Modeling and Simulating Human Activity

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

1. NASA/Ames Research Center
Mail Stop 269-3
Moffett Field, CA 90435-1000
{msierhuis, bclancey, rvhoof}@mail.arc.nasa.gov

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

Abstract

This paper describes an approach for modeling and simulating the activities of humans and systems in organizations we refer to as work practice modeling. We describe a simulation experiment of the work practice of the Apollo 12 astronauts during the ALSEP offload activity.

Introduction

In this paper, we report on the results of an experiment using a approach for modeling and simulating the activities of humans and systems in organizations. The understanding of work processes and workflow has become more and more an integral part in the development of support technologies for organizations. However, the common modeling approaches used for modeling work processes and workflow lack, in a fundamental way, the ability to include the intricacies of the actual practices of people and systems. In the Work Systems Design and Evaluation group at NASA Ames Research Center, we are developing theories and tools for modeling the practices in a work system. In particular, we have developed a new multi-agent modeling and simulation language and environment—Brahms[1]—to study these theories (Clancey et al. 1998).

Goals and Objectives

The goal of the experiment was to investigate the use of the Brahms-language in order to *describe an existing work practice*. The challenge we faced in this experiment was to investigate if our theory of modeling work practice, as implemented in the Brahms language, would be sufficient to describe the work practice in the chosen domain. The objectives of this first experiment 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.

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.

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).

A work practice is defined as the collective activities of a group of people who collaborate and communicate, while performing these activities synchronously or asynchronously. Most often, people view work merely as

[1] Business Redesign Agent-based Holistic Modeling System

the process of transforming input to output, i.e. a Tayloristic view of work. For example, when building a house the input and output of the work is well defined. Sometimes however, it is more difficult to describe the input and output of the work. For example, consider a soccer match between two professional soccer teams. It is difficult to define the input and output of this type of work, although most of us would agree that professional soccer players are working. To describe the work of a soccer team we quickly fall into descriptions about teamwork and collaboration on the field.

We claim that the individual activities that make up the work not only have to do with the transformation of input to output, but more importantly with the collaboration between individuals in action, in pursuit of a goal. Imagine soccer players who collaborate in their activities of kicking a soccer ball, in pursuit of scoring a goal. Just focussing on the in- and output of each individual activity of a soccer player would not only be very difficult, if not impossible, it would also miss the opportunity to understand what is really going on in this work. However, 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).

We take a different view, and 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 collaborative activities are performed, and identify the effects of these activities, as well as to understand the reasons why these activities occur in the way they do. Therefore, 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 that everyone understands. However, it is difficult to describe what a practice is. People notice when something is not a practice, and can often describe why. Although it can be said that a group of people has developed a practice, when asked to describe what that practice is, and what it consists of, we find it difficult to describe in words. As such, practice is part of our tacit knowledge (Polanyi 1983).

An ad hoc definition of the word 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 be able to describe practice at an epistemological level we call the *work practice level*. In the rest of this paper, we will discuss a representational language to represent models of work practice. These models can be simulated in order to show the effects of the activities of people and their communication, being situated in a geographical environment, and using tools and artifacts to perform their collaborative work.

Theory of Modeling Work Practice

We briefly describe our theory of modeling work practice. Representing how people do work can be done at many different levels. In the knowledge engineering and AI world, 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. In workflow models we 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 work gets done.

Work practice includes those aspects of the work process that make people behave a certain way in a specific situation, and at a specific moment in time. To describe people's situation-specific behavior we need to include those aspects of the situation that explain the influence on 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.

Activity Behavior

People's behaviors are determined by the "execution" of specific activities at certain moments. This means that 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 situational context that a person or system is in. Agents' behaviors are organized into activities, inherited from groups to which

agents belong. 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 a being on a business trip. When the phone rings in our hotel room, we get up and walk over to pick up the phone. This means that we interrupt the activity of reading our book, and start the activity of answering the phone. In a sense, we actually 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).

Context

People act based on the situation they are in. With this we mean that people behave based on their beliefs about what they experience (infer or detect) their context to be. Therefore, 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 objects and artifacts that 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 facts. 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. Implicit in the above example is the fact that people and objects are always located. Moving from one location to another is an activity that takes time.

Communication

In order for two or more people to collaborate they need to communicate. In the Speech Act theory of Searle 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 collaboration activity between sender and receiver, as well as the intention and meaning of the speech act. 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 speaker and the hearer. While this is important, it doesn't, for instance, take into account that a communication activity between two people works or does not work due to 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 collaboration is very much defined by 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.

