

Cognitive Modeling of Social Behaviors

William J. Clancey, Maarten Sierhuis, Bruce Damer, and Boris Brodsky

To appear in R. Sun, *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. New York: Cambridge University Press.

1. Introduction

The driving theme of cognitive modeling for many decades has been that knowledge affects how and which goals are accomplished by an intelligent being (Newell 1991). But when one examines groups of people living and working together, one is forced to recognize that whose knowledge is called into play, at a particular time and location, directly affects what the group accomplishes. Indeed, constraints on participation, including roles, procedures, and norms, affect whether an individual is able to act at all (Lave & Wenger 1991; Jordan 1992; Scribner & Sachs 1991).

To understand both individual cognition and collective activity, perhaps the greatest opportunity today is to integrate the cognitive modeling approach (which stresses how beliefs are formed and drive behavior) with social studies (which stress how relationships and informal practices drive behavior). The crucial insight is that norms are conceptualized in the individual mind as ways of carrying out activities (Clancey 1997a, 2002b). This requires for the psychologist a shift from modeling only goals and tasks — why people do what they do—to modeling behavioral patterns—what people do—as they are engaged in purposeful activities. Instead of a model that exclusively deduces actions from goals, behaviors are also, if not primarily, driven by broader patterns of chronological and located activities (akin to scripts).

This analysis is particularly inspired by activity theory (Leont'ev 1979). While acknowledging that knowledge (relating goals and operations) is fundamental for intelligent behavior, activity theory claims that a broader driver is the person's motives and conceptualization of activities. Such understanding of human interaction is normative (i.e., viewed with respect to social standards), affecting how knowledge is called into

play and applied in practice. Put another way, how problems are discovered and framed, what methods are chosen, and indeed who even cares or has the authority to act, are all constrained by norms, which are conceived and enacted by individuals.

Of special interest for the cognitive modeler, and emphasized in social theory (Lave 1988), is how norms are reinforced and shaped through behavior. Each enacting of a norm potentially reinforces the behavior pattern for the individual, as well as the group observing and relating to the behavior. But also, each action potentially changes the norm, including functional adaptations to the current circumstances as well as personal whim. One might refer to understanding of norms as an individual's "social knowledge"; but many or perhaps most norms are tacit—the patterns are not necessarily experienced or described. Of major interest for cognitive modeling is how individuals formulate situation-action rules of behavior (i.e., they develop models of norms) to deliberately accomplish goals in novel ways (i.e., they deduce how to relate and adapt available methods to permissible behaviors). For example, a leader may develop the group's capability by humorously violating a norm, reinforcing each individual's understanding of the group's structure and ways of interacting.

Our understanding of how to relate goals, knowledge, behaviors, and social concepts in a cognitive model has been developing over more than a decade in the Brahms modeling and simulation system (Clancey et al. 1998, 2002b; Sierhuis 2001). It has taken a long time to break out of the task analysis perspective to understanding the social notion of activity (Lave 1988; Suchman 1987) as a behavioral and not functional description, and to ground it in a cognitive architecture. The significant breakthroughs included:

- Understanding activities as patterns of what people do, when, and where, using what tools or representations;
- Representing activities in a cognitive model using a subsumption architecture (i.e., conceptualization of activities occurs simultaneously on multiple levels);
- Understanding that conceptualization of activities is tantamount to conceptualization of identity, "What I'm doing now," which is the missing link between psychological and social theory (Clancey 1997b, 1999; Wenger 1998).

- Simulating collective behavior in a multi-agent simulation with an explicit “geographic model” of places and facilities, using the Brahms tool.

A Brahms model is a way of formalizing (expressing, collecting, and organizing) field observations so they can be correlated, shared, and used in work system design (Sierhuis & Clancey 2002; Sierhuis et al. 2003; Seah, Sierhuis & Clancey 2005). The primary objective is not necessarily to construct a predictive model of human behavior, which is often emphasized in scientific modeling, including cognitive modeling, but to have a systematic way of relating disparate sources of information, including video, notes, and surveys.

To illustrate these ideas, this chapter presents an extract from a Brahms simulation of the Flashline Mars Arctic Research Station (FMARS), in which a crew of six people are living and working for a week, physically simulating a Mars surface mission (Clancey 2002a). This Brahms simulation of this mission is broadly described in Clancey (2002b); this chapter focuses on one part, the Brahms simulation of a planning meeting. How people behave during the meeting (e.g., standing at the table) exemplifies the nature of norms; this is modeled at the individual agent level in Brahms. The example shows how physiological constraints (e.g., hunger, fatigue), facilities (e.g., the habitat’s layout), and high-level events during the meeting interact. This chapter describes the methodology for constructing such a model of practice, from video and first-hand observation, and how this modeling approach fundamentally changes how one relates goals, knowledge, and cognitive architecture.

Relating physical behaviors to a meeting and producing a visual display with realistic timing involves integrating diverse information (topography, agent beliefs, posture, meeting structure). No attempt is made here to analyze or model the group dynamics of decision making in detail (e.g., raising one’s voice, misunderstandings, domination, digressions). Rather the effort here is intended to provide a framework within which such analysis could be meaningfully embedded. Specifically, we hypothesize that being able to model apparently superficial multi-agent behaviors, as we have here, is a necessary first step in understanding the cognitive and social nature of norms. Recognizing how norms are manifested, violated, adapted, etc. will enable us to subsequently use activity-based analysis to better analyze the quality of group decision making.

Following the analytic approach of Schön (1987), this research effort shifts from studying technical knowledge in isolation to modeling the context in which behavior occurs and how it unfolds over time through interactions of people, places, and tools. The resulting simulation model of practice is a powerful complement to task analysis and knowledge-based simulations of reasoning, with many practical applications for work system design, operations management, and training.

