MODELING AND SIMULATING WORK PRACTICES FROM APOLLO 12

Maarten Sierhuis[1,2]
William J. Clancey[1]
Ron van Hoof[1]
Robert de Hoog[2]

1. NASA/Ames Research Center
Moffett Field, CA 90435-1000
USA
{msierhuis,bclancey,rvanhoof}@mail.arc.nasa.gov

2. University of Amsterdam
Dept. of Social Science Informatics
The Netherlands
dehoog@swi.psy.uva.nl

ABSTRACT

This paper describes modeling and simulating the activities of humans and systems in organizations. We refer to this as work practice modeling. We describe a case study of the work practice of the Apollo 12 astronauts during the ALSEP offload activity.

1. INTRODUCTION

At NASA we develop new work systems for exploring extraterrestrial bodies, and scientific research in space. The work performed during a Space Shuttle flight is the result of a developed work system. The International Space Station will need a new and different work system, not yet fully developed. When eventually we will go to Mars or back to the Moon we need a different type of work system. A small group of humans will work together with sophisticated autonomous and semi-autonomous robots. Collaboration between the people on Earth and the people on Mars or the Moon will be of a different order than what we have been used to. On Mars, a forty-minute communication delay back to Earth changes the nature of collaboration.

In designing and implementing collaborative work places we need to understand the way humans and robots can work together, and analyze existing work practices in order to design improvements or changes. Work practice analysis, design methods and tools need to be developed to understand not only the process, but also the practice of an organization. To model the situated activities and collaboration of people, we need a theory and modeling language that incorporates the aspects of practice.

Work practice includes those aspects of a work process that make people behave a certain way. To describe people's situation-specific behavior we need to describe those aspects of the situation that explain the influence on the activity behavior of individuals (in contrast with problem-solving behavior). The important aspects that determine work practice are; the individuals and their activities, the context in terms of artifacts and tools, location and movement of artifacts and individuals, the interpretation of facts into an individual's beliefs on which actions are based, the communication between individuals, and the communication tools used.

We have implemented our theory of work practice modeling in the Brahms modeling and simulation environment. Brahms consists of a multi-agent modeling language, a discrete-event simulator, a history database of the simulation runs, and a tool for visualizing the activity-behavior of, and communication between agents and artifacts, as well as the geographical locations and movement (Clancey et al. 1998).

In this paper, we first explain the term work practice, and describe our theory of modeling work practice. Then, we describe a case study using our approach. We describe a small part of the work practice model of the astronauts deploying the Apollo Lunar Surface Experiment Package (ALSEP) during the Apollo 12 lunar mission.

2. WORK PRACTICE

Work practice is a concept that originates in socio-technical systems, business anthropology, work systems design, and management science.

The notion of "practice" is central to work systems design, which has its roots in the design of socio-technical systems, a method developed in the 1950s by Eric Trist and Fred Emery (Emery and Trist 1960). *Socio-technical systems design* sought to analyze the relationship of the social system and the technical system, such as manufacturing machinery, and then design a "socio-technical system" that leveraged the advantages of each. *Work systems design* extends this tradition by focusing on both the formal features of work (explicit, intentional) and the informal features of work (as it is actually carried out "in practice," analyzed with the use of ethnographic techniques) (Ehn 1988) (Greenbaum and Kyng 1991) (Pasmore 1993) (Weisbord 1987, chapter 16).

We are interested in describing work as a practice; a collection of psychologically and socially situated collaborative activities between members of a group. We try to understand how, when, where, and why activities occur. The central theme is to find a representation for *modeling work practice*. Many researchers in the social sciences use the word practice as if it is a well-defined concept everyone understands. However, it is difficult to describe a practice. People notice when something is not a practice, and can often say why. Although it can be said that a group of people has developed a practice, when asked to describe it we find it difficult. As such, practice is part of our tacit knowledge (Polanyi 1983).

In the past century, work has been defined as the transformation of input to output, starting with Frederick W. Taylor's view of work to Michael Hammer's view of business processes (Hammer and Champy 1993). For example, a manufacturing process has well-defined inputs and outputs for each step of the process. Sometimes, however, it is more difficult to describe the input and output of the work. Consider a soccer match between two professional soccer teams. It is difficult to define the input and output of this type of work, even though most of us would agree that professional soccer players are working. To describe the work of a soccer team we quickly fall into descriptions of teamwork and collaboration on and off the field.

