A Methodology for Developing Multiagent Systems as 3D Electronic Institutions

Anton Bogdanovych[1], Marc Esteva[1], Simeon Simoff[1], Carles Sierra[2], and Helmut Berger[3]

[1] Faculty of IT, University of Technology Sydney, Australia
{anton,esteva,simeon}@it.uts.edu.au
[2] Artificial Intelligence Research Institute (IIIA), CSIC, Campus UAB, Spain
sierra@iiia.csic.es
[3] ECommerce Competence Center, Vienna, Austria
helmut.berger@ec3.at

Abstract. In this paper we propose viewing Virtual Worlds as open Multiagent Systems and propose the 3D Electronic Institutions methodology for their development. 3D Electronic Institutions are Virtual Worlds with normative regulation of interactions. More precisely, the methodology we propose here helps in separating the development of Virtual Worlds based on the concept of 3D Electronic Institutions into two independent phases: specification of the institutional rules and design of the 3D interaction environment. The new methodology is supplied with a set of graphical tools that support the development process on every level, from specification to deployment. The resulting system facilitates the direct integration of humans into Multi-Agent Systems as they participate by driving an avatar in the generated 3D environment and interacting with other humans or software agents, while the institution ensures the validity of their interactions.

1 Introduction

The field of Multiagent Systems (MAS) focuses on the design and development of systems composed of autonomous entities which act in order to achieve their common or individual goals. Several methodologies based on the MAS paradigm have been proposed in the recent years (see [1,2,3] for reviews). Although humans can be seen as autonomous entities most of the MAS methodologies do not consider direct human participation. In general, human role is limited to acting behind the scenes by customising templates of the agents that participate in the system on humans' behalf. Moreover, existing MAS methodologies that consider direct human participation have not developed the necessary tools to facilitate human inclusion.

One of the few areas where direct human participation is considered is the domain of open systems [4], which with the expansion of Internet have been identified as the most important area of application of MAS [5]. Those are systems where participants are assumed to be heterogeneous and self interested

M. Luck and L. Padgham (Eds.): AOSE 2007, LNCS 4951, pp. 103–117, 2008.

and cooperative behaviour can not be expected from them. Hence, methodologies for open systems should not commit to a particular agent architecture or programming language, and should provide mechanisms to deal with agents with self-interested behaviours.

Two of the most prominent methodologies for open systems based on the MAS paradigm are Gaia [6] and Electronic Institutions [7]. In Gaia the system is designed as a set of organizations where agents participate playing different roles. However, Gaia methodology only covers the specification of the system and does not offer any technological support in regards to system execution. In Electronic Institutions the design of the system focuses on specifying a set of institutional rules which establish possible behaviour of the agents. The Electronic Institutions methodology covers all the steps from the specification, to the deployment and execution of the system. Furthermore, the steps of the methodology are supported by a set of provided software tools.

An Electronic Institution can be regarded as a mediator between participants that verifies the validity of their interactions against the set of rules, protocols and norms specified by the systems designers. No assumptions are made about the internal architecture of participating agents and it is only required for an agent to be able to connect to the institution and communicate with it. Thus, Electronic Institutions form a perfect playground for the development of human centered Multiagent Systems and open new horizons to the research in human-agent collaboration. Despite this fact participation of humans in Electronic Institutions have not been well studied and the facilities for their integration have not been properly developed.

In order to solve this problem we propose using 3D Virtual Worlds, which is one of a very few technologies that provides all the necessary means for direct human inclusion into software systems. 3D Virtual Worlds are software generated environments which follow the metaphor of architecture and emulate real world using 3-dimensional visualisation. Humans participate in those environments represented as graphical embodied characters (avatars) and can operate there using simple and intuitive control facilities, which are more or less similar throughout the whole variety of the different Virtual Worlds present on the market. We advocate that 3D Virtual Worlds technology can be successfully used for “opening” Multiagent societies to humans.

In this paper we present a methodology for 3D Electronic Institutions, a concept that appeared from the combination of Electronic Institutions and 3D Virtual Worlds. This methodology focuses on the development of normative environments inhabited by software and human agents. At this aim, the methodology extends the Electronic Institution methodology to generate a representation of the Electronic Institution in 3D Virtual Worlds and to define the necessary elements to successfully integrate both technologies. Hence, humans participate in the institution by controlling an avatar on the generated Virtual World. The methodology is supplied with a set of software tools which give support to all the stages of its development.