2. The Brahms Approach for Relating Cognitive and Social Processes

The Brahms simulation system was developed as a means of systematically relating information gained from the anthropological method of observing by participating in some activity, called “participant observation” (Spradley 1980; Clancey in preparation a, b). Being a participant allows the observer to detect and understand events that people would not otherwise report (e.g., a swimmer would probably not mention “you have to be in water; alive,” Wynn 1991, p. 49). Brahms’ patented design was conceived in 1992 to complement business process modeling tools by representing how work actually gets done. As a model of practice, in contrast with formal processes, Brahms simulations emphasize informal communications and assistance (i.e., actions that are not specified in task requirements or procedures), and circumstantial interactions (e.g., how placement of people and tools affects what information is shared or how long a job takes).

The Brahms modeling language enables representing and relating the following:

- people (as agents having beliefs, factual properties, and belonging to one or more groups)
- locations (as a hierarchy of geographic areas)
- tools and furniture (represented as objects having factual properties)
- computer systems (e.g., databases, represented as objects with stored beliefs that can be read or modified by agents)
- robotic systems (represented as agents)
- behavior of people and systems (represented as activities).

Activities are represented as prioritized situation-action rules called workframes and conditional inference rules called thoughtframes. Workframes have four parts:

- preconditions (matched against agent beliefs)

- actions (activities or primitive actions)
- detectables (conditions associated with actions, modeling perception of the world)
- consequences (changes to beliefs and the state of the world).

Primitive actions occur for a fixed duration (or the duration may be randomly generated from a specified interval). The simulation engine manages agent and object behaviors as a discrete event simulation.

The state of the world (physical properties of agents and objects) is modeled in Brahms as facts. Detectables match against facts, resulting in agent beliefs (which may be different from the facts), modeling how what is perceived is conditional on what an agent is doing. Changed beliefs then activate workframes for the activities in which the agent is currently engaged. Detectables may also abort or complete an activity. Thus, agent behaviors are largely data-driven within the context of activities. The language provides two special primitive actions: Move (to a specified location, taking a particular time) and Communicate (ask or tell another agent a belief matching a specified proposition, which applies as well to reading and writing beliefs to an object, e.g., a computer screen).

An agent is engaged in a hierarchy of activities at any particular time, constituting a subsumption architecture. For example, an FMARS crew member might be ParticipatingInPlanningMeeting while ConductingPlanningMeeting during the course of LivingOneDayinTheMarsHabitat. The agent is doing all of these activities at a particular moment, and thus a stack of activities is always active for every agent. The workframes and thoughtframes of these activities may activate, depending on the agents beliefs and the priorities of the workframes. Furthermore, any of the detectables on the current line of workframe activation may be triggered, according to the facts in the world that the agent encounters (subject to area and line of sight restrictions). Workframes may thus be interrupted or resumed as the agent behaves, gets new beliefs, and modifies the world. Furthermore, the initial beliefs of the agent, as well as the potential activities are inherited by group membership. Groups may belong to groups, providing an efficient way of representing beliefs and behaviors.

The Brahms language, architecture, and simulation engine are described in detail by Sierhuis (2001). Besides the original simulations of office work (Clancey et al. 1998) constructed for NYNEX (the former New York New England telephone company),

Brahms has been used to model NASA's mission operations, deployment of instruments on the lunar surface by Apollo astronauts (Sierhuis 2001), how procedures are followed on the International Space Station (Acquisiti et al. 2002), activities of scientists controlling the Mars Exploration Rovers (MER) (Seah, Sierhuis, & Clancey 2005), and teleoperations from earth of a proposed lunar rover (Sierhuis et al. 2003).

Before examining the Brahms model of the FMARS planning meeting, a few aspects of activity-based modeling should be emphasized:

- A model of activity is a model of practice, what people do. It should be contrasted with idealized or written models of procedures (what people are supposed to do).
- Tasks and activities are different units for viewing and describing human behavior. Like functional and behavioral models of artifacts such as electronic circuits, a task model can be related to, but does not strictly map onto an activity model. Most notably, many activities, such as eating, which can occur at any time during work, are omitted from task models. Simply put, a task model describes input and output relations as a kind of idealized specification of what should be accomplished. An activity model describes located, chronological behaviors and perceptual experiences. See Clancey (2002b) for extensive discussion and comparison of task analysis to Brahms, especially the historical relation to scripts.
- The emphasis on modeling behavior is not the same as behaviorism. Agent actions are totally driven by their perceptions, beliefs, and conceptualization of activities (represented by workframes and thoughtframes).
- Brahms activities are models of *conceptualizations*—which are largely non-verbal. Models of activities are quite different from the models of technical information and task-oriented procedures in knowledge-based systems.
- Brahms models are first and foremost investigators' models, not necessarily patterns articulated by the people being modeled. However, by incorporating agent beliefs (perhaps unarticulated), perception, conditional actions, and inferences, Brahms models have many characteristics of cognitive models.
- An agent's beliefs include how other people relate to activities, objects, and procedures, that is, social knowledge.
- Attitude, emotion, and personality are of fundamental importance in

understanding human activity, but are not included in the FMARS model. For example, the crew's attitude towards each other is revealed by their posture and spacing around the meeting table. These characteristics of people are essential for the application domain of long-duration space missions. In related work the FMARS data and simulation is being used to understand what aspects of personality for example are relevant in understanding the crew's behavior.

- Broadly speaking, a person's activities are identities. For example, one crew member was *simultaneously* being an American woman, a graduate student in geophysics at MIT, an FMARS crew member, and a person attending a planning meeting. These identities are dynamically composed and blended conceptions of "what I'm doing now," such that norms at each level are tacitly attended to and integrated (Clancey 1999; 2002b).
- Both formal structures (e.g., roles and procedures) and informal, emergent interactions (e.g., friendship) are part of the conceptualization of activity, but rules are always only consciously interpreted guides, not rigid controllers of behavior, as in computer programs.¹ Observing and documenting how preplanned procedures are adapted in practice is a central part of understanding the nature and role of cognition in the real world (Suchman 1987).

