent Organizations. In: Bordini, R.H., Dastani, M., Dix, J., El Fallah Seghrouchni, A. (eds.) PROMAS 2004. LNCS (LNAI), vol. 3346, pp. 181–198. Springer, Heidelberg (2005)

9. Hübner, J., Sichman, J., Boissier, O.: MOISE+: Towards a Structural, Functional and Deontic Model for MAS Organizations. In: Proc. AAMAS, pp. 501–502 (2002)
10. Argente, E., Botti, V., Julian, V.: GORMAS: An organizational-oriented methodological guideline for open MAS. In: Gleizes, M.-P., Gomez-Sanz, J.J. (eds.) AOSE 2009. LNCS, vol. 6038, pp. 32–47. Springer, Heidelberg (2011)
11. Argente, E., Botti, V., Carrascosa, C., Giret, A., Julian, V., Rebollo, M.: An Abstract Architecture for Virtual Organizations: The THOMAS approach. Knowledge and Information Systems 29, 379–403 (2011)
12. Deloach, S.A., Oyenan, W.H., Matson, E.T.: A capabilities-based model for adaptive organizations. Autonomous Agents and Multi-Agent Systems 16(1), 13–56 (2008)
13. Omicini, A., Ricci, A., Viroli, M.: Artifacts in the a&a meta-model for multi-agent systems. Autonomous Agents and Multi-Agent Systems 17(3), 432–456 (2008)
14. Lopes Cardoso, H., Oliveira, E.: Social control in a normative framework: An adaptive deterrence approach. Web Intelligence and Agent Systems 9, 363–375 (2011)
15. Centeno, R., Billhardt, H., Hermoso, R.: Persuading agents to act in the right way: An incentive-based approach. Engineering Applications of Artificial Intelligence 26(1), 198–210 (2013)
16. Centeno, R., Billhardt, H., Hermoso, R., Ossowski, S.: Organising MAS: a formal model based on organisational mechanisms. In: Proceedings of the 2009 ACM symposium on Applied Computing, pp. 740–746 (2009)
17. Ossowski, S. (ed.): Agreement Technologies. Low, Governance and Technologies Series, vol. 8. Springer (2013)
18. Martin, D., et al.: OWL-S: Semantic Markup for Web Services, W3C Member Submission (2004), `http://www.w3.org/Submission/OWL-S/`
19. Roman, D., Lausen, H., Keller, U. (eds.): Web Service Modeling Ontology (WSMO). Technical report, WSMO Final Draft (2005), `http://www.wsmo.org/TR/d2/v1.2/`
20. Kopecky, J., Vitvar, T., Bournez, C., Farrell, J.: SAWSDL: Semantic Annotations for WSDL and XML Schema. IEEE Internet Computing 11(6) (2007)
21. Sheth, A., Gomadam, K., Lathem, J.: SA-REST: Semantically Interoperable and Easier-to-Use Services and Mashups. IEEE Internet Computing 11(6), 91–94 (2007)
22. Kopecky, J., Vitvar, Zaremba, M., Fensel, D.: WSMO-Lite: Lightweight Semantic Descriptions for Services on the Web Full. In: Proceedings of 5th ECOWS, pp. 77–86 (2007)
23. Huhns, M., Singh, M., Burstein, M., Decker, K., Durfee, E., Finin, T., Gasser, L., Goradia, H., Jennings, N.R., Lakartaju, K., Nakashima, H., Parunak, V., Rosenschein, J., Ruvinsky, A., Sukthankar, G., Swarup, S., Sycara, K., Tambe, M., Wagner, T., Zavala, L.: Research directions for service-oriented multiagent systems. IEEE Internet Computing 9(6), 69–70 (2005)
24. Fernández, A., Ossowski, S.: A multiagent approach to the dynamic enactment of semantic transportation services. IEEE Transactions on Intelligent Transportation Systems 12(2), 333–342 (2011)
25. Fernández, A., Hayes, C., Loutas, N., Peristeras, V., Polleres, A., Tarabanis, K.: Closing the Service Discovery Gap by Collaborative Tagging and Clustering Techniques. In: 2nd International Joint Workshop on Service Matchmaking and Resource Retrieval in the Semantic Web (SMR2), Karlsruhe, Alemania (2008)
