

TOWARDS AGENT-BASED PRODUCT MODELLING

John S. Gero and Udo Kannengiesser

Key Centre of Design Computing and Cognition, University of Sydney

Abstract: This chapter presents an approach to product modelling, which aims to provide interoperability when a common data model is not available. It uses computational agents that represent product components. They autonomously organise themselves in a situated manner to form a society that represents the product.

Key words: product modelling, situated agents, self-organisation

1. INTRODUCTION

There has been an increased use of computational tools to support various tasks in product development. Examples include computer-aided drafting (CAD) and manufacturing (CAM) systems and a number of specialised tools for analyses such as finite element analysis (FEA) and spreadsheet analysis. These processes require fundamentally different data about the product and different representations of that data, even though the data is concerned with the same product. Most computational systems supporting these processes have been developed independently from one another to address the specific needs of each process and use different product data representations. However, industrial product development is a process that involves a complex network of interrelated activities, each of which needs information produced or manipulated by the other. Interoperability – the ability to move data from one representation of a product to another to allow other computational processes to operate on it – has become an area of growing concern as the cost of such interchanges increases (NIST 1999).

Product modelling (Eastman 1999) includes the notion that describes conceptual approaches to the issue of exchanging product data representations within and between companies. These approaches are commonly founded on a standard data model that is used to translate between the different native formats of the applications. Any object that needs to be made interoperable must be pre-defined in this model and encoded into a standard form. One of the best-known product models is the ISO 10303 standard, informally known as STEP (STandard for the Exchange of Product model data), developed by the International Standards Organization (ISO). The STEP product model consists of a number of partial models called Application Protocols (APs) describing classes of objects for specialised domains. All STEP APs are represented using the EXPRESS language (ISO 1994). Another well-known product modelling example is the International Alliance for Interoperability (IAI), an industry-led organization whose members have agreed upon a pre-defined product data schema called Industry Foundation Classes (IFCs) (IAI 2000), which are also represented in the EXPRESS format.

Although both STEP and IFCs have seen an increasing use in industry, there are still some unresolved practical issues. Many applications have specialised modelling capabilities that are not supported by existing STEP APs or IFCs. One of the reasons is that the development and extension of these standard models is a relatively slow procedure involving international consensus-seeking, which lags behind the advances in modelling technologies implemented in industry. In addition, many translators have been built to exchange specialised data between specialised tools. As a consequence, they are tailored to implement only a subset of the neutral model that would be capable of capturing the whole tool data. When data transfer is required between new combinations of the translators/tools, their data representations in the neutral model often mismatch leading to data loss. Other tools do not provide STEP or IFC translators at all.

Our research aims to lay the foundations of a framework to establish interoperability when a standard product representation such as provided by STEP or IFCs is not available. This framework is based on a pool of agents representing potential components that organise themselves in an autonomous fashion to represent the product in response to the current situation. This allows them to flexibly adapt their data representations to every particular design tool.

2. FOUNDATIONS

This research is founded on concepts drawn from recent developments in situated design agents and on work on multi-agent societies.

2.1 Situated Design Agents

Two concepts are central in the development of situated design agents. *Situated cognition*: This is founded on the view that “where you are when you do what you do matters”, i.e. an agent does not simply react reflexively in its environment but uses its knowledge and its interpretation of its current environment to produce a response (Clancey 1997). The foundational concepts of situatedness go back to the work of Dewey (1896) and Bartlett (1932). There is now considerable empirical evidence from design research to support this claim (Schön and Wiggins 1992; Suwa, Gero and Purcell 1999). The effect of this is that both the agent’s experience and its environment play important roles in determining the agent’s response. As a consequence the agent can be exposed to different environments and produce appropriate responses. A framework for situated cognition in a design agent has been developed by Gero and Fujii (2000). Based on this work, a number of design agents that embody situated cognition have recently been or are being implemented (Gero and Reffat 2001; Kulinski and Gero 2001; Smith and Gero 2001).