¹ In this form, situated cognition concerns the dynamic nature of human memory: Knowledge does not consist of stored structures such as rules and procedures that are indexed, retrieved, and subconsciously executed as in the von Neumann computer architecture. In general, social scientists promoting situated cognition in the 1980s did not present alternative neural arguments, and used sometimes confusing language (e.g., "The point is not so much that arrangements of knowledge in the head correspond in a complicated way to the social world outside the head, but that they are socially organized in such a fashion as to be indivisible," Lave 1988.) Some claims were absurdly interpreted by some researchers as "there is no knowledge in the head." For examples and discussion see Clancey (1993; 1994; 1995; especially, 1997b, "Remembering Controversies," chapter 3).

3. Simulation Model of Mars Crew Planning Meeting

Developing a Brahms model of a planning meeting exploited a unique opportunity and involved many steps:

- A crew of six people was living in the Mars analog mission for a week (at FMARS on Devon Island in the Canadian Arctic during July 2001).
- Clancey was selected to participate in the mission as a member of the crew (serving as journalist and meteorologist).
- The crew's activities were systematically observed and recorded.
- Time-lapse video was analyzed to map out patterns of what people did, when, and where.
- Selected multi-agent interactions were simulated (a planning meeting, filling the water tank, and preparing to work outside—an extra-vehicular activity, EVA).
- The Brahms simulation was integrated with a graphic rendering of agent postures, movements, object manipulations, etc. in the Brahms Virtual Environment (BrahmsVE) described in this chapter, implemented in Adobe Atmosphere™ (a commercially available, browser-based rendering engine).
- The simulation was refined by analyzing and further specifying the interaction of physiological, cognitive, and social structures (referring to the time-lapse video, photographs, and ethnographic field notes).

Over the course of a week, an FMARS participant observer can induce the typical pattern of the day, including what individuals do at different locations habitually. One approach is to keep an accumulating outline that is revised each day as part of the observer's field notes. The resulting Brahms model has a hierarchical activity structure, shown here chronologically:

LivingOneDayinTheMarsHabitat

 Sleeping

 GoingToRestroom

 MovingToArea

 GettingUp

 EatingBreakfast

HeatingWater
BringingBreakfast
DoingPersonalItemsAfterBreakfast
StartingPlanningMeeting
AnnouncingReadinessForPlanning
Gathering
ChattingBeforePlanning
AnnouncingStartOfPlanning
ConductingPlanningMeeting
ParticipatingInPlanningMeeting
CoveringAgendaItemWeather
CoveringAgendaItemWater
AnnouncingEndOfPlanningMeeting
ConductingEVAPreparation
DonningSuit
DepressurizingInChamber
ConductingEVA
EatingSnack
TakingNap

Many details in the model are omitted here, such as the steps in donning the suit and activities relating to specific roles and tasks (e.g., working with particular laboratory equipment).

The present model of the FMARS planning meeting does not attempt to replicate the conversational details of how people plan in a group by articulating and negotiating alternatives. As will become clear, there are many other issues to consider in simulating a planning meeting. The topics of the planning meeting, such as discussing the weather and reviewing the habitat's systems (power, water), are modeled as a sequence of events, with fixed durations. Even within such a restricted framework, individual agents can opportunistically change the topic (a subactivity) of the meeting or carry out a given subactivity in a way that changes what other agents are doing. For example, if there is a

fire alarm, the meeting will be interrupted and the activity of responding to the alarm would begin. This flexibility results from the combination of detectables, thoughtframes, communications, inheritance of activities through group membership, and the subsumption architecture for interrupting and resuming activities.

Subsequent sections explain in more detail how the planning meeting model is created and what its structure reveals about the relation of cognition and social behavior.

3.1 Planning Meeting Time Lapse

Using methods developed over several expedition field seasons (Clancey 2001), Clancey systematically recorded most of several days using a time-lapse apparatus. A quarter-frame (320x240 pixels) wide-angle view (Figure 1) was captured direct to computer disk every 3 seconds, such that the entire upper deck outside of the staterooms is visible. These frames were manually abstracted in a spreadsheet to show where people are (columns) at different times (rows).² From this, statistics and graphs are generated. Meetings such as the morning planning meeting are often video-recorded in full, so the conversations can also be analyzed.



² Foster-Miller, Inc. has been funded by NASA to develop the Crew Activity Analyzer, which uses image processing to automate most of the time-lapse analysis.

Figure 1. FMARS planning meeting of July 13, 2001, after KQ has moved from far left seat to standing on right. Commander sits on one long side of the table; Clancey is on the right. Ladder to lower deck is out of camera range on far left; staterooms are to far right.

The following are some typical observations about how people sit and stand at different places and times. These are all based on the time lapse of July 13, 2001. The identity of individuals is part of the public record (the meeting was filmed by the Discovery Channel); initials are used here.

- 1) [09:17:14] Everyone is at the table, and the meeting is started (then KQ and BC leave to get notebooks and clothing). Prior to this point there were never more than three people sitting at the table, although at different points in time the informal, pre-meeting conversation was joined by CC (at workstation), SB (at galley cabinet), and KQ (by the table).
- 2) Outside the formal meeting, SB rarely sits, while CC never leaves his workstation (aside from getting a drink). Those two appear to represent two ends of a volatility spectrum. CC works on one project, his paper; SB has many problems with the satellite network, walkie-talkies, power, etc. to resolve.
- 3) Later in the day, people spend relatively long times standing and pacing around the table: KQ ([11:54:49 - [11:55:52]); SB ([11:55:40] - [12:16:01]); CC ([12:06:53] - [12:08:59]). BC also has his notepad on the table, to which he returns periodically and makes notes while standing. The chairs are obviously still available, but they have been moved to workstation and the lab on the lower deck, where they “belong,” and nobody returns them to the table.
- 4) [15:01:24] VP sets up his laptop on the wardroom table, even though there is plenty of space available at workstation area (only two people are there). At [15:16:55], all but SB are sitting at their laptops.