Communities of Practice

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

Apollo 12 and the ALSEP Offload

One of the biggest objectives of the Apollo 12 mission was to deploy the Apollo Lunar Surface Experiments 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.*

Figure 1 shows the plan and start-time for the Apollo 12 ALSEP Offload.

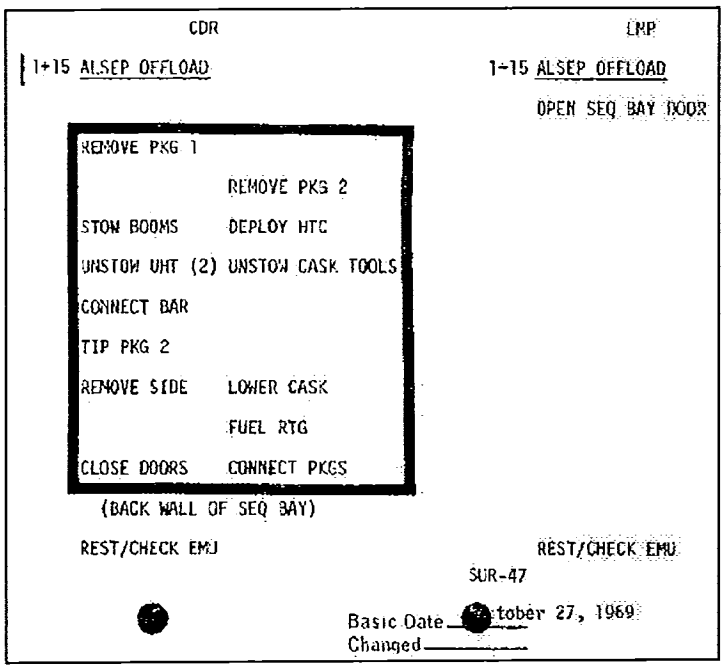

Figure 1. Apollo 12 Surface Checklist 47 for the ALSEP Offload

Even though this high-level task was planned and choreographed up front, the plan did not include the situational variations, the actual communication and collaborative activities between the astronauts, and the communication between and coordination of activities by the Manned Spaceflight Center (MSC) in Houston.

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 we discussed in the previous section, what constitutes the practice of the ALSEP Offload is the way the actual 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. It is situated action (Suchman 1987). A choreographed play "executed" during the performance, planned and trained, but always different.

In the next sections, we will describe how the ALSEP Offload work is modeled in a model of work practice. The model is not a model of the problem-solving knowledge of each individual involved in this task. Instead, it is a model of the behavior of the individuals. It describes how the collaboration, coordination, and communication between the three individuals happen, and make this a fluent event.

Agent Model

One of the most relevant design issues 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.

Designing an Agent Model is similar to the design of an Object Model in object-oriented design (Rumbaugh et al. 1998). Just as the class-hierarchy in an Object Model, we need to design the group-hierarchy in the Agent Model. As a rule of thumb, we identify the communities of practice 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 2 shows the Agent Model design. We start with defining our agents. Each agent represents a person in our domain, e.g. 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. 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. This way, the group represents the community of agents that can perform the ALSEP Offload task. For Apollo 12, both the CDR and the LMP trained for the ALSEP Offload activities, and both of them were able to perform all the ALSEP Offload tasks, and therefore belong to the group ALSEPOffloadGroup. Thus, the Commander and LunarModulePilot groups are members of the group AlsepOffloadGroup. Since both the CDR and the LMP were working on the surface there are tasks that both astronauts needed and/or could perform. As said before, the ALSEP Offload task was one of them, but there were others as well. All the activities that needed to be performed by all astronauts on the lunar surface are represented in the LunarSurfaceAstronaut group. Such activities include taking photographs and changing the cooling of their space suit.