26. Cong, Z., Fernández, A.: Enabling Web Service Discovery in Heterogeneous Environments. International Journal of Metadata, Semantics and Ontologies (in press)
27. Ehrig, M.: Ontology Alignment: Bridging the Semantic Gap. Springer, Heidelberg (2007)
28. Euzenat, J., Shvaiko, P.: Ontology matching. Springer, Heidelberg (2007)

29. Augusto, J.C.: Ambient Intelligence: the Confluence of Ubiquitous/Pervasive Computing and Artificial Intelligence. In: Schuster, A. (ed.) Intelligent Computing Everywhere. Springer (2007)
30. IST Advisory Group. The European Union report: Scenarios for Ambient Intelligence in 2010 (2001), ftp://ftp.cordis.europa.eu/pub/ist/docs/istagscenarios2010.pdf
31. Weiser, M.: The Computer for the 21st Century. Scientific American 265, 94–104 (1991)
32. Satyanarayanan, M.: Pervasive Computing: Vision and Challenges. IEEE Personal Communications 8(4), 10–17 (2001)
33. Ramos, C., Augusto, J.C., Shapiro, D.: Ambient Intelligence—the Next Step for Artificial Intelligence. IEEE Intelligent Systems 23, 15–18 (2008)
34. Want, R., Hopper, A., Falcao, V., Gibbons, J.: The Active Badge Location System. ACM Transactions on Information Systems 10(1), 91–102 (1992)
35. Skov, M., Hoegh, R.: Supporting information access in a hospital ward by a context-aware mobile electronic patient record. Journal of Perv. and Ubiq. Computing 10, 205–214 (2006)
36. Park, D., Hwang, S., Kim, A., Chang, B.: A Context-Aware Smart Tourist Guide Application for an Old Palace. In: Proc. of the 3rd International Conference on Convergence Information Technology, pp. 89–94 (2007)
37. Gu, T., Pung, H.K., Zhang, D.Q.: Towards an osgibased infrastructure for context-aware applications in smart homes. IEEE Pervasive Computing, 66–74 (2004)
38. Holvoet, T., Valckenaers, P.: Beliefs, desires and intentions through the environment. In: Proceedings of AAMAS 2006, pp. 1052–1054 (2006)
39. Schilit, B., Theimer, M.: Disseminating active map information to mobile hosts. IEEE Network 8(5), 22–32 (1994)
40. Pynadath, D.V., Scerri, P., Tambe, M.: Towards Adjustable Autonomy for the Real World. Journal of Artificial Intelligence Research 17, 171–228 (2002)
41. Dorais, G., Bonasso, R.P., Kortenkamp, D., Pell, B., Schreckenghost, D.: Adjustable autonomy for human-centered autonomous systems. Working notes of the Sixteenth International Joint Conference on Artificial Intelligence, Workshop on Adjustable Autonomy Systems, pp. 16–35 (1999)
42. Atencia, M., Schorlemmer, M.: An interaction-based approach to semantic alignment. Web Semantics: Science, Services and Agents on the World Wide Web 12-13, 131–147 (2012)
43. Tordai, A.: On Combining Alignment Techniques. Editorial Vrije Universiteit, Amsterdam (2012)
44. Koster, A., Schorlemmer, M., Sabater-Mir, J.: Engineering trust alignment: Theory, method and experimentation. Journal of Human-Computer Studies 70(6), 450–473 (2012)
45. Koster, A., Sabater-Mir, J., Schorlemmer, M.: Trust Alignment: A Sine Qua Non of Open Multi-agent Systems. In: Meersman, R., et al. (eds.) OTM 2011, Part I. LNCS, vol. 7044, pp. 182–199. Springer, Heidelberg (2011)
46. Abras, S., Ploix, S., Pesty, S., Jacomino, M.: A multi-agent home automation system for power management. In: Cetto, J.A., Ferrier, J.-L., Pereira, J.D., Filipe, J. (eds.) Informatics in Control Automation and Robotics. LNEE, vol. 15, pp. 59–68. Springer, Heidelberg (2008)