Apart from opening MAS to humans the research in 3D Electronic Institutions can also benefit the Virtual World community, which is looking for mechanisms to incorporate social rules into Virtual Worlds in order to control and structure participants' interactions. The design and development of Virtual Worlds has emerged as a phenomenon shaped by a home computer user rather than by research and development activities at universities or companies. As a result, Virtual Worlds are more or less unregulated environments and continue to be developed on an ad hoc basis. Despite the fact that active support of human interactions is one of the key characteristics that set Virtual Worlds apart from other technologies, there are no flexible facilities to control these interactions. As the number of inhabitants of the artificial societies established in Virtual Worlds grows, the level of immersion increases and participants become more and more involved with the experience, the need for structuring their interactions becomes more explicit. Lacking clear mechanisms for doing it, users try introducing some of the convenient social rules from the real world. Doing so, however, in a system that was built without a methodology centered on structuring users' interactions is a very challenging task. One of the consequences of this is that Virtual Worlds are mostly used in computer games, where structuring the interactions of participants is not necessarily useful and the consequences of errors in the code are not dramatic. In order to extend the scope of Virtual Worlds technology to be applied to a wider range of problems, exploit the benefits brought by the Virtual Worlds and deal with their growing complexity, methodologies that regulate the interactions of participants and improve the reliability and security issues need to be applied. We believe that Virtual Worlds have much greater potential and can be used for a broader spectrum of problems. New economical circumstances and conceptual similarity with open systems create a need for Virtual Worlds to be used in domains like E-Commerce, online travel etc. The aforementioned problems of the Virtual Worlds can clearly be solved by applying the 3D Electronic Institutions methodology to their development.

The remainder of the paper is structured as follows. In section 2 we present the conceptual model behind the 3D Electronic Institutions metaphor. Section 3 outlines the steps that 3D Electronic Institutions Methodology utilization requires to be followed and gives the detailed overview of the technical aspects and tools supplied with 3D Electronic Institutions. Next, in section 4 we describe the deployment architecture, while in section 5 we summarize the contribution and present some concluding remarks.

2 3D Electronic Institutions

Conceptually speaking, *3D Electronic Institutions are Virtual Worlds with normative regulation of interactions.* More precisely, we propose to separate the development of 3D Electronic Institutions into two independent phases: specification of the institutional rules and design of the 3D Interaction environment. Such separation is widely used in architecture [8], whose metaphor inspires Virtual Worlds. We are convinced that having it in Virtual Worlds would also be highly beneficial.

For the purpose of the rule specification we suggest employing the Electronic Institutions methodology [9], which is able to ensure the validity of the specified rules and their correct execution. In contrast to Electronic Institutions and Gaia, the normative part of a 3D Electronic Institution does not represent all the activities that are allowed to be performed in a Virtual World. The normative part can be seen as defining which actions require institutional verification assuming that any other action is allowed. Not every Virtual World requires such an approach as well as not every institution requires 3D visualization. Only systems that have a high degree of interactions and those interactions need to be structured in order to avoid violations may need institutional intervention. And only the institutions where 3D visualization of active components is possible and beneficial should be visualized in Virtual Worlds.