An ad hoc definition of a practice is: *The (collaborative) performance of situated activities in real life situations, by making use of knowledge previously gained through experience in performing similar activities.*

In short, practice is doing in action (Suchman 1987). Scientists have described how a practice develops, like Wenger, who defines the creation of a practice as follows (Wenger 1997):

> Being alive as human beings means that we are constantly engaged in the pursuit of enterprises of all kinds, from ensuring our physical survival to seeking the most lofty pleasures. As we define these enterprises and engage in their pursuit together, we interact with each other and with the world and we tune our relations with each other and with the world accordingly. In other words, we learn. Over time, this collective learning results in practices, which reflect both the pursuit of our enterprises and the attendant social relations. These practices are thus the property of a kind of community created over time by the sustained pursuit of a shared enterprise.

Everybody knows what Wenger means when he says, "this collective learning results in practices", but what is it that results? Can it be described? Can it be modeled? To do this, we need to describe practice at an epistemological level we call the *work practice level*. In the rest of this paper we discuss a computational language for modeling work practice. Computational work practice models can be simulated to show the emergent effects of the activities of people and their communication, situated in a geographical environment, as well as the use of tools and artifacts.

3. THEORY OF MODELING WORK PRACTICE

We briefly describe our theory of modeling work practice. Representing how people work can be done at many different levels. In knowledge engineering and artificial intelligence (AI), people's work has been described in terms of their problem-solving expertise. The theory is that we can model people's problem-solving behavior by representing this behavior in a computational model that is able to duplicate some of this behavior. Work process models, such as Petri-Net models of a work process, describe what tasks are performed and when. Workflow models describe how a specific product "flows" through an organization's work process. This describes the sequential tasks in the work process that "touch" a work-product. All these modeling approaches describe the work in an organization at a certain level of detail. However, what is missing from all these types of modeling approaches is a representation of how the work actually gets done.

Work practice includes those aspects of the work process that make people behave a certain way in specific situations, and at specific moments in time. To describe people's situation-specific behavior we need to include those aspects that explain the *activity behavior* of individuals (in contrast with problem-solving behavior). Following is a brief description of the important aspects that determine an individual's situation-specific behavior.

Modeling Activity Behavior

People's behaviors are determined by the "execution" of specific activities at certain moments. A person or system cannot be "alive" without being in some kind of activity. Even "doing nothing" is described in terms of a "do-nothing" or idle activity. Furthermore, what activity is being performed depends on the context. Agent behaviors are organized into activities. Most importantly, activities locate behaviors of people and their tools in time and space, so that resource availability and informal human participation can be taken into account (Vygotsky 1978, Originally published in Rusian in 1934).

Activities can be subsumed by other activities in a hierarchical structure (Nakashima et al. 1996) (Brooks 1991). With this we mean that a person can be in multiple activities at once. For example, we can be in the activity of reading a book, while at the same time be in the higher level activity of being on a business trip. When the phone rings in our hotel room, we get up and walk over to pick up the phone. We interrupt the activity of reading our book, and start the activity of answering the phone. In a sense, we never stop being in the activity of reading our book, but we suspend the activity to focus on a new activity continuing with the suspended activity when the phone call is over.

A model of activities doesn't necessarily describe the intricate details of reasoning or calculation, but instead captures aspects of the social-physical context, including space and time in which reasoning occurs (Clancey 1997).

Modeling Context

People act based on their interpretation of a situation. With this we mean that people behave based on their beliefs about what they experience (infer or detect) as their context. Different people can/will have different beliefs about a similar context. If we want to model work practice, we need to be able to separate the context from people's different interpretation of that context. In order to do so, we describe context in terms of the objects and artifacts people observe and use within their environment (Agre 1995). We also describe the geographical locations of people and artifacts (Kirsh 1995). What describes a context is known as world-facts or simply the facts about the context. Facts represent factual information about the three-dimensional world people live in. People do not automatically have "knowledge" about those facts, and if people have "knowledge" about those facts it might not be correct. For example, you can believe that your car is parked in the garage, whereas in reality someone has taken the car to go out. So, the fact is that the location of the car is wherever it has been taken, while you believe that the location of the car is the garage. You will have that belief until either someone tells you about the actual location (or wrong location) of the car, or until you go to the garage and observe (i.e. detect) that the car is not there. Of course, if the car returns before any of this takes place you will never know the car had been gone. In other words, although facts are global (the car can only be in one location), not every person can get "access" (i.e. get a belief) about that fact.

Modeling Communication

In order for two or more people to collaborate they need to communicate. In Speech Act theory the meaning and intent of certain speech acts are formalized (Searle 1969). Using this type of communication analysis, we can model the sequence of (communication) actions in a collaborative activity, as well as the intention and meaning of the speech acts. However, in analyzing the way collaboration occurs in practice, we also need to analyze communication in terms of how it actually happens in the real world, thereby modeling collaboration as it really occurs. Speech Act theory analyzes communication in terms of patterns of commitment entered into by the sender and the receiver. While this is important, it does not take into account the communication tools used in the situated speech act. Today, communication is more and more efficient and certain communication tools are used globally. Phones, voice mail, e-mail, and fax, are communication tools that are more and more taken for granted in the way that we use them. However, it should not be taken for granted that we all have created our own practice around the use of these tools in certain situations.