A graph of the planning meeting (Figure 2) reveals some surprising patterns and provides a basis for characterizing behavior in terms of norms.

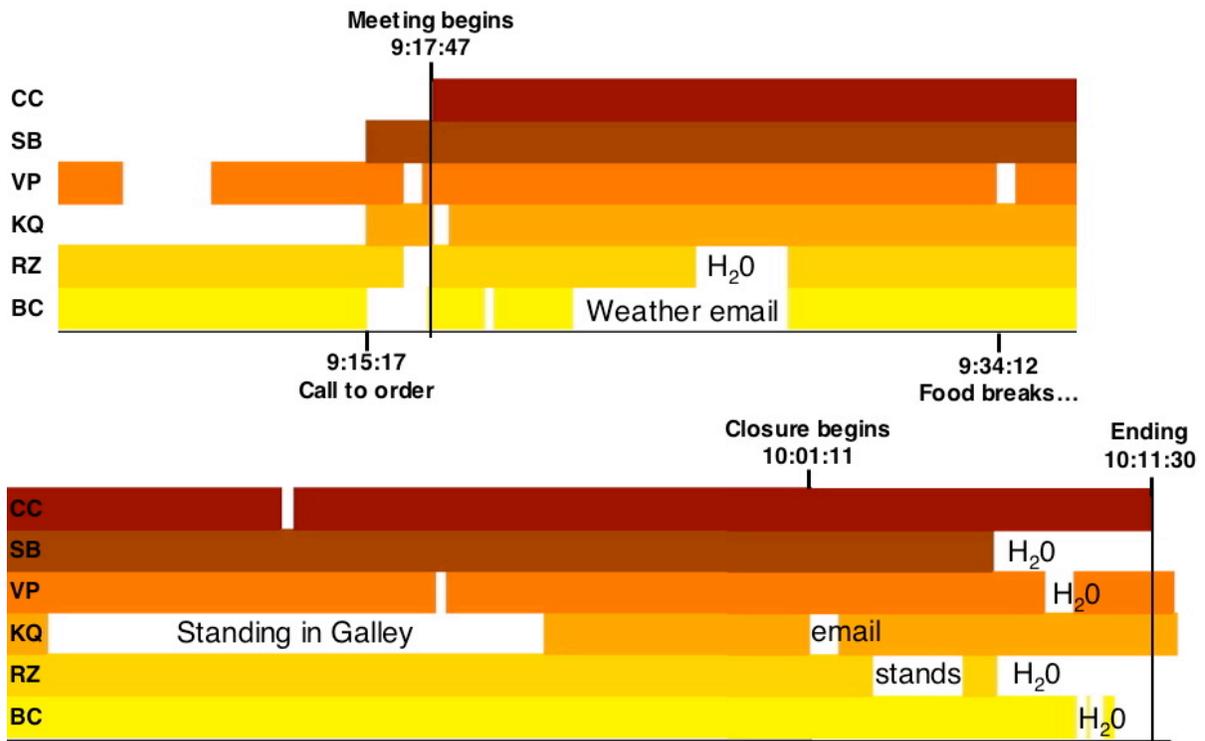


Figure 2. Location of crew members during planning meeting. The timeline is broken into two parts, starting at top left. Shading indicates seated at table, otherwise the activity is indicated. See text for analysis.

To understand what one needs to know about the structure of the meeting in order to simulate it, consider the problem of representing the locations and postures of the individual agents. At a first-order evaluation of simulation fidelity, before the model can be used to explain what is happening, the interacting agent behaviors must visibly resemble real life. This means that the graphic simulation must appear plausible to someone familiar with such settings. For example, it would be implausible to have the six people taking their chairs simultaneously or leaving at the same moment—any crew member knows that this never occurs.

The chart reveals what kinds of events are plausible, though they may still be unexpected to analysts because people do not necessarily reflect on even highly structured social behavior. Thus one observes a kind of “vetoing” of the meeting start when BC leaves his chair, just as RZ calls the meeting to order, which is the moment

when SB and KB have sat for the first time. Shortly after, RZ (meeting organizer) and VP leave. RZ begins the meeting when BC returns; simultaneously CC spins his chair around (waiting to the last moment to leave his personal work). Equally interesting is that KQ stands during about a third of the meeting, after reheating her drink in the microwave. This establishes a norm for the group: It is permissible to stand during the meeting, at least near the food area. At the very end of the meeting RZ stands and holds his chair in a way that appears to signify an ending. If someone were to stand and hold his/her chair in the same way in the middle of the meeting, it might appear that they are planning to leave for a moment, for example to go to the bathroom. VP & BC return to table after checking water (signifying that the meeting is not over). CC turns his chair around as the meeting ends, although two people remain at the table.

In short, modeling how individual agents carry out a group activity, as conditional actions organized into activity conceptualizations, begins to reveal how collective (social) behavior relates to individual cognition (involving perception, motive, and action). However, common sense knowledge about social behavior is far more complex than has been modeled in Brahms. In addition, social theorists (e.g., Lave 1988) suggest that every action within a group involves learning for all participants: Norms are being reinforced through their reproduction, but also adapted and even purposefully violated (e.g., for humor to confirm or deny emotional relationships). The FMARS simulation does not represent this learning (i.e., reinforcement or adaptation of workframes). Other social analyses suggest (Wenger 1998) that activity conceptualizations involve dynamic blending of identities, another aspect of learning that occurs as action that may not be deliberately planned. For example, FMARS crew members are always improvising their roles, as seen through their prior conceptualizations (e.g., “being a scientist on an expedition” “being a NASA representative”). In some respects, the interleaving of actions in different parallel activity conceptualizations models this blending in Brahms.