Constructive memory: This is another fundamental concept for the development of situated design agents, as Gero (1999) pointed out. It is best exemplified by a quote from Dewey via Clancey: “Sequences of acts are composed such that subsequent experiences categorize and hence give meaning to what was experienced before”. The implication of this is that memory is not laid down and fixed at the time of the original sense experience but is a function of what comes later as well. Memories can therefore be viewed as being constructed in response to a specific demand, based on the original experience as well as the situation pertaining at the time of the demand for this memory. Each memory, after it has been constructed, is added to the agent’s knowledge and is now available to be used later, when new demands require the construction of further memories. These new memories can be viewed as new interpretations of the agent’s augmented knowledge.

Any system that uses agents embodying these two concepts is no longer static, as they are responsive to their environment. The situated design agents developed to date have produced the classes of results expected of them, i.e. that they are responsive to both their experiences and their environments in ways that static agents are not. As a consequence, all

representations are generated for the particular need in the current situation, rather than stored independently of their use.

2.2 Multi-Agent Societies

Societies of agents have been studied in several areas within computational agents research. Multi-agent system (MAS) research has traditionally focussed on the coordination of agents to accomplish a common goal. Its main concerns are practical issues like task allocation among agents, communication protocols, coherent decision making, reasoning about agents, conflict resolution and implementation issues (Bond and Gasser 1988). In recent years, there has been some focus on providing agents with social abilities (Castelfranchi 1998; Panzarasa and Jennings 2001). All these efforts aim at facilitating coordination within existing societies of agents rather than creating these societies.

Artificial Life (ALife) is the field concerned with simulating bottom-up phenomena, where high-level behaviour emerges from low-level interactions (Langton 1989). Typical examples include swarm behaviour and biological and socio-economic evolution (Kauffman 1993; Resnick 1994; Epstein and Axtell 1996). The emergent global patterns achieved by these systems are the result of simple local rules encoded in the agents. The simplicity of these agents does not allow for feedback effects from the top down, which are needed, particularly in the design domain, to avoid randomness by guiding the process of emergence in useful directions. Most ALife-inspired design configuration systems (e.g. Campbell et al. 1998) have used evolutionary computation to address this issue. Some design examples comprise self-configured robots (Lipson and Pollack 2000) and manufacturing systems (Vaario and Ueda 1996).

Behaviour-based Artificial Intelligence (AI) similarly focuses on agent architectures that emerge complex behaviour from a collection of simple entities interacting with the environment. Prominent examples are Brooks' (1986) subsumption architecture and Maes' (1989) action selection system. The problem of how to emerge useful and coherent behaviour has been solved in these approaches by defining the possible interactions of the agents and encoding their relationships.

A synthesis of the research in cooperating agents forming societies and the concept of situatedness is still inadequately developed to date. An approach to product modelling based on societies of situated agents has the potential to be used to produce representations appropriate to the specific need that has been set as the goal. Each new need may produce a different representation appropriate to that need.

3. TOWARDS A FRAMEWORK FOR DYNAMIC SOCIETIES OF SITUATED AGENTS

In this paper we propose a conception of product modelling that is founded on a distributed set of situated agents representing potential components of a product. These representations are not fixed; rather, they are created in response to the specific goal in the specific situation.

3.1 Different Societies for Different Requirements

The situation that is relevant in a certain design task is governed by the specific set of product requirements. Our aim is a system of agents (within an agent base) that responds to the given functional requirements by producing an appropriate structural representation of the product. As the agents are specialised to represent specific artefacts that are the components of the required product, they cannot fulfill this task individually. They have to form a society to represent the product as a whole. Figure 1 depicts how different product models, formed by societies of agents, emerge according to the current situation produced through specific product requirements.

Every emergent product model is the external representation of a distinct society of agents, however the agents themselves remain located at all times within the agent base. Consequently, the agent base contains both “active” agents that currently form a society and “passive” agents that are currently not in use but may become “active” when changes in the situation occur. As Figure 1 shows, agents that participate in one society (or sub-society) can simultaneously participate in another society for a different set of requirements.

Figure 2 depicts an example for our agent-based product modelling approach. Here two different societies of agents as sub-assemblies of mechanical products emerge from a given set of 10 agents. Not all of these agents are needed in the formation of the two societies (valve agent and spring agent). Some agents are exclusively used in one society (gear agent, nut agent, housing agent, roller agent and sleeve agent), whereas others participate in both societies (shaft agent, bearing agent and seal agent) by representing themselves differently. Agents can virtually also be present several times in a society by producing multiple representations (bearing agent, seal agent and sleeve agent).