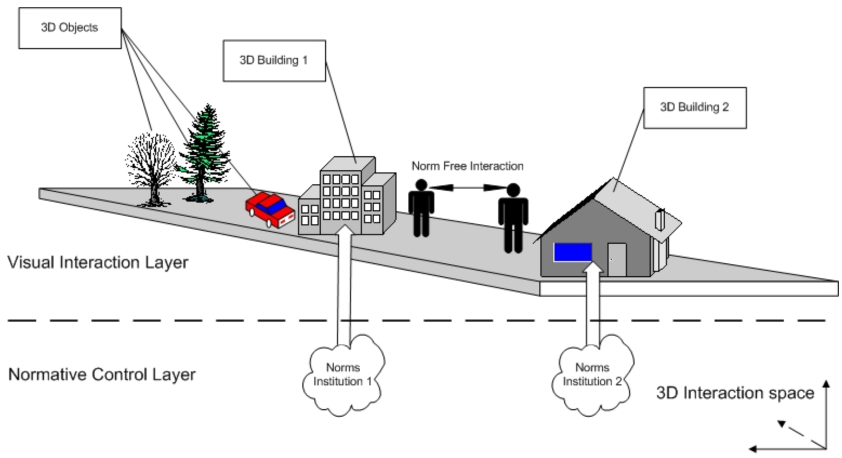

Fig. 1. 3D Electronic Institutions Concept

For those systems that can benefit from both 3-dimensional visualization and institutional control of the specified rules we suggest using the following *conceptual model*. A 3D Electronic Institution is visualized in terms of a 3D Virtual World. We call this Virtual World a 3D Interaction Space. Inside the 3D Interaction Space an institution is represented as a building, and participants are visualized as avatars. Once they enter the building their actions are validated against the specified institutional rules. The institutional buildings is divided into a set of rooms, which are separated from each other by doors. The doors are open or closed for a participant depending on the role she/he is playing, the institutional rules and the current execution state. Figure 1 outlines the brief idea behind the 3D Electronic Institutions concept presented so far. Next we describe the components of the conceptual model in more details.

3D Interaction Space. It represents the generated 3D Virtual World, and there is no possibility for participants to move beyond it. The only way to leave it is by disconnecting from the Virtual World. Once someone enters it, he/she will become embodied as an avatar and will be physically located inside. To enhance the believability of the visualization the space is usually populated with a number of various 3D Objects. The most typical case is that a 3D Interaction Space is decorated as a garden, where the objects enhancing the believability are trees, bushes, cars etc. A special type of objects within the space are the buildings. Each of the buildings metaphorically represents an institution. Anywhere outside the institutional building interactions among participating avatars are not regulated and every event that happens inside this space is immediately visualized without any prior validation.

Institutional Building. An institution is represented as a building in the 3D Interaction Space, and the interactions within the building are regulated by the specified institutional rules. Every event that a participant requests by pressing keys on the keyboard or operating with the mouse are first sent to the institutional infrastructure for validation. If the institution permits event's execution – the corresponding action is visualized, otherwise the event is ignored. It is also possible for the institution to provide context based explanations of the reasons why a particular event can not be processed. The institutional building has a single entrance door, through which the participants can enter it.

Avatars. The participants of the 3D Interaction Space are visualized as avatars. We distinguish between the following two types of avatars: avatars for users and avatars for the institutional employees. For the users' avatars an initial set of default appearances is provided, but those appearances can be changed later. The institutional employees are usually represented by autonomous agents that play internal roles in the corresponding Electronic Institution. They are assumed to have similar appearance which goes inline with the dress code of the institution they are employed with. While outside the institutions the avatars are free to execute any possible actions and their communication is not moderated by any of the institutions. Once they enter an institutional building they can only execute the actions that are permitted by the corresponding institution. In some of the rooms it is allowed by the institution to split the user into several alteroids (avatars), to participate simultaneously in different activities. Each time a new alteroid is created the user should decide which to choose to control and a new autonomous agent is executed to take control over the other ones. This functionality allows a user to employ autonomous agents for performing some routine tasks on user's behalf, while the user may be involved into some other activities.

Rooms. Every institutional building consists of a set of rooms, each one representing a different activity. The number of rooms within a building and the activity going on in each one is determined by the institution specification. Rooms are represented as a set of rectangular boxes closed by walls from every side. Agents can enter and leave a room by traversing one of the doors embedded in their walls connecting it with other rooms.

Doors. The Doors are used to connect different rooms in the institutional building. The institutional rules and the execution state determine which agents depending on their role can progress through the door. This is strictly controlled by the Electronic Institution.

Map. In order to simplify the navigation of the users, every institution is supplied by the map of the building. The map usually appears in the upper-right corner of the screen as a semitransparent schematic plan. Each of the available rooms is shown on the map and the human-like figures show every user the positions of all the alteroids a user is associated with. While moving through the institution the positions are updated accordingly.

Backpack with obligations. While acting in an institution a user may acquire some commitments. An example of such a commitment may be that a user who just won the auction will not be able to directly leave the institution, but is committed to visit the payment room before leaving. These commitments are expressed in the specification of the underlying Electronic Institution and are fully controlled by the system. In order to have a simple way to present those commitments to a user we use the metaphor of a backpack used in many computer games. The backpack is usually present in the lower right part of the screen and a user may decide to hide it or show it back after hiding. Clicking on the backpack will result in a user being presented with the textual list of the acquired commitments.