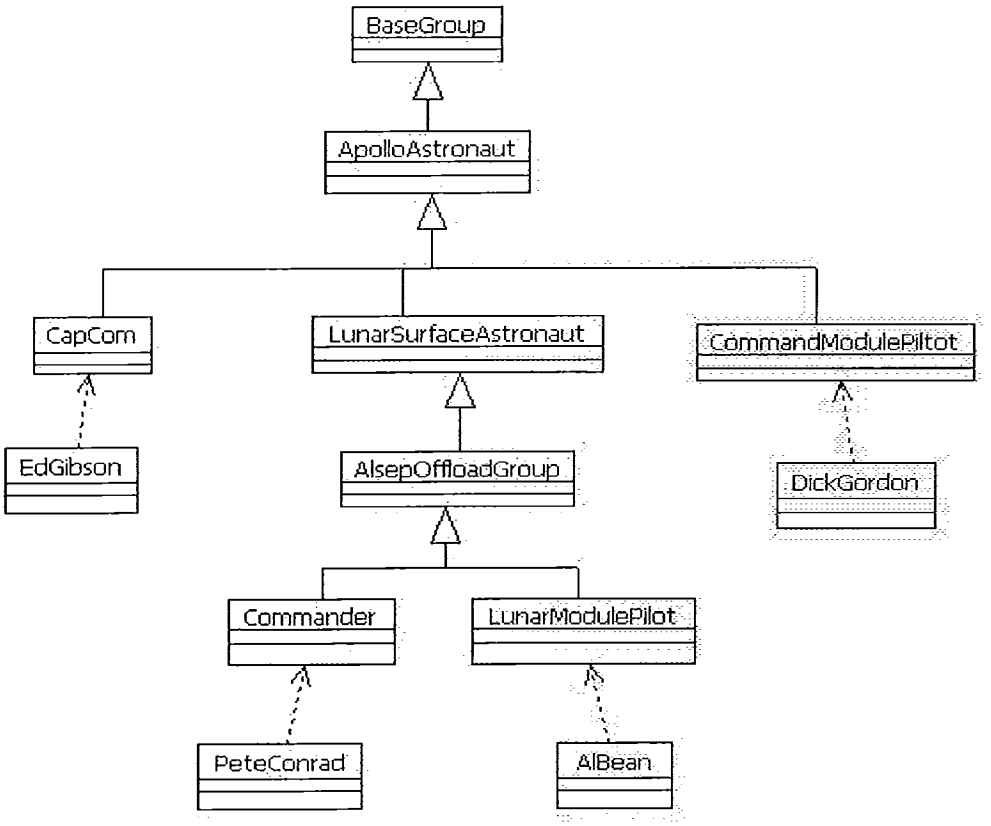

Figure 2. Apollo Agent Model Design

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 3 shows the Object Model design for the Apollo 12 domain objects and artifacts. As with the Agent Model, the root-class of the class hierarchy is the class *BaseClass*. All other classes and objects inherit from this BaseClass class.

Table 1 shows the Brahms model source code for the LM and SEQBay objects. Both the LM and SEQBay objects are instances of BaseClass.

Besides representing the corresponding artifacts on the Apollo 12 mission, the source code also specifies the initial location of the object within the Geography Model (see next section). Both objects are located in the SEQBayArea area. Furthermore, the objects