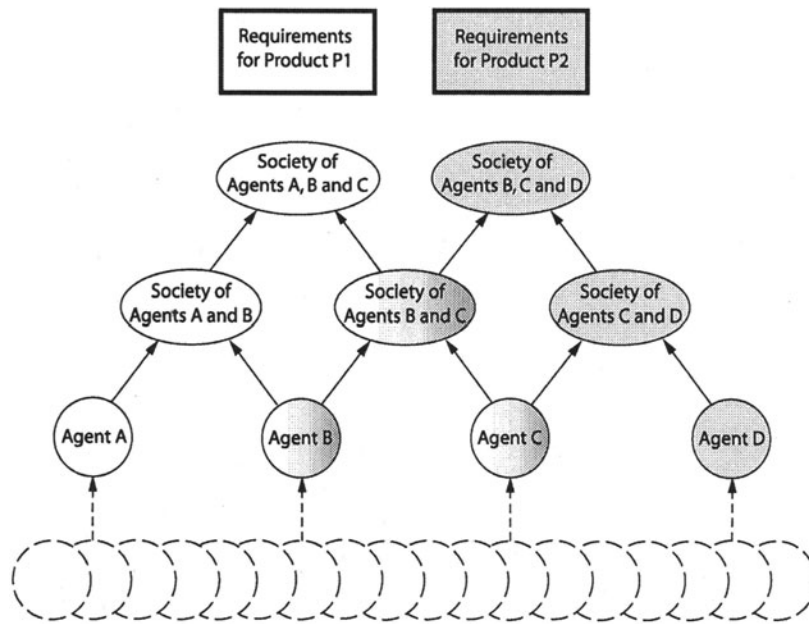


Figure 1. Individual agents, within an agent base, form societies in response to specific product requirements




Figure 2. Different societies out of one set of agents

3.2 Different Representations for Different Purposes

Once the agents have formed a society and thus a product, they represent this society/product according to the specific informational needs of the users or software tools. As Figure 3 exemplifies, the marketing person is provided with different product data, e.g. economic data, than the production planner and the designer, who are provided with a component list and technical drawings, respectively. As situated agents produce representations in response to their current situation, in which the required representational purpose plays an important role, this approach has the potential to suit the different needs of re-representation of product data.