Events/Actions/Messages. Although, we anticipate that the users may use all sorts of different devices for navigating virtual worlds, in a standard case a participant of a 3D Interaction Space is able to control the avatar and change the state of the Virtual World by pressing keyboard buttons, moving a mouse or clicking mouse buttons. Those physical actions executed by a user in the real world generate events inside the Virtual World, which are then visualized as actions executed within the 3D Interaction Space. The events that a user is trying to execute inside an institutional building are not directly visualized. Before visualization every event is transformed into a message understandable by the institution and send to the institutional infrastructure for validation. Only if the given message is consistent with the current state of the institution and the institutional rules, the action is performed and visualised.

3 3D Electronic Institutions Methodology

In the previous section we presented the metaphor of 3D Electronic Institutions. Here we describe the methodology that facilitates their development and show how this new methodology embeds the Electronic Institutions methodology. We want to remark that this methodology covers the development of a single institution. In order to have an Interaction Space populated by several institutions, the methodology has to be applied to each one of them.

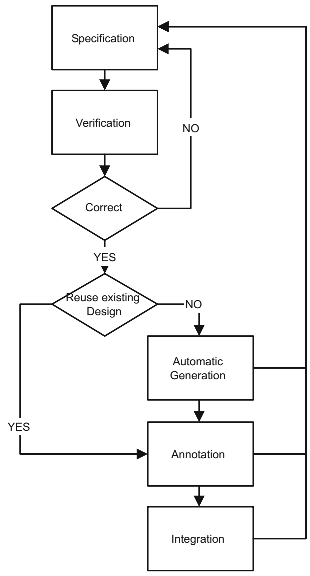

Fig. 2. Methodology steps

Applying 3D Electronic Institutions methodology requires 5 steps to be accomplished:

1. Specification of an Electronic Institution using ISLANDER [7].
2. Verification of the specification.
3. Automatic Generation of the corresponding 3D environment (if needed).
4. Annotation of the Electronic Institution specification with components of the 3D Virtual World.
5. Integrating the 3D Virtual World into the institutional infrastructure.

Figure 2 presents the overview of each of the steps and their sequence. The detailed explanation of each of them follows next.

Step 1 – Specification. The specification step is the same as in the Electronic Institutions methodology [7]. It establishes the regulations that govern the behaviour of the participants. This process is supported by ISLANDER which permits to specify most of the components graphically, hiding the details of the formal specification language and making the specification task transparent. The institutional regulations are established by three types of conventions.

Conventions on language, the *Dialogical Framework*. It defines a common ontology and communication language to allow humans with different cultural backgrounds, as well as, agents to exchange knowledge. This ontology and language for humans will be further transformed into actions that are allowed to be executed in the Virtual World. Those actions are connected to 3D models in the environment, the affordances of which will help in eliminating the cultural barrier. Due to the further provided translation of the communication language

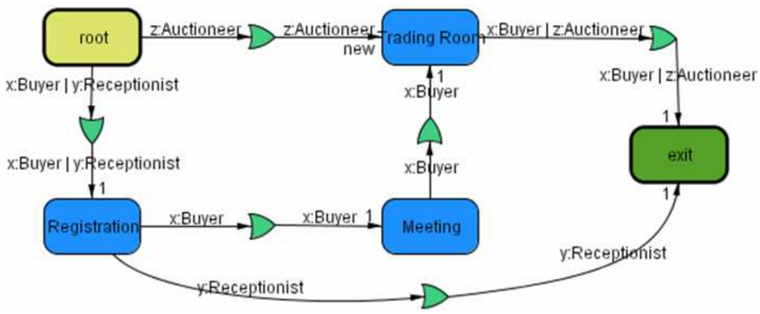

Fig. 3. Trading Institution Performative Structure

into actions and vice-versa, the agents will be able to interact with humans and understand their actions. The dialogical framework also fixes the organizational structure of the society, that is, which roles can participants play.