This emphasizes the point that communication is very much dependent on our practice surrounding our communication tools, and that we therefore need to include the use of communication tools in modeling how people actually coordinate their collaboration in the real world. We need to include a model of the workings of communication tools, and how they are used in practice.

Modeling Communities of Practice

In order to describe different people performing the same activities given the same context, we borrow the term *community of practice* (CoP) from the social sciences (Wenger 1997). People belong to many different communities. One way to distinguish one community from another is by the different activities they perform. For instance, at NASA we can distinguish the community of Apollo astronauts from the rest of the communities at NASA. We describe a particular community as a separate "group." Members of a group can perform the group's activities. Thus, we can describe people's activity behavior in terms of the groups they belong to.

4. CASE STUDY OBJECTIVES

The goal of the case study we describe next was to investigate the use of the Brahms-language to *describe an existing work practice*. The objectives were:

- Being able to represent the people, things, and places relevant to the domain.
- Represent the actual behavior of the people, second by second, over time.
- Show which of the tools and artifacts are used when, and by whom, to perform certain activities.
- Include the communication between co-located and distributed people, as well as the communication tools used, and the effects of these communication tools on the practice (e.g. communication delay using the Earth to Moon voice-loop)

The domain we chose for this experiment, and we will describe in this paper, is the work practice of the Apollo 12 astronauts in the deployment of the Apollo Lunar Surface Experiments Package (ALSEP) on the Moon.

5. APOLLO 12 AND THE ALSEP OFFLOAD

One of the biggest objectives of the Apollo 12 mission was to deploy the Apollo Lunar Surface Experiment Package (ALSEP). The ALSEP consisted of a number of independent scientific instruments that were to be deployed on the moon. The instruments were data collection devices for different scientific experiments about the moon's internal and external environment. By deploying similar ALSEP instruments over multiple Apollo missions (A12, 14, 15, 16 and 17), the ALSEP deployments created an array of data gathering instruments at different locations on the lunar surface.

To deploy the ALSEP on the lunar surface the astronauts had to accomplish three high-level tasks. First, they had to *offload* the ALSEP from the Lunar Module (LM). Second, they had to *traverse* with the ALSEP packages to the deployment area, away from the LM. Third, they had to *deploy* each ALSEP instrument onto the surface. In this paper, we discuss the development of a work practice model for the first task, the *ALSEP Offload*. Even though this high-level task was planned and choreographed up front, the plan did not include the situational variations and the actual communications.

The work practice of the ALSEP Offload, or any work practice for that matter, consists of more than the sequence and distribution of tasks. As discussed in the previous sections, what constitutes the practice of the ALSEP Offload is the way the plan is carried out. The situational activities of the collaborators, the way they react to their environment, the way they communicate, what is said, the way they "know" how to do their tasks given the situation.

In the next sections, we will briefly describe how the ALSEP Offload work is represented in a model of work practice.