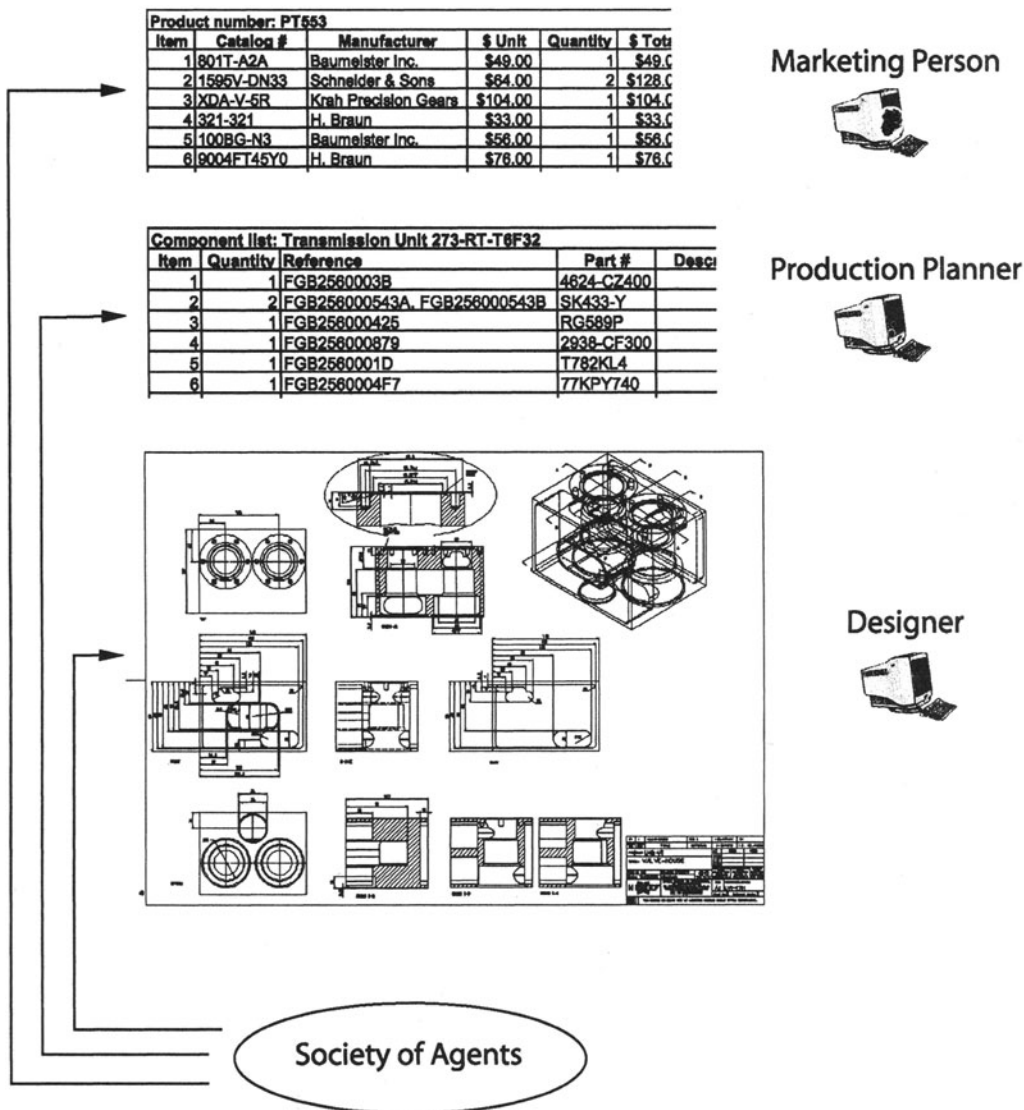


Figure 3. The same society of agents represents itself differently for different purposes

4. THE ISSUE OF SELF-ORGANISATION

Section 3 has outlined the fundamental concepts underpinning product modelling using a set of situated agents. However, the issue arises of how a loose collection of situated agents can form a society. As our situated agents are autonomous and rational (Smith and Gero 2000), this cannot be achieved by pre-defining their relationships, which would restrict their autonomy. They need to self-organise their societies autonomously, and this ability must reside in the agents themselves, i.e. in their rationality. This involves social reasoning and requires that the individual agents are able to construct concepts of possible societies as well as of other agents. This is enabled primarily through the interpretation of the current external representations of other agents.

4.1 An Example of a Self-Organisation Task

Figure 4 depicts an example of a sequence of external representations that a set of three different agents produces in order to form a society. The sequence contains six “snapshots” of a self-organisation task and starts with the given design requirements (1) that can be seen (interpreted) by all the agents in the agent base. The shaft agent responds first in our example and contributes to a potential design solution by representing two shafts (2). The bearing agent, by using its interpretation about the shaft agent’s external representation, then produces an appropriate representation of two bearings that guide the two shafts (3). This is perceived by the shaft agent, which modifies its representation to create shoulders to retain the bearings in position. The next “snapshot” (4) shows how the gear agent comes into play by representing two pairs of gears. This is followed by the shaft agent modifying its representation again to introduce another shoulder for two of these gears. The other two gears are to be connected to a second shaft, whose representation is subsequently produced by the shaft agent (5). The bearing agent reacts to the introduction of an additional shaft by representing two new bearings for the shaft (6), followed by appropriate modifications of the shaft.

This introductory example of self-organisation has been kept very simple and does not show other types of representations than geometry. Yet in most cases, representations that are tailored to the specific informational needs of particular agents have to be produced. In the context of engine components, for example, the crankshaft agent is likely to be interested in the average speed of the piston rather than in its diameter. The piston agent, aware of this specific need (through receiving or anticipating a respective request from the crankshaft agent), will then provide the crankshaft agent with the desired

information. This makes obvious that every external representation may be viewed as the result of a *social action*, which is a type of action whose “subjective meaning takes account of the behaviour of others and is thereby oriented in its course” (Weber 1968). This social character of all inter-agent representations has not been emphasised in Figure 4.