Conventions on activities, the *Performative Structure*. This dimension determines in which types of dialogues users can engage. For each different activity, interactions among participants are articulated through group meetings, called scenes, which follow well-defined interaction protocols. The protocol of each scene is specified by a graph where the states represent the different interaction states and arcs are labeled with messages of the communication language or timeouts. Participants in a scene can change over time and at this aim, a set of access and exit states per role are defined. Finally, role populations are specified by establishing the minimum and maximum number of participants that must or can play each role. More complex activities are specified by establishing connections among scenes. The resulting, network of scenes, the *Performative Structure*, defines how agents can legally move among scenes depending on their role. This transit of participants between scenes is regulated by special (simple) scenes called transitions, which allow expressing synchronization, parallelization and choice points. In their transit among scenes users are allowed to change their role. The Performative Structure contains two special scenes, the initial and final scene, which does not model any activity and must be regarded as the institution entrance and exit. Participants entering the institution are initially placed in the initial scene, while reaching the final scene means leaving the institution.

Conventions on behavior, the *Norms*. Norms determine the consequences of user actions. These consequences are modeled as commitments that participants acquire as a consequence of their actions and have to fulfill later on. These commitments may restrict future activities of the users.

In order to illustrate the different steps of the methodology, we use a very simple Trading Institution. This institution can be enacted by the agents playing the *receptionist*, *auctioneer* and *buyer* roles. Figure 3 shows the performative structure of the Trading Institution, where rectangles represent different scenes and triangular shapes are transitions. Apart of the root and exit scenes which just represent the entrance and exit, the institution contains the *Registration*, *Meeting* and *Trading Room* scenes. In the *Registration* scene buyers register by communicating their login and password to an agent playing the receptionist role. In the *Meeting* scene buyers can meet and freely interact, while in the *Trading Room* buyers can acquire products auctioned by an agent playing the auctioneer role. The arcs connecting scenes and transitions are labeled with the roles that can progress through them. Notice that buyer agents are required to go to the Registration scene before moving to the meeting room to interact with other buyers. From the Meeting scene, they can proceed to the Trading Room to participate in the auctions. Receptionist can only go to the Registration scene, while agents playing the auctioneer role can only access the Auction Room from the root scene.

Step 2 – Verification. One of the advantages of the formal nature of the 3D Electronic Institutions methodology is that the specification produced on the previous step can be automatically verified for correctness by ISLANDER. The tool verifies the scene protocols, the role flow among the different scenes and the correctness of norms. This verification starts with the validation of the correctness of the protocol defined by each scene. This includes checking that for each state there is a path from the initial state to a final state that passes through the current state, and that the messages of the arcs are correct with respect to the communication language. At the Performative Structure level it is verified that agents will not get blocked at any scene or transition. Thus, it is checked that from each scene and transition users have always a path to follow, that each of them is reachable from the initial scene and that from each scene and transition exists a path to the final scene that will allow participants to leave the institution. Finally, ISLANDER checks that norms are correctly specified and that participants can fulfil their commitments. As commitments are expressed as actions that users have to carry out in the future, it is verified that those actions can be performed by agents.

The verification permits to detect errors in the specification before starting the design and development of the 3D visualization. If such errors are found, the developers should go back to step 1 to correct them. If the specification contains no errors, there are two options. If the 3D Visualization of the environment is already created (reuse of the existing design) then the developers may skip the next step and continue with Step 4. Otherwise, the generation step, Step 3, should be executed.

Step 3 – Generation of the visualization. The institutional specification does not only define the rules of the interactions, but also helps to understand which visualization facilities are required for participants to operate in the institution. Most elements of the specification have conceptual similarities with basic concepts of 3D Virtual Worlds, which makes it possible to create an automatic mapping between those. In our metaphor scenes and transitions, correspond to rooms, connections (arcs) in the Performative Structure graph become doors, and the number of participants allowed in a scene determines the size of a room. The performative structure corresponds to the map of the institution and the

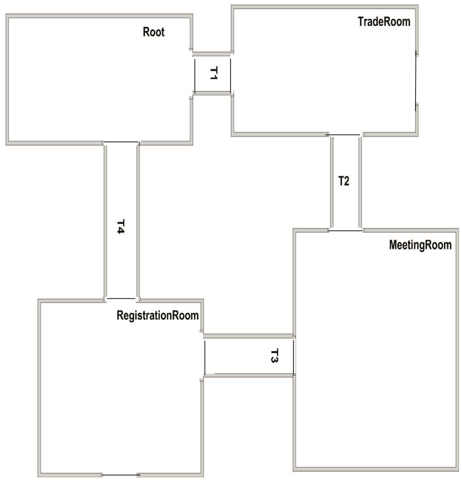

Fig. 4. Trading Institution floor plan

backpack with obligations is a visual way to communicate the normative obligations to the users.