3.2 Planning Meeting Model Details

To create a model of the planning meeting, Brodsky and Clancey analyzed the time lapse video and wrote elaborate descriptions of the chronology of events. The following excerpt uses formatting to indicate the *located activities* of **AGENTS** using **objects**:

RZ requests weather info from **BC**. (They need it to decide whether to go for EVA). **BC gets up** from his chair, **walks** to workstation area, to his laptop (in a subarea), and checks weather report (for ~7 min; **sitting** facing laptop). After **BC** is done, he **walks back to** wardroom table area, approaches his chair area, and **sits down** on his chair. He then communicates the weather data to **RZ**.

Shortly after **BC** goes to check the weather, **RZ gets up** from his chair, **walks** to water tank area, **climbs** the water tank ladder, and **checks water level** (by looking into the **water tank** –standing on the ladder at the upper rim of **water tank level**, facing it).

On this basis, Brahms locations, agents, activities, and objects are related by declaring group-agent-activity relationships and writing workframes. For example, one part of the above sequence of events is modeled by this workframe (Brahms language constructs appear in **bold**):

```
workframe CheckWaterLevel
when (unknown(current.timeToFillWaterTank))
detectable DetectWaterLevel {
    detect((WaterTank.waterLevel = 0))
    then continue;}
do {
    Getup();
    Walk(GalleyLadderArea);
    Upladder(WaterTankArea);
    CheckWaterLevel();
    Downladder(GalleyLadderArea);
    Walk(WardroomTableArea);
    conclude((current.waterLevelChecked = true)); }
```

The subactivities in the **do** part are defined by other workframes, most of which use the **move** primitive activity.

After the simulation is run, the modeler may display agent actions using the AgentViewer (Figure 3). While RZ is checking the water level, BC is checking the weather report. Figure 4 shows this moment graphically using the Brahms Virtual

Environment (BrahmsVE; Damer 2004).

In the 2002 implementation, the simulation output is recorded in a database and mapped by BrahmsVE onto graphic primitives and scripts. The scripts generate short, agent-specific movements or gestures, such as walking up the ladder. In general, the scripts are created by analyzing photographs and videos, then developing storyboards, as if creating a cartoon or movie (Figure 5). These were reviewed for accuracy and plausibility, based on the ethnographer’s memory and records of events. For example, whether people would be able to or choose to squeeze between CC and the table instead of walking around is a matter of practice and should be rendered accurately. In general, the simulation might generate interactions that are not based on specific events; these must be evaluated for plausibility based on similar known events.

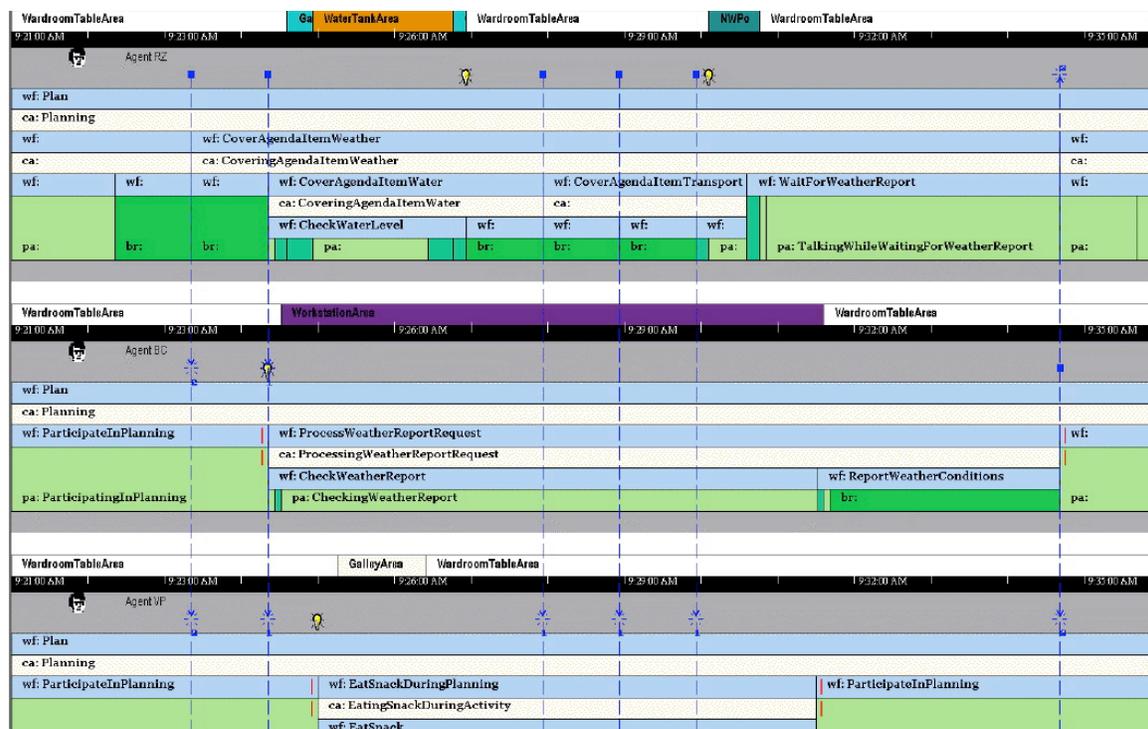


Figure 3. Brahms AgentViewer showing actions, communications, and inferences of agents RZ, BC, and VP during the first part of the planning meeting. At the time RZ does CheckWaterLevel, he is simultaneously engaged in Planning, Covering AgendaItemWeather and Covering AgendaItemWater.



Figure 4. Frame 3:24 from animation showing RZ checking the water level while BC is reading the weather report at his workstation [9:25:19]. Developed by DigitalSpace Corporation.

To illustrate the interface between Brahms simulation engine and the rendering system, consider the simple example of RZ doing the action: **move** Upladder(BaseAreaDef loc) {**max_duration**: 5; **location**: loc;} where loc is GalleyLadderArea. A program called OWorld Service converts this simulation event into the following scheduled animation:

```
activity|move|164|169|projects.  
fmarsvre.RZ|Upladder||projects.  
fmarsvre.GalleyLadderArea|  
projects.fmarsvre.WaterTankArea
```

Another program, OWorld Parser (implemented as Javascript in Adobe Atmosphere™), sends this scheduled animation to the BrahmsVE agent object queue. The RZ agent's Upladder action script executes the movement details. Figure 4 shows one of the frames.