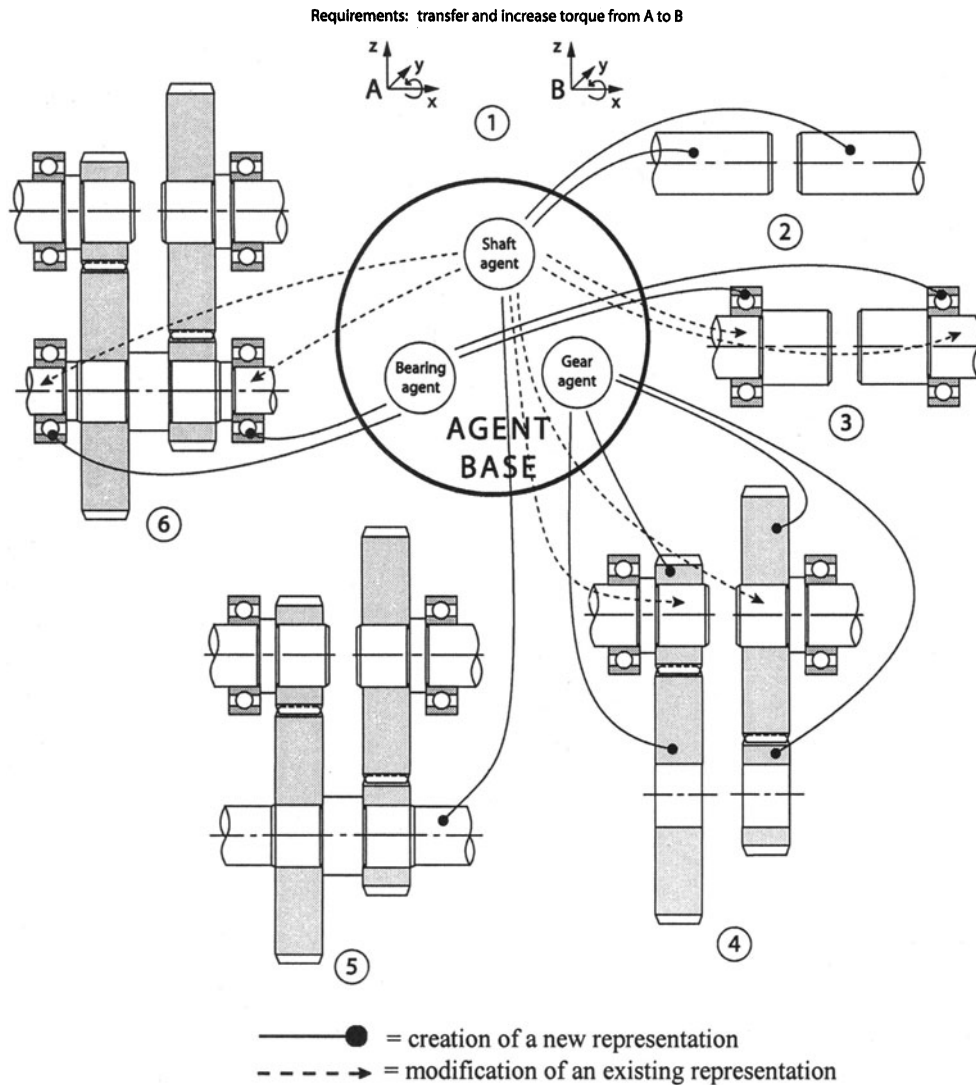


Figure 4. An example of a self-organisation process

The depicted example for self-organisation also looks very linear, which is partly due to the small number of agents. For more complex assemblies, it is unlikely that individual agents can self-organise merely by interpreting and reacting upon external representations of other agents. A more direct form of communication among the agents is necessary. One agent might not

comprehend or even agree with another agent's intention that is inferable from its current representation. In this case, the agents have to communicate to remove any misunderstandings and disagreements and come to a common solution, which corresponds to their self-organisation.

4.2 Self-Organisation as Increasing Consensus

The idea of organisations (societies) based on the consensus among individuals has been well expressed by Habermas (1984). In particular, he views communication as a crucial means for achieving consensus and defines *communicative action* as "a social action oriented to reaching understanding". Here the notion of understanding corresponds to the mutual, intersubjective agreement on several levels, including the society level. This view conforms to our purpose of autonomous agents coming to a consensus based on their individual view of the world.