Having this mapping serves two conceptual purposes: explaining the Electronic Institutions metaphor in terms of Virtual Worlds using the concepts familiar to most of the humans and explaining the Virtual Worlds in terms of the institutional specification of the underlying processes. Practically it helps to generate a part of the visualization in a fully automatic way (see [10] for details). The generation can function in two different modes: Euclidean and non-Euclidean. In the first case the rooms on the generated floor plan are positioned so, that each scene and transition connected in the Performative Structure are physically placed next to each other and there is a door between them. In the non-Euclidean case the rooms may be located anywhere and are not necessarily involved in any sort of spatial relationship. The movement between connected rooms in the non-Euclidean approach will then be conducted using teleportation[1].

Next, all the rooms are resized to be able to include the maximum number of participants allowed in the corresponding scene. Another outcome of this step is the schematic plan (map) of the institution.

Figure 4 depicts the automatically generated floor plan for the Trading Institution. Notice that there is a room for each scene and transition of the performative structure shown in Figure 3, except for the output scene (the output scene does not model any interaction and it only represents the exit point of the institution). This is expressed in the figure by the doors (represented by thin lines) in the registration and trading room that are not connected to any other room. Once an agent traverses one of these doors – the agent leaves the institutional building and appears inside the uncontrolled part of the 3D Interactions

[1] The process of moving objects from one place to another instantaneously, without passing through the intervening space.

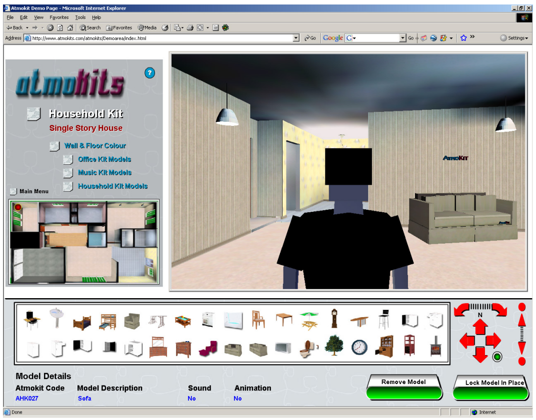

Fig. 5. Annotating the rooms with Atmokits

Space. Transitions are the rooms with names Ti, where i ranges from 1 to 4. The connections among rooms are established through doors, positions of which are determined by the arcs connecting scenes and transitions in the performative structure of the Trading Institution. Although two rooms can be connected by a door, only the agents playing the roles that label the corresponding arc will be allowed to progress through the corresponding door.

Step 4 – Annotation. Although a part of the visualization of a 3D Electronic Institution can be automatically generated there is usually not enough information present in the specification to produce an appealing visualization. To enrich the generated visualization we use the Annotation Editor tool. This tool helps to change the textures, colors and add additional objects inside each of the rooms. In the current implementation we use the Atmokits software[2] for this purpose. It is supplied with a set of standard objects and textures that can be used to enrich the design of the rooms. Figure 5 shows the interface of Atmokits. Left side of the figure shows the map of the institution and the right part displays the 3-dimensional representation of one of the rooms with an avatar inside. The bottom part outlines a set of objects that can be inserted and the control buttons that are used for precise object positioning.

After the annotation step the user can return to step 1, if for any reason he/she wants to modify the specification, or move to step 5.

Step 5 – Integration. On the integration step the execution state related components are specified. This includes the creation of the set of scripts that control the modification of the states of the 3D Virtual Worlds and mapping of those scripts to the messages, which change the state of the Electronic Institution. Firstly, the scripts that correspond to the messages from the agent/institution protocol need to be defined. These include entering scene, leaving scene, entering transition, leaving transition etc. Next, the scripts that correspond to the

[2] http://www.atmokits.com

specific messages that are defined in the ontology on the specification step must be created. If there is a need to map the data types in the ontology to 3D objects in the Virtual World it should also be done on this step. At the end, the correspondences between the messages and scripts (actions) are created by filling in the Action/Message table. The Action-Message table for the trading institution is presented in Table 1. The table specifies, for example, that when an avatar collides with a door this action is mapped to an *ExitScene* message, while the action of raising a hand is mapped to a *Bid* message. Furthermore, the table is also used to map the messages of the institutional infrastructure to actions in the Virtual World that change the visualization. For instance, if a *EnteredAgentInstitution* message is received a new avatar will be shown in the initial scene of the institution.