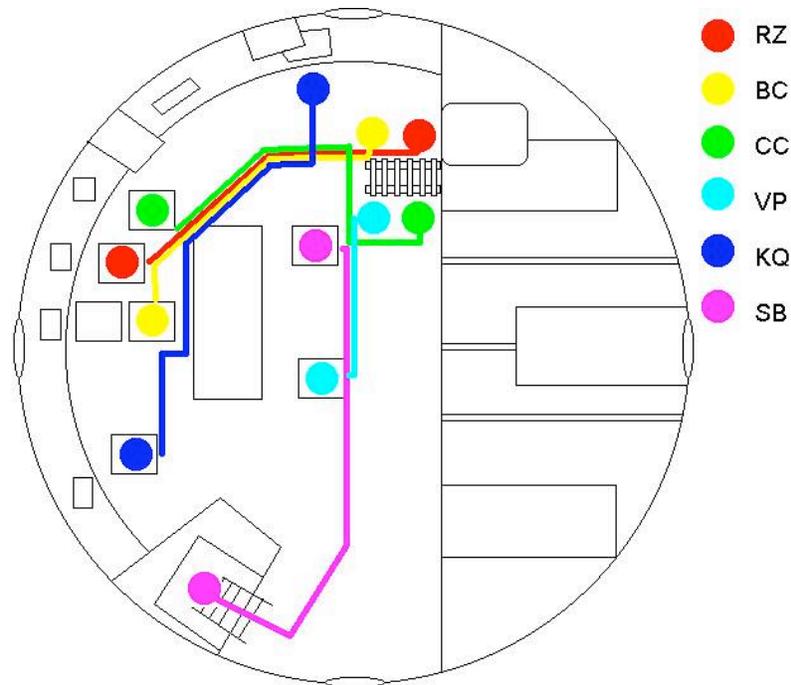


Figure 5. Initial storyboard showing ending of the planning meeting (DigitalSpace Corporation).

All together, three complex FMARS scenarios are simulated in BrahmsVE: the planning meeting (requiring 200 OWorld scripts), filling the water tank (67 scripts), and the EVA preparation (gathering equipment and helping each other don space suits, 423 scripts).

In this implementation, the rendering occurs in batch mode, after the simulation is completed. The timings of primitive motions and renderings are adjusted dynamically by the individual scripts, so they properly add up to the durations of Brahms activities. For example, a primitive activity in Brahms such as moving to the Galley Ladder Area, would require seven animation scripts, for getting out of a chair and walking, which together should total the five seconds declared in the Brahms model:

- Head Track Horizontal
- Head Track Vertical
- Stand Up From Chair
- Walk
- Turn While Walking

- Idle Standing(s)
- Breathe

The idle animations (e.g., shifting weight, moving arms) are random within the available time. Timing of primitive motions and renderings are not hard-coded in scripts, rather scripts are designed to play faster or slower to take the amount of time the Brahms model requires. An animation such as walking may take five seconds in real time, but if told by Brahms to take two, it will be accelerated, or it could be slowed if necessary. Waypoints must be specified by the graphic designers (one purpose of the storyboards), so the agents don't run through objects or into each other. Primitive motions refer to the waypoints in a general way, so they needn't be encoded in the script itself.

Using BrahmsVE, an analyst can now visualize postures and layout of the planning meeting. For example, one can see how RZ sits alone on one long end of the table (Figure 1), which is not visible in the AgentViewer. In effect, the graphic scripts of BrahmsVE represent part of the practice of the activity—the details of how people sit and move.

3.3 Modeling Biological Motives and Behaviors vs. Goals

Developing a multi-agent model of a day in the life of the FMARS crew naturally leads to including biological drivers of behavior, such as fatigue, hunger, and the need to use the bathroom. Such aspects of human behavior are ignored by most cognitive models (but see CLARION in Part 2 of this volume), but are emphasized by the discipline of psychology and design called *human factors* (e.g., Kantowitz and Sorkin 1983). Thus, a Brahms activity model provides a way to relate human factors concerns to cognition.

An activity model necessarily reveals that how people accomplish tasks within an activity (e.g., recording data while working at the computer in the workstation area) is affected by biological concerns (e.g., interrupting work in order to put on a sweater). At the same time, activities such as eating are interleaved with group activities (such as the planning meeting) and how they are carried out reflects the group's norms (e.g., one may get something to eat during an FMARS meeting, but would do this in a business office setting only if the food were already laid out for the participants in the meeting room).

In the FMARS simulation, biological needs are modeled in a simple way; the initial research objectives did not require replicating the state of the art of physiological

modeling. Each factor is represented by a single parameter (physical energy, hunger, urine in the bladder) that accumulates over time and is reset by a compensating action (rest, eating, elimination).