This broad meaning of understanding for an agent requires in the first instance that it can construct concepts relating to societies, which are, in terms of design, (sub-)assemblies and the roles (functions) that the components play in them. Formally speaking, every agent must be able to form a design state space not only of the artefact that the agent is specialised for, but also a design state space of (sub-) assemblies including all their components. The union of all the design state spaces of a particular agent X we will denote as \mathcal{D}_X , which represents the agent's individual view of the world in terms of its current design goals or hypotheses.

Further, as understanding is established through mutual agreement on only one common view of the world, an autonomous agent must be able to distinguish between its own and another agent's view. Formally speaking, agent X must be able to construct a belief about the design state space of agent Y (which can be agent X itself), which we will denote as $\mathcal{B}_X(\mathcal{D}_Y)$. For three agents A, B and C, we get:

- for agent A: $\mathcal{B}_A(\mathcal{D}_A)$: A's own design state space
- $\mathcal{B}_A(\mathcal{D}_B)$: A's belief about B's design state space
- $\mathcal{B}_A(\mathcal{D}_C)$: A's belief about C's design state space
- for agent B: $\mathcal{B}_B(\mathcal{D}_B)$: B's own design state space
- $\mathcal{B}_B(\mathcal{D}_C)$: B's belief about C's design state space
- $\mathcal{B}_B(\mathcal{D}_A)$: B's belief about A's design state space
- for agent C: $\mathcal{B}_C(\mathcal{D}_C)$: C's own design state space
- $\mathcal{B}_C(\mathcal{D}_A)$: C's belief about A's design state space
- $\mathcal{B}_C(\mathcal{D}_B)$: C's belief about B's design state space

Reasoning about other agents' goals has been introduced in computational MAS many years ago (Sycara 1988). Applied to our model,

an agent can evaluate how far its own design goals are consistent with those of other agents. Figure 5 illustrates this for the three agents A, B and C. Consistent design state spaces can be assumed by an agent X (which can be agent A, B or C) to exist for some states of affairs among the whole set of agents (area abc), for other states of affairs among partial sets of agents (areas ab, bc and ac) and for certain states of affairs not at all (areas a, b and c). The ultimate aim of every agent is to increase inter-agent consensus and to eliminate dissension, either by changing its own design state space or by trying to change those of the other agents via negotiation.

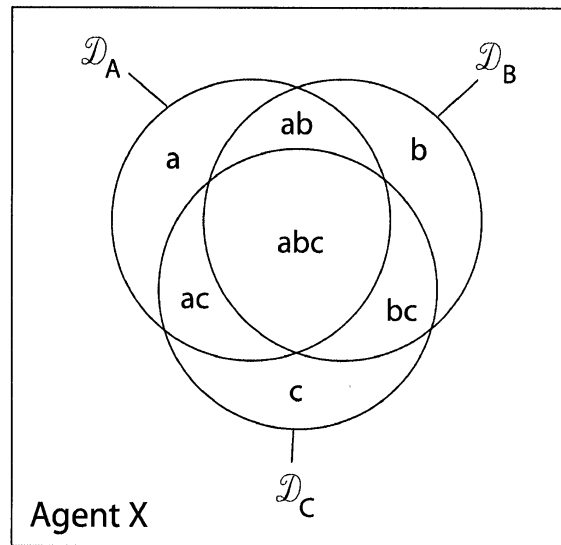


Figure 5. Evaluating the consistency of different design state spaces

5. COMMON GROUND AS A BASIS FOR COMMUNICATION AND SELF-ORGANISATION

Modelling the goals of other agents is necessary but not sufficient for our purposes. For instance, if an agent is to evaluate the current state of self-organisation, a one-sided comparison of the agent's own design state space with its beliefs about other agents' design state spaces might not suffice: What might look at first sight for agent A like a consensus with agent B could turn out not to hold for B if B's belief about A's design state space differs from A's actual design state space. Therefore, if an agent is to evaluate consensus with another agent, then it has to construct not only its individual model of both agents' goals (as in Figure 5), but also a belief about what the other agent considers as consensus.