Making the integration a separate step of the methodology stimulates the development of the scripts in the form of design patters, that are generic enough to be reused in other systems.

After accomplishing this step the generated 3D Virtual World is ready to be visualized and the 3D Electronic Institution infrastructure will be executed to take care of the validity of interactions between participants, verify the permissions of participants to access different scenes and will make sure that all the institutional norms and obligations are imposed.

Table 1. Action-Message Table

Action	Message
addNewAvatar	EnteredAgentInstitution
doorCollideFromScene	ExitScene
roomEnter	MoveToScenes
doorCollideFromTransition	ExitTransition
transitionEnter	MoveToTransition
raiseHand	Bid
removeAvatar	ExitedAgentInstitution

4 Deployment

For deployment of 3D Electronic Institutions created by the proposed methodology, we use a 3-layered infrastructure presented in Figure 6.

First layer is the *Electronic Institution Layer*. It uses the AMELI system [7] for enforcing the institutional rules established on the specification step. AMELI keeps the execution state of the institution and uses it along with the specification to guarantee that participants' actions do not violate any of the institutional constraints.

Second layer is the *Communication Layer*. Its task is to causally connect [11] the institutional infrastructure with the visualization system transforming the actions of the visualization system into the messages, understandable by the institutional infrastructure and the other way around. This causal connection

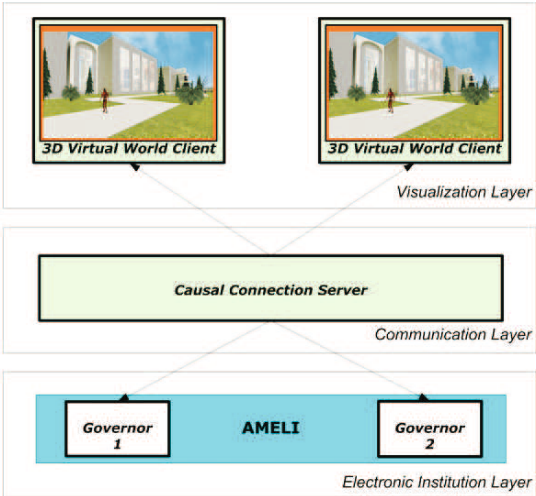

Fig. 6. Runtime Architecture

is done via the Causal Connection Server, which uses the Action-Message table created on the integration step. The causal connection is happening in the following way: an action executed in the 3D Virtual World (that requires institutional verification) results in a change of the institutional state in the AMELI layer, as well as every change of the institutional state is reflected onto the 3D Virtual World and changes its state. The Communication layer conceptually and technologically connects two metaphors: Electronic Institutions and Virtual Worlds and we see it as one of our major scientific contributions.

The third layer is called *Visualization Layer*. It is used to visualize the 3D Virtual World for the users. Currently, we are employing the Adobe Atmosphere[3] technology for this task, however, due to the fact that it was discontinued we are making a transition to Second Life[4]

A clear separation of the runtime architecture into three different layers has a number of advantages:

1. The interactions inside the 3D Virtual World become structured, secure and predictable, as everything that needs control is verified by AMELI and will happen as specified.
2. The Visualization Layer can be easily replaced (i.e. when a more advanced visualization platform appears on the market) with minimal changes in the rest of the system.
3. The changes in the Electronic Institution Layer will be automatically reflected onto the Visualization layer or will require minimal manual adjustment.

[3] http://adobe.com/products/atmosphere
[4] http://secondlife.com

4. A number of different visualization platforms (possibly implemented via different technologies) can be simultaneously connected to the Causal Connection Server and share the same institution.
5. Some participants (i.e. software agents) can bypass the 3D Virtual World and directly connect to the institution via the Electronic Institutional layer, while other participants (humans) will be able to observe their presence and actions in the 3D Virtual World.