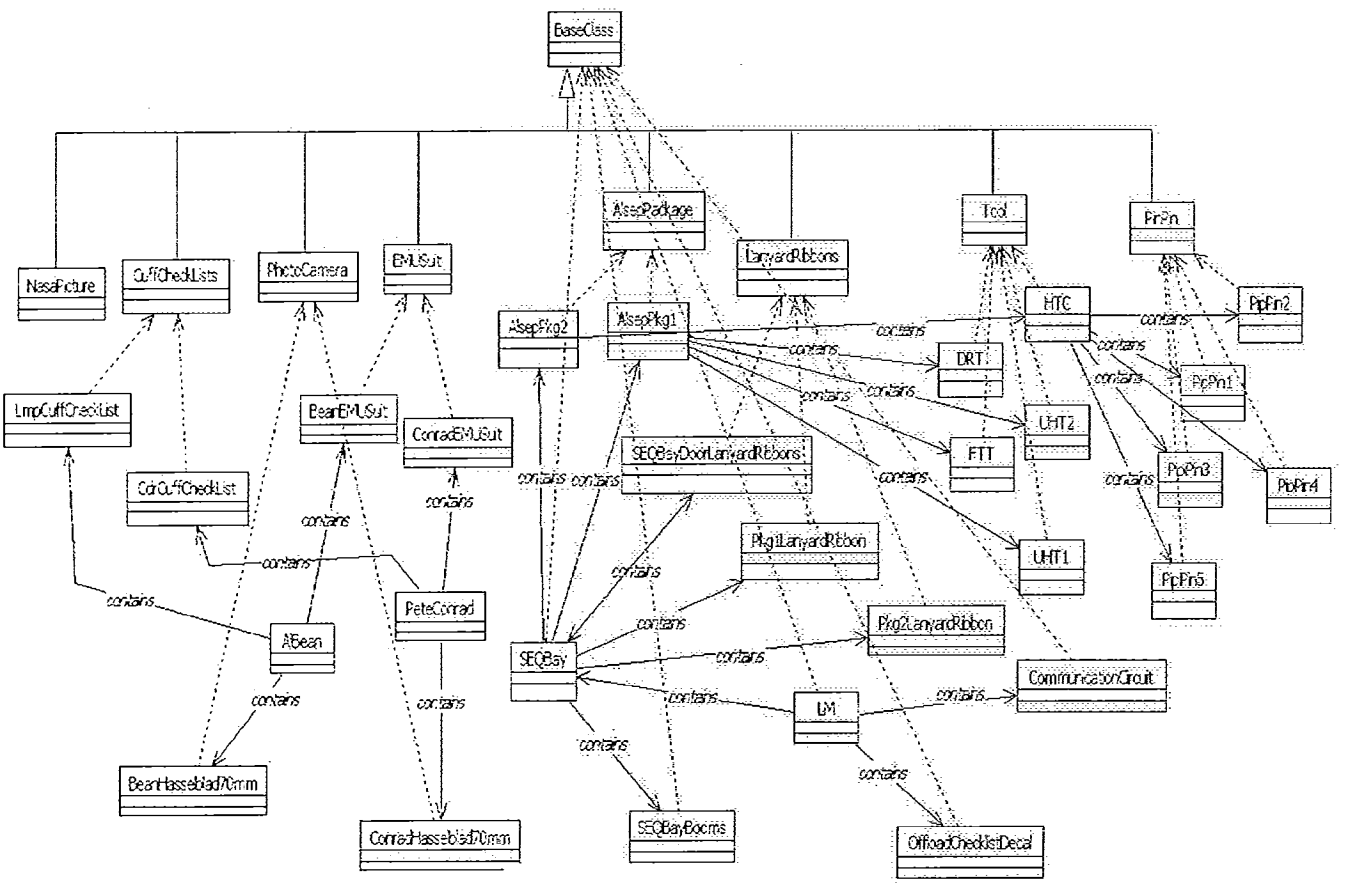

Figure 3. Apollo Object Model Design

declare the attributes with which we describe the different aspects of these objects. Although we could describe any number of aspects of an object, such as the color, height, etc, we only declare those attributes that are relevant. To model the fact that the astronauts inspect the LM and the SEQ Bay's exterior appearance after the landing, we declare the attribute *exteriorAppearance* as a type symbol attribute. Using this attribute we can represent the state of the exterior of these objects (Agre 1995).

Both the LM and the SEQBay objects have a fact describing the state of their exterior appearance after the landing on the moon as an initial fact for the simulation, e.g.

*(current.**exteriorAppearance** = LooksGood)*

```
// Apollo 12 objects
object LM instanceof BaseClass {
        display: "Intrepid";
        location: SEQBayArea;
        attributes:
                public symbol exteriorAppearance;
        initial_facts:
                (current.exteriorAppearance = LooksGood);
                (current contains SEQBay);
}

object SEQBay instanceof BaseClass {
        location: SEQBayArea;
        attributes:
                public symbol door;
                public symbol exteriorAppearance;
        initial_facts:
                (current.exteriorAppearance = LooksGood);
                (current.door = closed);
                (current contains AlsepPkg1);
                (current contains AlsepPkg2);
                (current contains OffloadChecklistDecal);
                (current contains SEQBayDoorLanyardRibbons);
                (current contains Pkg1LanyardRibbons);
                (current contains Pkg2LanyardRibbons);
                (current contains SEQBayBooms);
}
```

Table 1. Apollo 12 LM and SEQ Bay Brahms objects

The status of the door of the SEQBay is modeled with the *door* attribute of type symbol. The door is in the initial state (i.e. an initial fact) of being closed, e.g.

*(current.**door** = closed)*

This represents the door of the SEQ Bay being closed at the start of the ALSEP offload. Next, we model the objects that are located within the LM and SEQ Bay. This is represented with the contains relation (see Figure 3). This relation is declared in the BaseClass class, and inherited by the LM and SEQBay objects. The fact that the SEQBay is located on the outside of the LM is represented as an initial fact in the LM object, i.e.