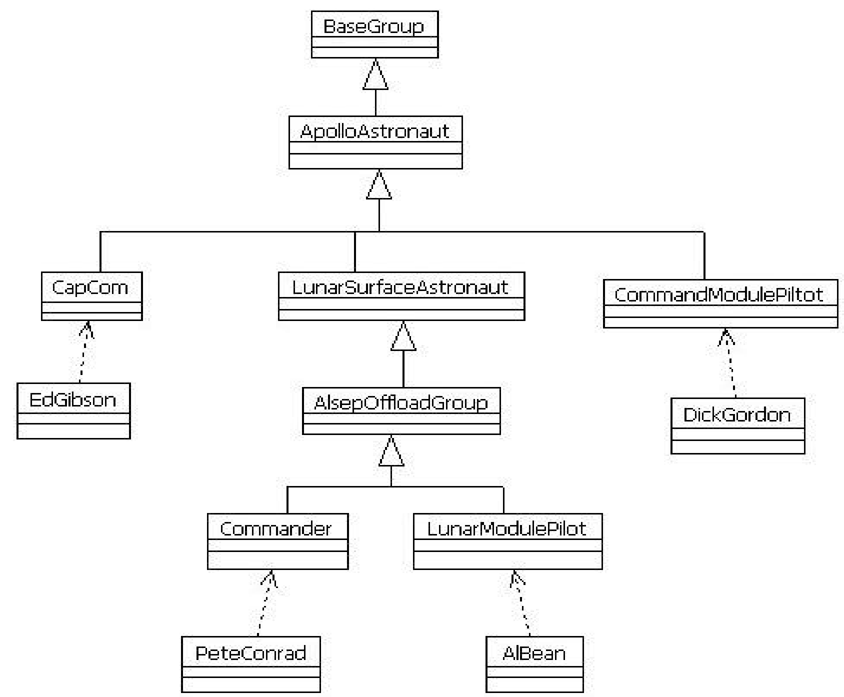

Figure 1. Apollo Agent Model Design

6. AGENT MODEL

One of the most relevant design for any Brahms model is the design of the agents and the groups they belong to. The Agent Model describes to which groups the agents belong and how these groups are related to each other.

As a rule of thumb, we identify the communities of which the agents in the model are members, and abstract them to a common denominator for all agents. It should be noted that groups and agents could be members of multiple groups.

Figure 1 shows the Agent Model design. We start with defining our agents. Each agent represents a person in our domain, i.e. Ed Gibson, Pete Conrad, Al Bean, and Dick Gordon. We generalize the community all four agents belong to as the group of ApolloAstronauts. We represent the role of each of the astronauts as a group (CapCom, Commander, LunarModelulePilot, CommandModulePilot). This way we can represent role specific attributes and activities at the group level. The AlsepOffloadGroup is a functional group in the sense that it doesn't specify a specific role, but a task of the agent. This group represents all work activities and attributes that have to do with the ALSEP Offload task in one group. The group represents the community of astronauts that trained to perform the ALSEP Offload task. For the Apollo missions, both the Commander (CDR) and the Lunar Module Pilot (LMP) trained for the ALSEP Offload activities, and both of them were able to perform all the ALSEP Offload tasks. Therefore, the Commander and LunarModulePilot groups are members of the group AlsepOffloadGroup. Since both the CDR and the LMP were working on the lunar surface, there are tasks that both astronauts needed and/or could perform. The ALSEP Offload task was one of them, but there were others as well. The other activities that needed to be performed on the lunar surface are represented in the LunarSurfaceAstronaut group. These activities include taking photographs and changing the cooling of their space suit.

7. OBJECT MODEL

After the Agent Model, the next model that needs to be designed is the Object Model. In this model we design the class-hierarchy of all the domain objects. Figure 2 shows the Object Model design for the ALSEP Offload domain objects and artifacts. The root-class of the class hierarchy is the class *BaseClass*. All other classes and objects inherit from this BaseClass class.

Figure 2 also shows the *containment* relation. This relation describes agents and agent containing other objects and agents. When an agent or object moves, all its contained objects and agents move with it. For example, when an astronaut agent moves the tools that are being carried move with the agent.