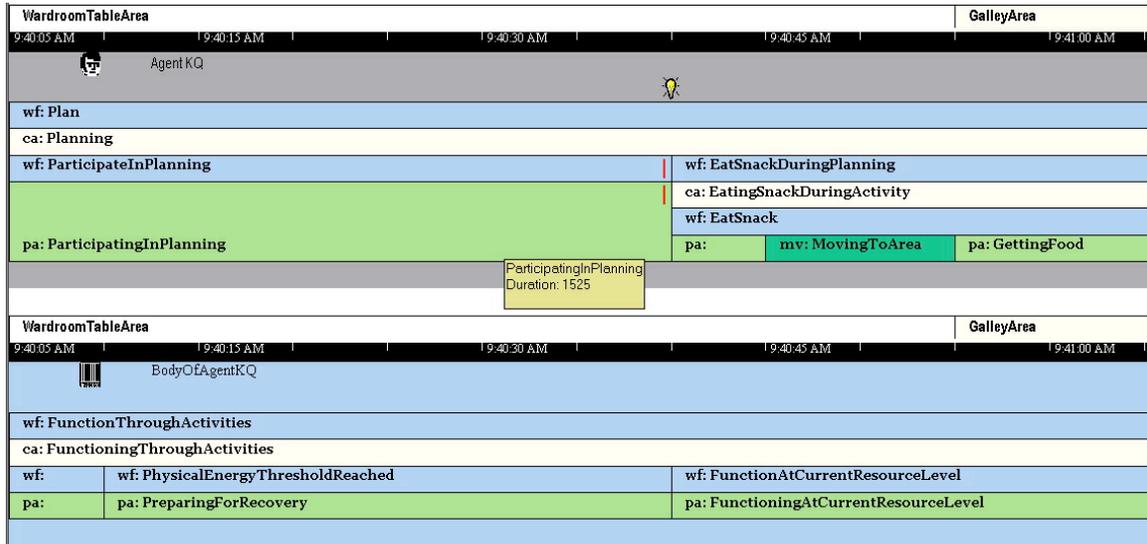


Figure 6. Brahms AgentViewer display of KQ Eating Snack during the planning meeting. This is a workframe within the Planning activity. Eating the snack is itself a composite activity, involving a sequence of conditional actions, including moving and getting food.

The inclusion of biological motives in explaining human behavior provides an interesting problem for cognitive modeling. For example, consider KQ warming her drink in the microwave and then standing by the side of the table (Figures 1 & 6). There are many explanations for this behavior: Her drink may be cold; she might be cold; her back may hurt; she may be bored with the meeting; someone at the table who hasn't had a shower in a week may smell, etc. One doesn't know her goals, aside from, perhaps, warming her drink. Even this may be a kind of convenient cover for accomplishing her "real intention."

Perhaps most interesting, the single action of standing to the side may be satisfying for several reasons, none of which need be conscious (i.e., deliberately reasoned to create a plan that the action carries out). Behavior may be determined by many physiological, personal, and social functions at the same time, and these need not be articulated or

distinguished by the person. A functional (goal-based) analysis tends to ascribe a single purpose to an action. A broad analysis of a day-in-the-life of the FMARS crew shows that of course all human activity is purposeful, but not every activity accomplishes a task (i.e., the work of the crew) nor can it easily be assigned to a single goal (i.e., a conscious proposition). This follows especially from the subsumption architecture in which multiple activity conceptualizations on different levels are affecting behavior by inhibiting, enabling, or blending actions (e.g., people in a meeting conventionally wait for an appropriate moment to use the bathroom). In contrast, when the crew discusses what EVA to do on this day, including where to go for what purpose, including what equipment and who should go, they are clearly engaged in goal articulation and planning. What is revealing is how much else is occurring that is modulated by perception of the environment and each other, physiological needs, and relationships (e.g., how people sit at the table, who chooses to remain silent)—modeled in the FMARS simulation without reasoning about goals and alternative plans of action.

Conventional goal/task analysis is a descriptive abstraction of human behavior, imposed by an observer, which may be an agent doing an activity. Goal/task analysis has implied that every human behavior has a direct goal as its cause (i.e., knowledge explains behavior). In contrast, the subsumption architecture in Brahms represents a conceptual nesting of activities, each of which has many implicit goal structures, so any behavior may make sense from multiple perspectives. It is far from clear whether KQ stands for one reason or five. (Notice how the “rationalist” framework suggests analyzing behavior in terms of *reasons* or reasoning.) Did a combination of activations cause her to stand at that time or were other satisfying relations emergent in the action (e.g., standing aside, she discovered that she related to the conversation better as an observer than as a direct participant)? It is highly problematic (if not theoretically impossible) to uniquely explain by subgoals behaviors that have not been *deliberately planned*.³ Instead, a Brahms

³ One can distinguish sensory stimuli (e.g., an odd feeling in the body), perceptual categorization (e.g., recognizing hunger), and conceptualization of a goal (e.g., “I will get something to eat”). Some perceptual categorizations may be reactive and not conceptually coordinated, as occurs in the stroop task, where the meaning of a word and

activity model represents the context in which the behavior occurs and (ideally) descriptively captures all gross movements, sequences, and communications. A goal-subgoal analysis can always be imposed later, and certainly a task analysis is necessary for designing layouts, procedures, work flow tools, etc.

4. Discussion: Lessons about Activity Modeling

This section considers lessons about the use of the virtual environment interface, methodology of constructing a Brahms model, how individual behaviors reflect and reinforce group dynamics, the relation of cognitive modeling and social interaction, and what can be learned by reconsidering Newell's social band framework. The section concludes with some remarks about applications of multi-agent simulations like FMARS to failure analysis.

4.1 Use of the virtual environment interface

The most important finding about the graphic interface is that it is not merely a display, but rather constitutes a second simulation—of the physical world—that must be integrated with the perceptual and action multi-agent model. That is, the modeler relegates to the virtual world simulation the physics of the real world influencing where and how agents and objects move (e.g., the microgravity of the International Space Station), line of sight, auditory range, and placement of objects on surfaces. In general, one would incorporate an anthropometric (human body) model, representing reachability and physical coordination in moving and holding objects. Work is underway to integrate the BrahmsVE with the agent simulation engine such that primitive actions with fixed durations and location would be modified during the physical simulation in the virtual environment. This is important not only for computing appropriate motion paths, but also to enable interruption of movements, for example, to allow to agents to encounter each other on the ladder and have a conversation. In effect, the notion of a primitive activity is fully open in Brahms, both to the purposes of the model (e.g., is fidelity in modeling the

its physical color conflict. Goal conceptualizations may also form reactively (which is one interpretation of KQ's standing), without reasoning about alternative motives, subgoals, or methods (i.e., deliberative planning).

hand required?) and the possible interactions that may occur between objects, agents, and the facility (e.g., an open stateroom door enables calling someone from outside).