*(current **contains** SEQBay).*

Now that the agents and artifacts are represented, the next section describes the geography model in which the agents and artifacts are located during the simulation.

Geography Model

In Brahms we model geographical locations using two concepts, area-definitions and areas. Area-definitions are user-defined types of areas. Areas are instances of area-definitions. Thus an area is an instance of 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 important; Earth, the Manned Spaceflight Center[2] (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 4 shows the geography model design.

Initial Locations

Each agent and object has an initial location in one of the lowest-level areas, (CommandModule, AwayFromTheSEQBayArea, AlsepDeploymentArea, LandingSite, SEQBayArea, or MissionControlCenter). Initial locations are locations in which an agent or object is placed during the initialization phase of the simulation. This way each agent and object starts out being located in a geographical location (an area). To define an initial location for an agent or object the modeler uses the location attribute (see Table 1).

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 both agents arrive at their new location area they will immediately detect facts about the location of other agents and objects that are also in the area they arrive at. The simulation engine automatically creates beliefs for the agent from the facts about other objects and agents that are in the same location. The agents already in that location will get the belief that the agent that arrived is now also in

[2] During the Apollo days the NASA center in Houston was called the Manned Spaceflight Center (MSC). Today it is referred to as Johnson Space Center (JSC).

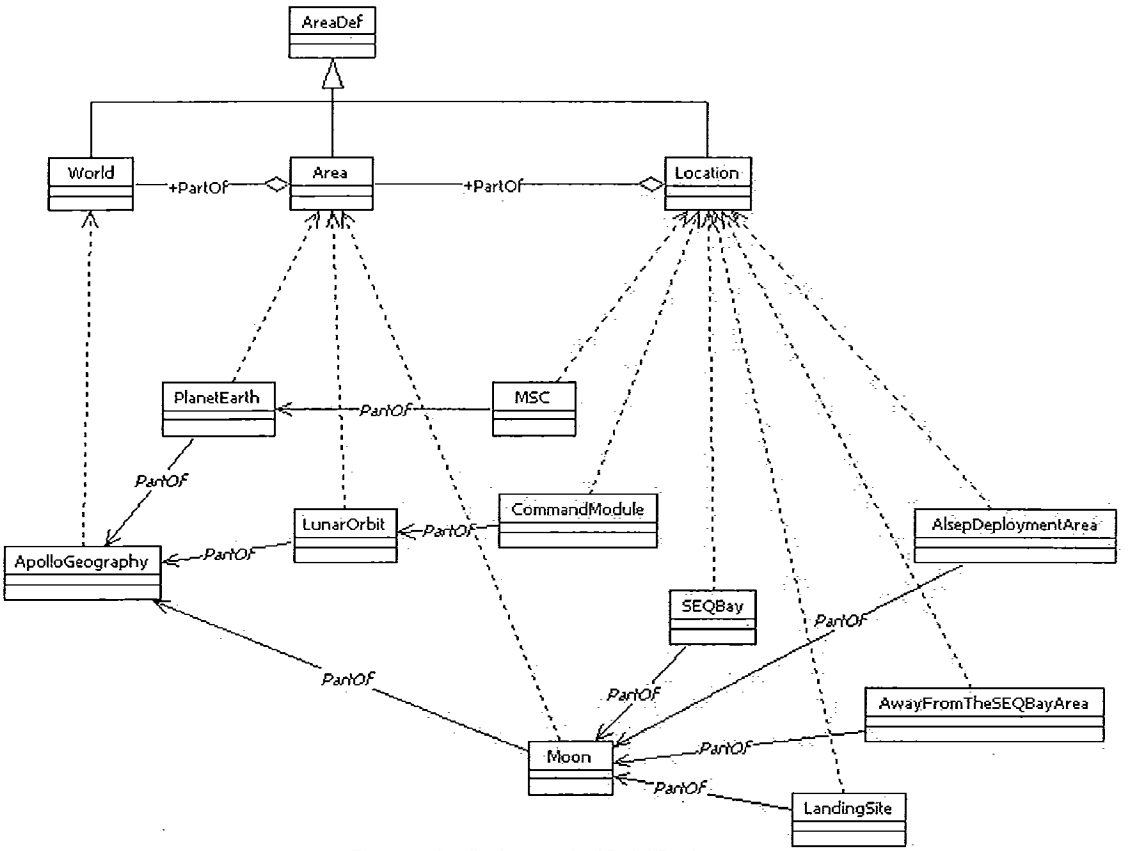

Figure 4. Apollo Geography Model Design

the location. This way, agents will always notice other agents and objects that are in the location the same area.