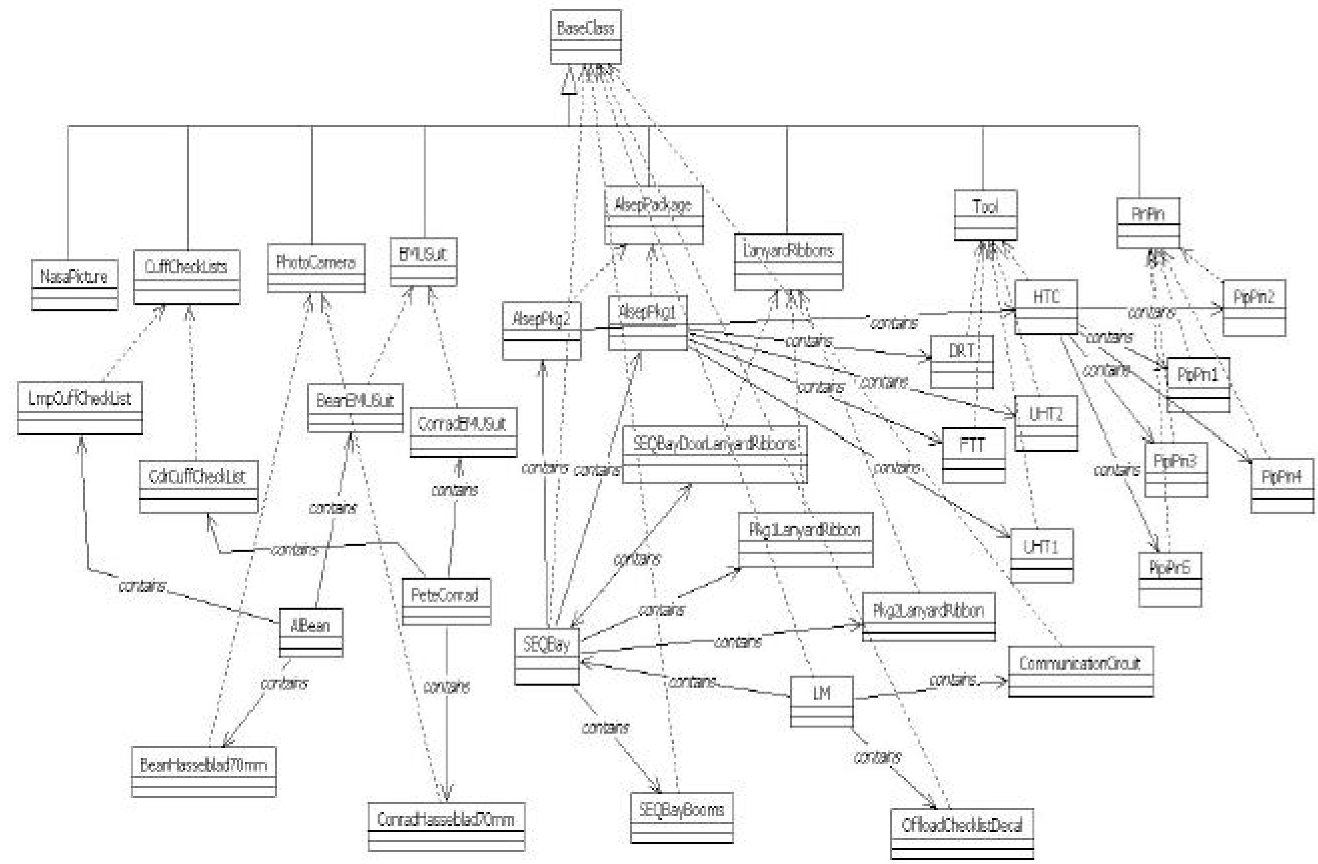

Figure 2. Apollo Object Model Design

8. GEOGRAPHY MODEL

In Brahms we model geographical locations using two concepts, area-definitions and areas. Area-definitions are user-defined types of geographical locations. Areas are instances of area-definitions. An area is a specific location in the real world that is being modeled. Furthermore, areas can be part-of other areas. With this representation scheme we can represent any location at any level of detail.

For the Apollo 12 ALSEP Offload activity, the following locations are modeled; Earth, the Manned Spaceflight Center (MSC), the Moon, the Apollo 12 Landing-site ("Surveyor Crater"), the area where the SEQ Bay is located, the ALSEP deployment area, an area away from the SEQ Bay to place artifacts after offloading, and last, the lunar orbit and the Command Module ("Yankee Clipper"). Figure 3 shows the geography model design.

Movement

Agents and objects can move from one location to another. Moving from one location to another removes the agent from the starting location and moves the agent to the new location. This is accomplished by having the agent perform a move-type activity. The time the activity is active (i.e. the activity duration-time) determines how long it takes the agent to move from location A to location B.

Detecting Agents and Objects

As agents arrive at their new location they will immediately detect facts about the location of other agents and objects that are also in the location. The simulation engine automatically creates beliefs for the agent from the fact that there are other objects and agents in the same location. The agents already in that location will get the belief that the agent that arrived is also in the location. This way, agents notice other agents and objects that are in the location.

9. ACTIVITY MODEL

In Brahms we have developed an organization of human behavior that does *not* use the same type of goal-structures as referred to by Anderson and Lebiere (Anderson and Lebiere 1998, p. 39). The organizing principle of human behavior used in Brahms has its roots in Activity Theory (Vygotsky 1978, Originally published in Rusian in 1934). The most important concept is the concept of an activity. An *activity* is a function performed over a period of time. An activity execution has a well-defined beginning and end. Although the definition of an activity looks similar to that of a