This leads to an issue that is known in cognitive science as mutual knowledge or common ground (Clark 1992). Common ground describes a presupposed shared background that can sufficiently account for the success of an interaction. The content conveyed by a communicative action has to fit into this shared background, otherwise communication is likely to fail resulting in agents not understanding or misunderstanding one another. It has been shown that agents take their common ground into account when producing and comprehending communicative actions (Clark and Murphy 1982; Fussell and Krauss 1989). In contrast to pre-defined shared ontologies used in AI research (Gruber 1993), common ground is based on situated constructions of individual agents. During the course of a communicative interaction, the common ground among the agents constantly increases rather than being fixed from the outset.

The problem of an agent using common ground for communication has been recognised as one of infinite regress, involving recursive constructions of beliefs about other agents' beliefs (Clark and Marshall 1981). However, situated agents can avoid this recursion by "satisficing" (Simon 1969) common ground by constructing adequate models of one another according to the needs of the current interaction. We have proposed elsewhere (Gero and Kannengiesser 2003) a representation schema for agents, which supports this view of common ground as emerging from satisficed individual constructions. This schema provides a set of constructs that account not only for the beliefs and knowledge but also for the behaviour and the teleology of agents. We expect that the use of this schema will make communication both more efficient and more effective.

6. DISCUSSION

ISO STEP and IFCs offer a good solution to the interoperability problem, They both use the concept of a neutral format file into which the source data formats are converted. In order to be interoperable any system must be able to read from that neutral format file as well as write into it. Once the translators to and from the neutral format file have been written then interoperability is assured and maintained. Writing these translators is an arduous task as it can involve obtaining the agreement of the majority of representative users in a particular segment of an industry. This potentially creates a barrier to the development and introduction of those novel systems where the marginal cost of producing interoperability in this way is high. This paper addresses this issue by proposing an approach that generates an appropriate translation format when none is available.

The approach suggested makes use of situated cognitive agents that communicate amongst themselves to form a society that is the product and its interaction with a tool. This social behaviour allows for a common ground to develop between the product and the tool as the components of the product organise themselves in such a way as to present themselves in an appropriate format for the tool. In this way a new system or tool can be used. This approach would then provide the beginning of an ontology to support an IFC.

ACKNOWLEDGEMENTS

This research is supported by a Sesqui Research and Development grant from the University of Sydney and by an International Postgraduate Research Scholarship. Computing resources are provided by the Key Centre of Design Computing and Cognition, University of Sydney.

REFERENCES

- Bartlett FC (1932 reprinted in 1977) *Remembering: A Study in Experimental and Social Psychology*, Cambridge University Press, Cambridge.
- Bond AH and Gasser L (eds) (1988) *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann Publishers, San Mateo, CA.
- Brooks RA (1986) A robust layered control system for a mobile robot, *IEEE Journal on Robotics and Automation* 2(1): 14-23.
- Campbell MI, Cagan J and Kotovsky K (1998) A-design: Theory and implementation of an adaptive, agent-based method of conceptual design, in JS Gero and F Sudweeks (eds), *Artificial Intelligence in Design '98*, Kluwer, Dordrecht, pp. 579-598.
- Castelfranchi C (1998) Modelling social action for AI agents, *Artificial Intelligence* 103(1-2): 157-182.
- Clancey WJ (1997) *Situated Cognition*, Cambridge University Press, Cambridge.
- Clark HH (1992) *Arenas of Language Use*, University of Chicago Press, Chicago, IL.
- Clark HH and Marshall CR (1981) Definite reference and mutual knowledge, in AK Joshi and B Webber (eds), *Linguistics Structure and Discourse Setting*, Cambridge University Press, Cambridge, pp. 10-63.
- Clark HH and Murphy GL (1982) Audience design in meaning and reference, in JF Le Ny and W Kintsch (eds), *Language and Comprehension*, North-Holland Publishing Company, Amsterdam, pp. 287-299.
- Dewey J (1896 reprinted in 1981) The reflex arc concept in psychology, *Psychological Review* 3, 357-370.
- Eastman CM (1999) *Building Product Models: Computer Environments Supporting Design and Construction*, CRC Press, Boca Raton, FL.
- Epstein JM and Axtell R (1996) *Growing Artificial Societies*, MIT Press, Cambridge, MA.