Activity Model

ACT-R and Soar are production systems that all use some form of goal-structure as part of their architecture. Anderson and Lebiere reflect that the lesson that has been learned in properly modeling human cognition is that goal-structures are necessary to represent the system's "current purpose" and to organize its behavior. (Anderson and Lebiere 1998, p. 39).

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. 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 like a function a person performs over a period of time. An activity execution has a well-defined beginning and end, and thus takes an amount of time. This is different than the concept of a goal. A goal is not a function. Cognitive scientists often talk about a "goal state." The use of the notion of a "goal state" suggests that it means that when a person or system has reached the "goal state," it is done with achieving the goal. However, in cognitive science it most often means that the person or system is in the *state of pursuing a goal*, and thus has not reach the goal yet. Pursuing a goal takes time and effort. This is what we call *being in the activity of pursuing the goal.*

Although the definition of an activity looks similar to that of a 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 still in the mitts of performing the higher-level activity. Each higher-level activity subsumes all its lower-level activities. This has led us to implement the activity-structure of an agent in an agent-subsumption architecture (Brooks 1991).

In this 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.

There are three separate areas where people are located during the Apollo ALSEP Offload activity. The CapCom is located at MSC. His main function is to listen to and communicate directly over the voice-loop with the astronauts. The CDR and LMP are the astronauts on the lunar surface and are located at or near the area of the SEQ Bay, which is located on the backside of the Lunar Module

(LM) "Intrepid". The CMP is orbiting around the moon in the Command Module (CM) "Yankee Clipper."

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 LSJ. The actual names of the activities and sub-activities are more or less arbitrary, and conceptualize the modeler's *interpretation* of the observations of the Apollo 12 communication data and the Apollo 14 video data. However, these data and observations are strong evidence that these are the actual activities that are performed during the Open SEQ Bay Door activity.

The activities from Table 2 are implemented in the Brahms model as the *OpenSEQBayDoor* composite-activity. Figure 5 shows this activity, its sub-activities and workframes.

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.

As the first activity during the ALSEP offload, the CDR and LMP start walking to the area of the SEQ Bay. Walking to the SEQ Bay area to start opening the SEQ Bay door is modeled by the Move activity, as can be seen at the top of Figure 5.

The first workframe to fire is the *OffloadingAlsep* workframe, which executes the *AlsepOffload* activity. Executing the *AlsepOffload* activity enables all the workframes in it to potentially fire for the agent. Each of the workframes will execute lower-level activities, which are said to be subsumed by the higher-level 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.	

Table 2. Open SEQ Bay Door Activity

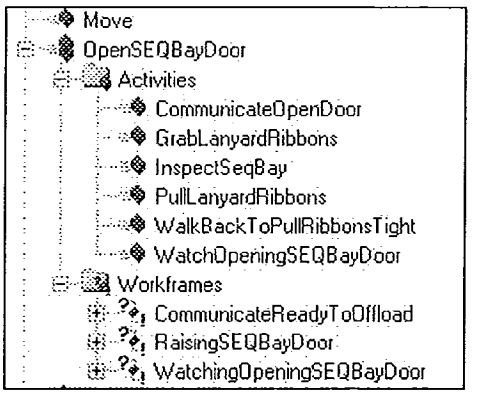

Figure 5 The OpenSEQ BayDoor composite-activity, sub-activities, and workframes

We can represent the relationship between workframes executing activities, containing other workframes that execute activities, etc, in a workframe-activitiy subsumption hierachy as shown in Figure 6.

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. This means that 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 to be 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).

Viewing the simulation results

Figure 7, shows the AlsepOffload activity performed by the AlBean and PeteConrad agents, as well as the communication between the two agents. While performing the AlsepOffload composite activity, both agents are within the OpenSEQBayDoor activity. While AlBean is performing the activities within the CommunicateReady- and the RaisingSEQBayDoor workframe, the PeteConrad agent is performing the activities within the WatchingOpenSEQBayDoor workframe. The grain-size of the simulation is one second. This means that the simulation engine updates every agent and object every second. We can therefore say that the simulation is a second by second model of the work practice of the lunar surface astronauts. Figure 7 also shows the location the agent when performing the activity. As an overlay, the (blue) dotted arrows show the communication of beliefs between agents AlBean and PeteConrad. The direction of the arrows show the direction in which the beliefs are being communicated, while the little square blue box at the start of the arrow shows the agent that is performing the communication.