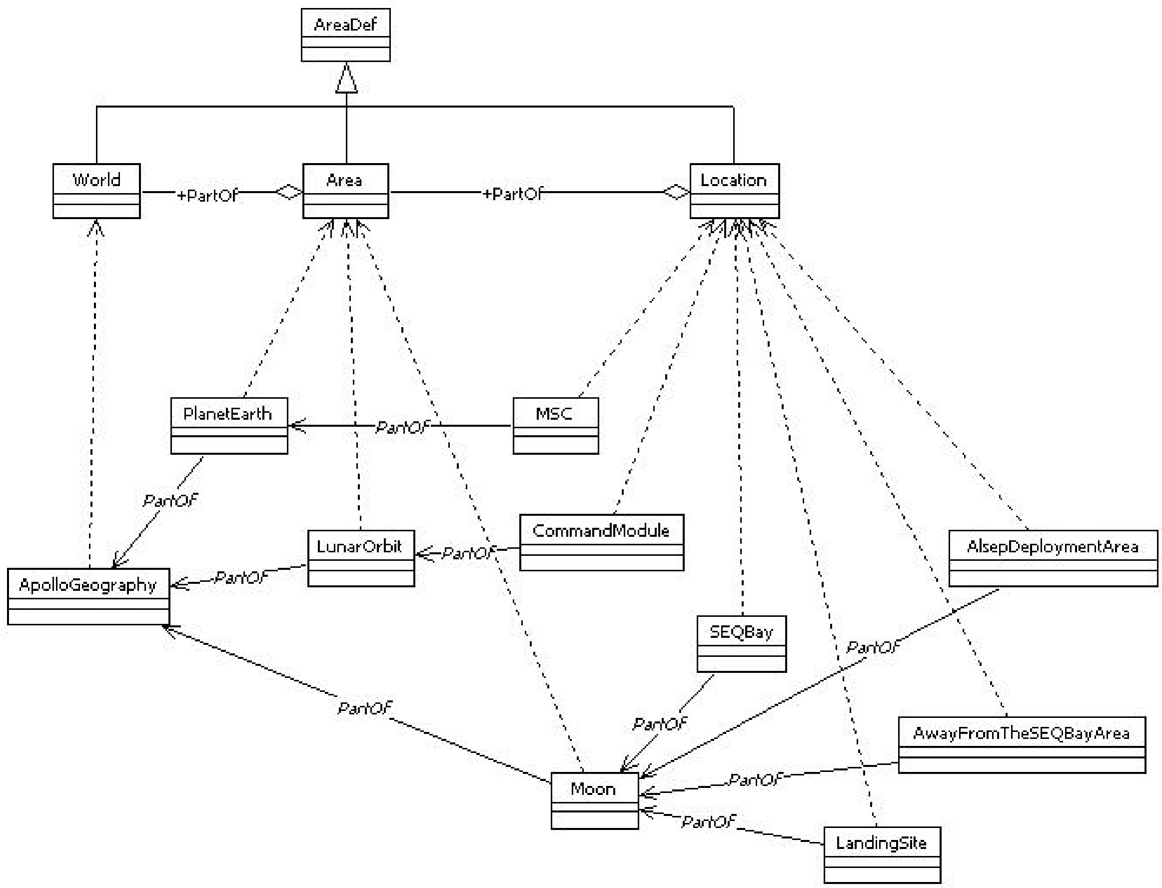

Figure 3. Apollo Geography Model Design

goal or goal state, it is actually closer related to Newell's idea of a *response function* (Newell 1990, p. 44-43).

Just as goals, activities can be decomposed into sub-activities, and thus creating a hierarchical activity-structure that organizes behavior of an agent. An activity-structure is a subsumption hierarchy. This means that while an agent is performing a sub-activity, it is also performing the higher-level activity. Each higher-level activity subsumes all its lower-level activities. The activity-structure in Brahms is implemented as a reactive agent-subsumption architecture (Brooks 1991).

In the next section we describe the *Open SEQ Bay Door* activity performed during the ALSEP Offload. This model represents a part of the work practice of the Apollo 12 lunar surface astronauts as they performed the ALSEP Offload activity.

Open SEQ Bay Door Activity

Table 2 shows the activities and sub-activities of the *Open SEQ Bay Door* activity for both LMP and CDR, mapped onto the communication transcribed in the Apollo Lunar Surface Journal (Jones 1997). The actual names of the activities and sub-activities are more or less arbitrary, and conceptualize the modeler's *interpretation* of the Apollo 12 communication data and the Apollo 14 video data. However, these data and observations is strong evidence that these are the actual activities that were performed during the Open SEQ Bay Door activity.

Each (sub-)activity is "executed" by a *workframe*, which means that when an agent executes the workframe the activity is performed within the context of that workframe. An agent has an individual set of workframes inherited from the groups it belongs to. A workframe is a *production-rule* with preconditions matching the beliefs of an agent. When the preconditions of a workframe match with beliefs of the agent it becomes available. The simulation engine schedules the next activity of an agent based on her set of current, available and interrupted workframes.

Table 1. Workframe

```
workframe MovingToSEQBay {
  when (not(current.location = SEQBayArea))
  do {
      Move(SEQBayArea);
  }
}
```

Table 1 shows the Brahms source code for a workframe. The semantics are as follows: When the agent does not believe it is in *SEQBayArea* location—does not have the belief *(current.location = SEQBayArea*—the simulation engine will make the workframe available for execution. If there is no other workframe to execute first, it executes (fires) the *MovingToSEQBay* workframe.