- Fussell SR and Krauss RM (1989) The effects of intended audience on message production and comprehension: Reference in a common ground framework, *Journal of Experimental Social Psychology* **25**: 203-219.
- Gero JS (1990) Design prototypes: a knowledge representation schema for design, *AI Magazine* **11**(4): 26-36.
- Gero JS (1999) Constructive memory in design thinking, in G Goldschmidt and W Porter (eds), *Design Thinking Research Symposium: Design Representation*, MIT, Cambridge, pp. 1.29-35.
- Gero JS and Fujii H (2000) A computational framework for concept formation for a situated design agent, *Knowledge-Based Systems* **13**(6): 361-368.
- Gero JS and Kannengiesser U (2003) Function-Behaviour-Structure: A model for social situated agents, *Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions, International Joint Conference on Artificial Intelligence 2003*, Acapulco, Mexico, (to appear)
- Gero JS and Reffat R (2001) Multiple representations as a platform for situated learning systems in design, *Knowledge-Based Systems* **14**(7): 337-351.
- Gruber TR (1993) A translation approach to portable ontologies, *Knowledge Acquisition* **5**(2): 199-220.
- Habermas J (1984) *The Theory of Communicative Action*, Vol. 1, Beacon Press, Boston.
- IAI (2000) *Industry Foundation Classes*, International Alliance for Interoperability, http://eande.lbl.gov/btp/iai/copyright_ifc2x.html.
- ISO (1994) *ISO 10303-11 Product Data Representation and Exchange – Part 11: Description methods: The EXPRESS language reference manual*, International Standards Organization, Geneva, Switzerland.
- Kauffman SA (1993) *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, Oxford.
- Kulinski J and Gero JS (2001) Constructive representation in situated analogy in design, in B de Vries, J van Leeuwen and H Achten (eds), *CAADFutures 2001*, Kluwer, Dordrecht, pp. 507-520.
- Langton CG (1989) Artificial life, in CG Langton (ed.), *Artificial Life*, Addison-Wesley, Reading, MA, pp. 1-47.
- Lipson H and Pollack JB (2000) Automatic design and manufacture of robotic lifeforms, *Nature* **406**, 974-978.
- Maes P (1989) The dynamics of action selection, in M Kaufmann (ed.), *Proceedings of the International Joint Conference on Artificial Intelligence-'89*, Detroit, MI, pp. 991-998.
- NIST (1999) Interoperability cost analysis of the US automotive supply chain, *Planning Report #99-1*, NIST Strategic Planning and Economic Assessment Office, Gaithersberg, Maryland.
- Panzarasa P and Jennings NR (2001) The organisation of sociality: A manifesto for a new science of multi-agent systems, in *Proc 10th European Workshop on Multi-Agent Systems (MAAMAW-01)*, Annecy, France.
- Resnick M (1994) *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*, MIT Press, Cambridge, MA.
- Schön DA and Wiggins G (1992) Kinds of seeing and their functions in designing, *Design Studies* **13**(2): 135-156.
- Simon HA (1969) *The Sciences of the Artificial*, MIT Press, Cambridge, MA.
- Smith G and Gero JS (2000) The autonomous, rational design agent, in H Fujii (ed.), *Workshop on Situatedness in Design*, Artificial Intelligence in Design'00, Worcester, MA, pp. 19-23.

- Smith G and Gero JS (2001) Situated design interpretation using a configuration of actor capabilities, in JS Gero, S Chase and MA Rosenman (eds), *CAADRIA 2001*, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, pp. 15-24.
- Suwa M, Gero JS and Purcell T (1999) Unexpected discoveries and s-inventions of design requirements: A key to creative designs, in JS Gero and ML Maher (eds), *Computational Models of Creative Design IV*, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, pp. 297-320.
- Sycara K (1988) Resolving goal conflicts via negotiation, in *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, St. Paul, MN, pp. 245-250.
- Vaario J and Ueda K (1996) Self-organization in manufacturing systems, in K Stelson and F Oba (eds), *Proceedings of the 1996 Japan/USA Symposium on Flexible Automation*, Boston, MA, pp. 1481-1484.
- Weber M (1968) *Economy and Society: An Outline of Interpretive Sociology*, Bedminster Press, New York.