Figure 7 is a screen shot from the AgentViewer application[3]. The AgentViewer application takes as input a Brahms Simulation History database[4]. This history database contains the historical situation-specific model data of a particular simulation run. The AgentViewer

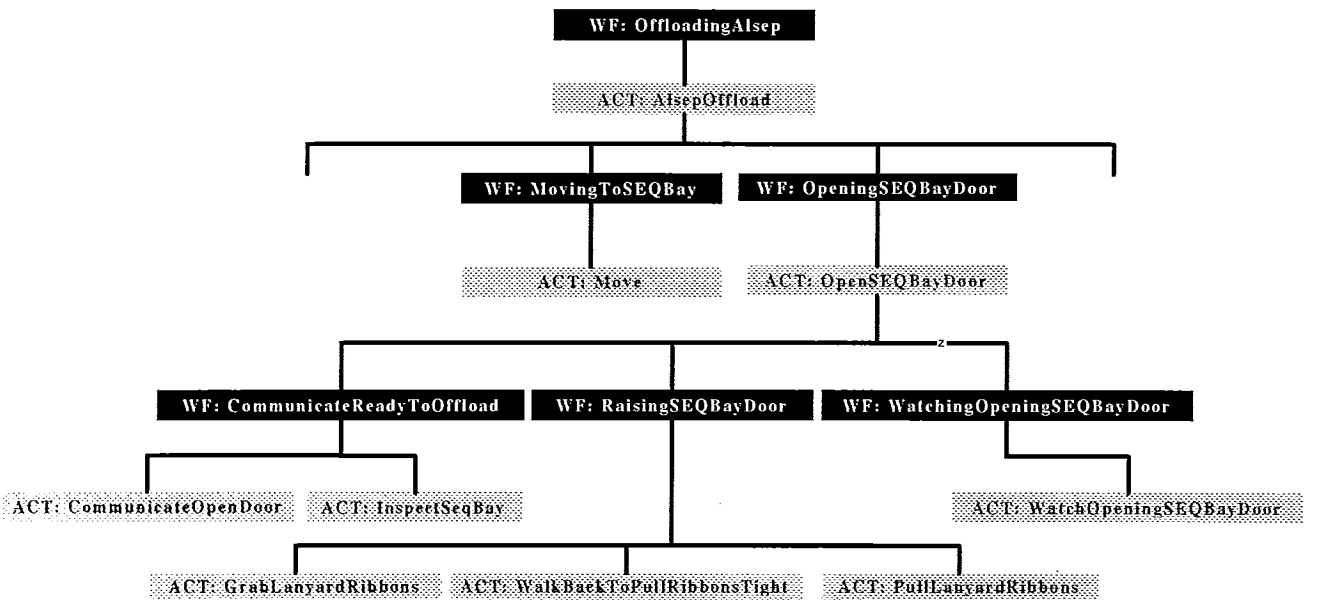

Figure 6. AlsepOffload Workframe-Activity Subsumption Hierarchy

[3] The AgentViewer application is a stand-alone Visual Basic application we developed for viewing the results of a simulation.

[4] The history database is a complex relational database containing the simulation data preserving their relationships.