Table 2. Open SEQ Bay Door Activity

LMP		CDR	
Communicate Ready To Offload		**Watching Opening SEQ Bay Door**	
Activity	**Communication**	**Communication**	**Activity**
Communicate Open Door	116:31:34 Bean: Okay. And we'll off-load the ALSEP. (Garbled).		*Watch Opening SEQ Bay Door*
Inspect SEQ Bay		116:31:39 Conrad: Nope. (Pause)	
	116:31:42 Bean: We ought to be able to move out with this thing.		
		116:31:44 Conrad: Okay.	
	116:31:48 Bean: The experiment bay looks real good.		
		116:31:49 Conrad: Yup.	
Raising SEQ Bay Door			
Activity	**Communication**		
Grab Lanyard Ribbons	116:31:50 Bean: The LM exterior looks beautiful the whole way around. Real good shape. Not a lot that doesn't look the way it did the day we launched it.		
Walk Back To Pull Ribbons Tight			
Pull Lanyard Ribbons		116:32:02 Conrad: (Possibly pulling a lanyard to open the SEQ bay doors) Light one. (Pause)	
	116:32:12 Bean: Okay. Here we go, Pete. Ohhhhh, up they go, babes. One ALSEP. (Pause)		
		116:32:22 Conrad: There it is.	

This means that the agent starts the *Move* activity. The agent performs this activity for as long as it takes to move from the location it is in to the *SEQBayArea* location. At the end of the *Move* activity the agent is in the *SEQBayArea* location. The agent notices the other objects and agents, such as the *SEQBay* object, and it will get beliefs about their location. New beliefs can trigger more workframes to be fired. The agent continues to perform activities, based on changes in its beliefs over time.

We can represent the relationship between workframes executing activities, containing other workframes that execute activities, etc, in a workframe-activity subsumption hierarchy as shown in Figure 4. Only one activity can be active at any given time (i.e. at any clock-tick), consequently only one workframe is "being worked on" at any given time. The order in which workframes at the same level in the hierarchy fire depends on two things: first, the conditions of the workframe that are matched to the beliefs of the agent, and second, the priority of the activities within the workframes. There are a number of important language constructs (such as detectables, consequences, and thoughtframes) we are leaving out from the discussion in order to keep the length of the paper within the necessary limits. For a more detailed description of the Brahms language we refer the reader to (van Hoof and Sierhuis 2000).

10. CONCLUSION

In this paper we have shown how the Brahms multi-agent modeling and simulation environment is used to model an activity of the Apollo 12 astronauts. We argue that in order to model and simulate the work practices of people we need to include several aspects of human behavior other than the reasoning behavior.

To simulate the work practice of humans, we need to represent the behavior of people at the *work practice level*. At this level we represent the way human agents act, react and interact with each other and their environment. The most important concept at this level is the situated *activity* that takes time and is constrained by the agent's beliefs about the specific situation.