5 Conclusion

In this paper we presented the 3D Electronic Institutions methodology, which supports human integration into MAS-mediated environments and provides all the necessary technological support for them to actively participate and interact with other humans or autonomous agents. This methodology is supplied with a set of tools that facilitate the design, development and execution of such environments. We would like to stress that, to our knowledge, 3D Electronic Institutions is the first methodology that is specifically concerned with the developments of Virtual Worlds with normative regulations of interactions. Its formal nature has a number of advantages. Firstly, it forces the designer to follow a structured and formal approach, having to analyse the system in detail before implementing it. This permits designers to detect the critical points and possible problems at an early stage. Furthermore, the methodology clearly distinguishes between the design of the institutional rules and the design of its visualization in Virtual Worlds, which proved to be an efficient way to develop real world designs. Another advantage of using this methodology is that the supplied tools make the development faster, helping to achieve some tasks automatically. Moreover, due to the distributed architecture possible updates of the system can be accommodated in an easy way. Notice, that the development process is independent of the particular Virtual Worlds technology used for the visualization of the system. This in combination with the execution infrastructure permits a quick and easy portability of the system to new visualization platforms.

The proposed architecture also supports an efficient collaboration between humans and agents. There is always a software agent assigned to every human participating in the institution and either of them can control the avatar. When the human is driving the avatar the agent observes and records the actions of the principal. This information is used later on, when the agent is in control of the avatar for achieving its goals and expressing believable human-like behaviour. The immersive nature of 3D Virtual Worlds creates better possibilities to observe human behaviour without a need to overcome the embodiment dissimilarities, while institution control of the interactions helps the agent to reduce the number of possible behaviours and hence, to learn faster. Furthermore, when the agent is driving the avatar the human is supplied with convenient interface to observe its actions and intervene when necessary. In this way the behaviour of the agent acting on user's behalf can be easily controlled, increasing the trust and confidence of the humans in the agent.

Acknowledgements

The research reported in this paper is partially supported by the ARC Discovery Project DP0451692 "The Evolution of Business Networks in Virtual Marketplaces" and the Spanish projects "Autonomic Electronic Institutions" (TIN2006-15662-C02-01) and "Agreement Technologies" (CONSOLIDER CSD2007-0022, INGENIO 2010).

References

1. Iglesias, C., Garijo, M., González, J.: A Survey of Agent-Oriented Methodologies. In: Rao, A.S., Singh, M.P., Müller, J.P. (eds.) ATAL 1998. LNCS (LNAI), vol. 1555, pp. 317–330. Springer, Heidelberg (1999)
2. Wooldridge, M., Ciancarini, P.: Agent-Oriented Software Engineering: The State of the Art. In: Ciancarini, P., Wooldridge, M.J. (eds.) AOSE 2000. LNCS, vol. 1957, pp. 55–82. Springer, Heidelberg (2001)
3. Gómez-Sanz, J., Pavón, J.: Methodologies for developing multi-agent systems. Journal of Universal Computer Science 10(4), 359–374 (2004)
4. Hewitt, C.: Offices are open systems. ACM Transactions on Office Information Systems 4(3), 271–287 (1986)
5. Wooldridge, M., Jennings, N.R., Kinny, D.: A methodology for agent-oriented analysis and design. In: Proceedings of the third annual conference on Autonomous Agents (AGENTS 1999), pp. 69–76. ACM Press, New York (1999)
6. Zambonelli, F., Jennings, N., Wooldridge, M.: Developing multiagent systems: The gaia methodology. ACM Transactions on Software Engineering Methodology 12(3), 317–370 (2003)
7. Arcos, J.L., Esteva, M., Noriega, P., Rodriguez-Aguilar, J.A., Sierra, C.: An Integrated Developing Environment for Electronic Institutions. In: Agent Related Platforms, Frameworks, Systems, Applications, and Tools. Whitestein Book Series, Springer, Heidelberg (2005)
8. Maher, M., Simoff, S., Mitchell, J.: Formalizing building requirements using an activity/space model. Automation in Construction 6, 77–95 (1997)
9. Esteva, M.: Electronic Institutions: From Specification to Development. PhD thesis, Institut d'Investigació en Intel.ligència Artificial (IIIA), Spain (2003)
10. Bogdanovych, A., Drago, S.: Euclidean Representation of 3D electronic institutions: Automatic Generation. In: Proceedings of the 8th International Working Conference on Advanced Visual Interfaces (AVI 2006), pp. 449–452 (2006)
11. Maes, P., Nardi, D.: Meta-Level Architectures and Reflection. Elsevier Science Inc., New York (1988)