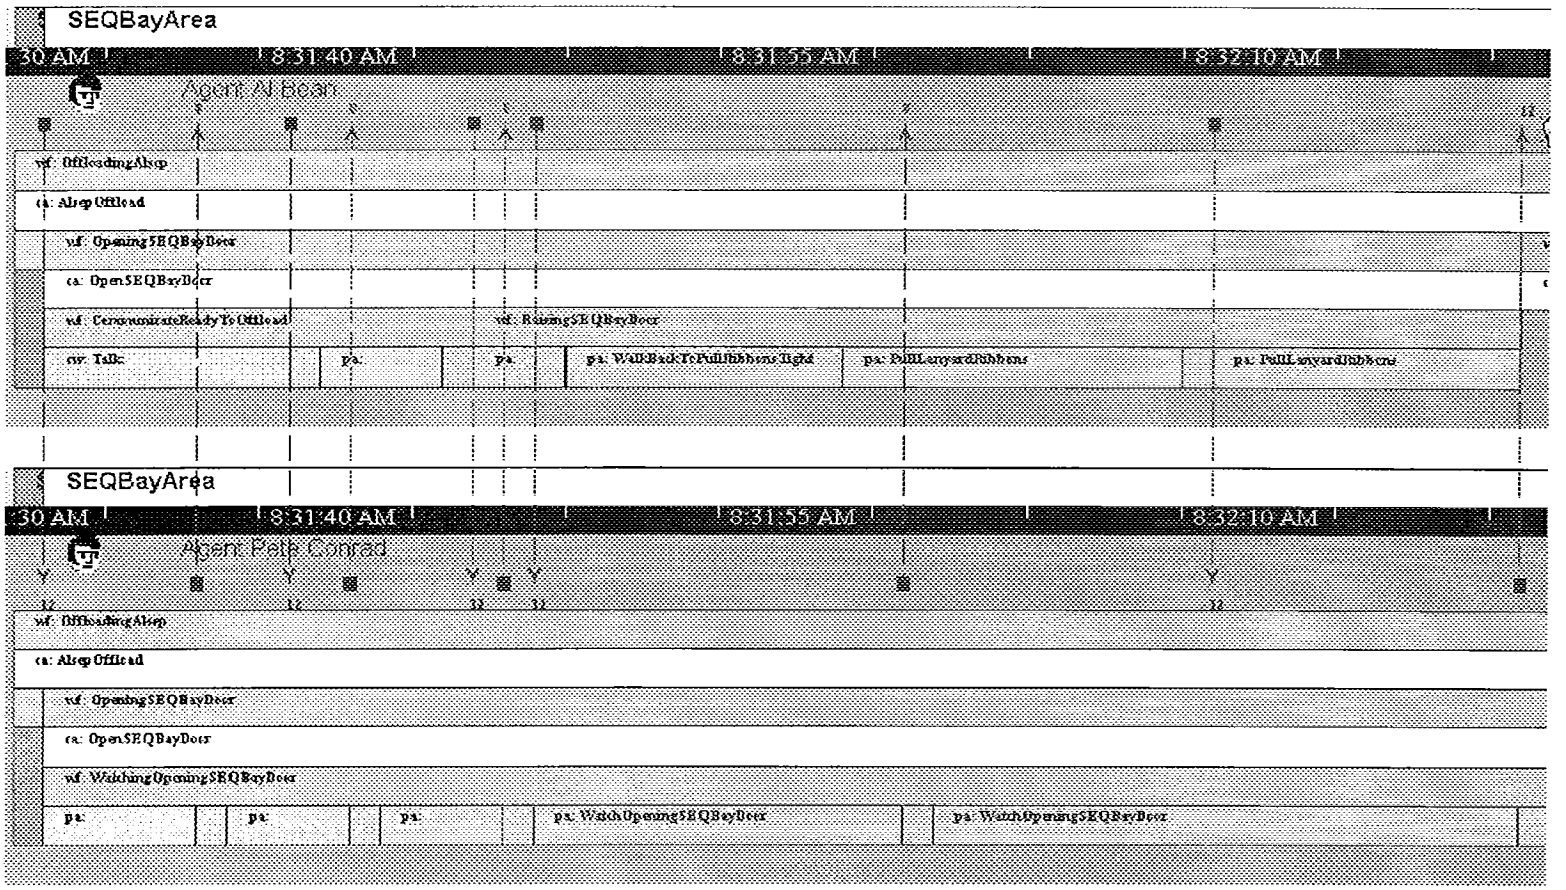

Figure 7. AlsepOffload Activity Agent Timeline

application creates a graphical representation of the activity of agents and objects during a simulation.

Discussion

In this paper we have shown how the Brahms multi-agent modeling and simulation environment is used to simulate the situated activity behavior of human actors and artifacts in a real-life setting. The experiment shows 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.

We argue that a model of human behavior in the real world needs to include aspects that can explain the way people actually behave, including collaboration, communication, and interaction with the environment.

We have shown that 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. Application of problem-solving knowledge occurs within a situated activity, and is therefore constrained by the situation in which the agent performs the activity.

Describing the behavior in terms of what actually happens in the world does not necessarily lead to a description of the individual's problem-solving behavior. Rather, it leads to a higher-level description of the emergent total system behavior, in terms of the individual interactions and responses to the other elements in the system (people and artifacts), as well as the emergent sequence of individual activity.

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.

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 Fliexible, High Performing Organization*, Wiley, New York.

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

Rumbaugh, J., Jacobson, I., and Booch, G. (1998). *The Unified Modeling Language Reference Manual*, Addison-Wesley.

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.