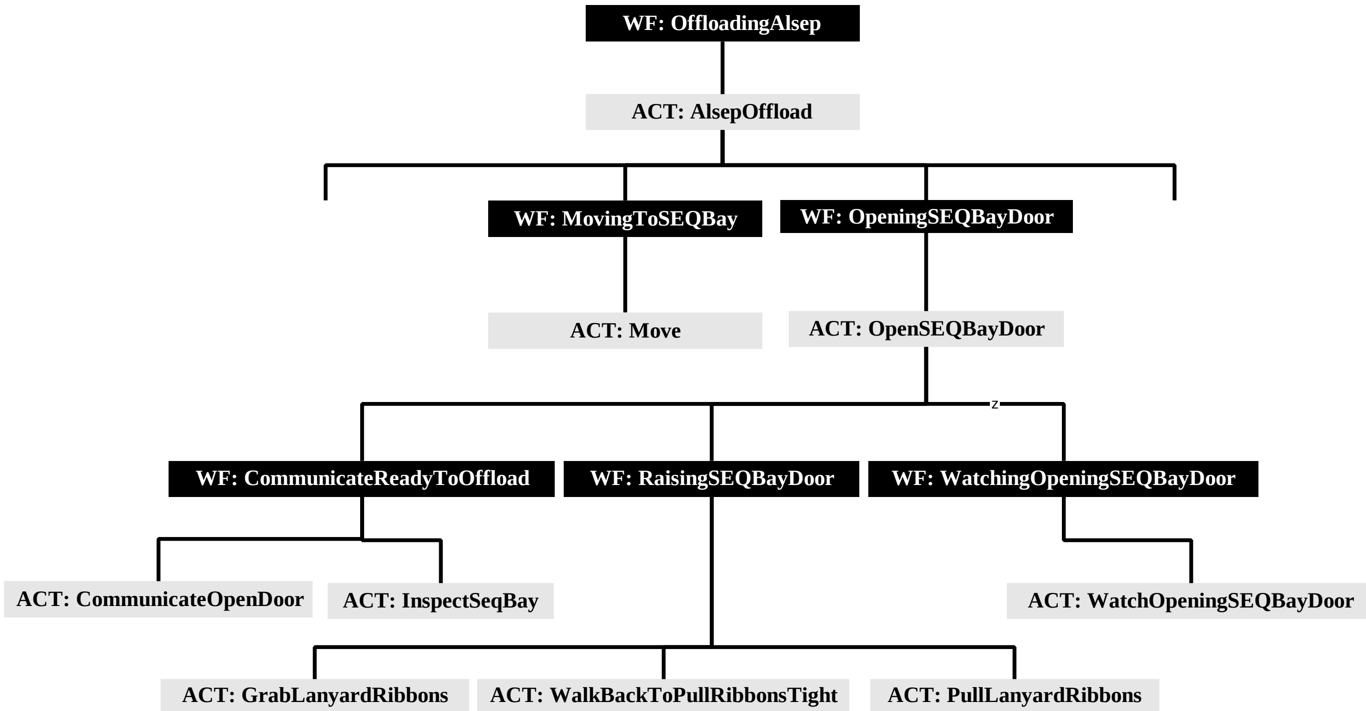

Figure 4. AlsepOffload Workframe-Activity Subsumption Hierarchy

11. REFERENCES

Agre, P. E. (1995). “Computational research on interaction and agency.” *Artificial Intelligence*, 72(1-2), 1-52.

Anderson, J. R., and Lebiere, C. (1998). “The atomic components of thought.” , Lawrence Erlbaum Associates, Mahwah, NJ.

Brooks, R., A. (1991). “Intelligence without representation.” *Artificial Intelligence*, 47, 139-159.

Clancey, W. J. (1997). “The Conceptual Nature of Knowledge, Situations, and Activity.” Human and Machine Expertise in Context, P. Feltovich, R. Hoffman, and K. Ford, eds., The AAAI Press, Menlo Park, CA, 247-291.

Clancey, W. J., Sachs, P., Sierhuis, M., and van Hoof, R. (1998). “Brahms: Simulating practice for work systems design.” *International Journal on Human-Computer Studies*, 49, 831-865.

Ehn, P. (1988). *Work-Oriented Design of Computer Artifacts*, Arbetslivcentrum, Stockholm, Sweden.

Emery, F. E., and Trist, E. L. (1960). “Socio-Technical Systems.” Management Sciences, Models and Techniques, C. W. C. a. o. (eds.), ed., Pergamon, London.

Greenbaum, J., and Kyng, M. (1991). “Design at Work: Cooperative design of computer systems.” , Lawrence Erlbaum, Hillsdale, NJ.

Hammer, M., and Champy, J. (1993). *Re-engineering the Corporation*, Harper Collins Publishers, Inc., New York, NY.

Jones, E., M. (1997). “The Apollo Lunar Surface Journal.” *URL: http://www.hq.nasa.gov/office/pao/History/alsj/*, National Aeronautics and Space Administration.

Kirsh, D. (1995). “The intelligent use of space.” *Artificial Intelligence*, 73(1-2), 31-68.

Nakashima, H., Noda, L., and Handa, K. “Organic Programming language GAEA for multi-agents.” *Proceedings of the Second International Conference on Multi-Agent Systems*, Kyoto, Japan, pp. 236-243.

Newell, A. (1990). *Unified theories of cognition*, Harvard University Press, Cambridge, MA.

Pasmore, W. A. (1993). *Creating Strategic Change: Designing the Flexible, High Performing Organization*, Wiley, New York.

Polanyi, M. (1983). *The Tacit Dimension*, Peter Smith, Magnolia, MA.

Searle, J., R. (1969). *Speech Acts*, Cambridge University Press, Cambridge, UK.

Suchman, L. A. (1987). *Plans and Situated Action: The Problem of Human Machine Communication*, Cambridge University Press, Cambridge, MA.

van Hoof, R., and Sierhuis, M. (2000). “Brahms Language Reference.” , http://www.agentisolutions.com.

Vygotsky, L. S. (1978, Originally published in Rusian in 1934). *Mind in Society: The Development of Higher Psychological Processes*, Harvard University Press, Cambridge, MA.

Weisbord, M. R. (1987). *Productive Workplaces: Organizing and Managing for Dignity, Meaning, and Community*, Jossey-Bass Inc., Publishers, San Francisco, CA.

Wenger, E. (1997). *Communities of Practice; Learning, meaning, and identity*, Cambridge University Press.