The virtual environment itself was first conceived as an appropriate way to both construct and view Brahms simulations. The browser-based, distributed nature of the interactive 3D Adobe Atmosphere™ platform enables collaborative design and engineering, by which a common virtual world (e.g., FMARS) incorporates avatars (Damer 1997) that may interact with simulated agents, objects, and each other. In general, this could be a suitable framework for teleoperating teams of robots, especially with astronauts present, such as constructing and maintaining a lunar base. A more futuristic application would involve uploading agents to deep space such as to Mars or asteroids, where a time delay prevents conversation with Earth. Astronauts could converse with simulated agents, surrogates for human counterparts on Earth (e.g., the remote science team and specialized engineers), serving as coaches or assistants in real-time during Mars operations. The resulting interactions could be transmitted back to Earth and replayed to analyze and improve the work system design.

4.2 Methodology of constructing a Brahms model

The experiment of constructing a day-in-the-life FMARS model has reinforced the view that a Brahms model is a way of stating and organizing information about a work system. For example, after creating the model, Clancey received from CC a paper (Cockell et al. 2003) about doing biology in FMARS. The paper includes CC's view of his daily schedule. Using the full-day FMARS model, one could verify whether his summary fits what was observed (including time-lapse data).

CC distinguishes in his experience between an EVA day and a sample analysis day. A typical day includes an EVA, but not everyone goes out every day, and the model does not include what CC does on "an analysis day." The lesson learned is that simulating a sequence of multiple days is a heuristic for capturing work practices. Also, a simple interview may have revealed this distinction; one could ask, "Do spend your time in the same way every day?"

CC gives details about his scientific work that were not recorded or modeled (e.g., the names of his tools and their parts, and the lab equipment is in sequential order for sample

processing). He says he performed a procedure 100 times in two weeks; to verify this claim, another recording method is required, such as time-lapse on the lower deck or a log book near CC's microscope. He also tells us that he sent images to a colleague, an activity that was not observed, but might be learned by examining his email record.

The idea of modeling "a day in the life" is a starting point. The FMARS day simulated in Brahms is not intended to be a particular day, but a pastiche, something generalized from the available data, a typical day. The next step might be to refine the overall pattern to characterize types of typical days. Certainly modeling a sequence of days is as important for real applications (e.g., instruction and developing work flow tools) as having a full-day model.

Cockell et al. (2003) relate that CC had to abort his analysis work at one point to provide support for an EVA team, indicating how he detected the need for assistance: "during the science activity it is necessary for the scientist to be concentrating but aware of other activities...having an EVA radio close-by." This shows how an overarching activity (being the EVA support person) blends with a familiar activity (writing a paper), so it is carried out in a different modality. Furthermore, he says he was "constantly shuttling" between the decks. Time-lapse data provides the frequency on some days. If that were in the model, the statistics could be provided to Cockell for his own report.

Related work by Clancey during NASA's Haughton-Mars Project in 2003 showed that people were not accurate in estimating how often they were interrupted and for how long (e.g., a group stopped to navigate during an EVA every 3.7 minutes on average, while they estimated they drove for ten minutes between stops). These data suggest that people in highly interactive settings prone to interruption are not aware of the broader structure of emergent patterns, including the frequency of events. The analysis and simulation of group behavior is obviously of great value for capturing and visualizing these patterns.

Developing a model of social behavior consequently has a special challenge that conventional cognitive modeling may not—patterns are often undetected by participants who are immersed in the setting, and even an observer may miss the regularities. A striking example is Clancey's (2001) analysis of the Haughton-Mars expedition in 1999, revealing that what people called the "work tent" was most often visited for less than two minutes, and was in fact primarily a place for storing things. This pattern was not

detected while working inside the tent, but was only clear from the statistical analysis of the group's behavior over a day, which time-lapse video allowed. Thus, some means is required for capturing located behaviors over time, so that what individuals are doing becomes visible. The statistical patterns (e.g., frequency of interruption) may be emergent in the simulation as it is run for many simulated hours, but one must somehow learn what activities are occurring. An observer working in a "work tent" will not easily see all the people coming and going, because they are part of the background and tuned out like so many gnats. In contrast, a conventional cognitive model is constructed from a task whose parameters are fully defined by the modeler, and all that must be observed are operations for transforming the materials or describing the situation.

In summary, a fundamental problem in constructing a model of social behavior is knowing what everyone is doing at all times. A Brahm's model provides a way of organizing observations (and redesigns), so particular information can be easily viewed and brought into juxtaposition and related. Conventional ethnographic text (e.g., field notes or analytic memo) does not enable relating data in this way. As the examples illustrate, it is particularly interesting to attempt to discover and replicate frequencies of recurrent events, such as how often people are interrupted in their work setting.

4.3 How individual behaviors reflect and reinforce group dynamics

Throughout the FMARS analysis we have been struck by how individual behavior ranging from seconds to minutes is sensitive to other people's interpretations and actions. A good example is the process by which individuals stop what they are doing and arrive at the meeting. As known from common experience, groups tolerate varying degrees of lateness, and in a situation where communication is possible, as in the FMARS habitat, one may negotiate the start of the meeting ("I just need a few minutes to finish [photographing this rock slice]").

More interesting is how people notice, through their peripheral awareness of the group arriving at the table, that they must hurry. For example, someone on the lower deck can hear the difference between four people at the table and two, and may notice that he/she is now alone. Whether a meeting starts on time and how an individual may cause others to wait is a paradigmatic norm for the group. More broadly, how individuals balance their

own agendas as scientists (with papers and sponsors to satisfy) against the group's objectives and imposed responsibilities (e.g., chores) is starkly revealed when individual work is simulated within a day-in-the-life context. This is a rich phenomenon for further investigation. How are individuals rationalizing their actions, and where do they draw the line in compromising or adapting their original plans as problems such as resource constraints develop within the group?

Finally, the effort to graphically render the FMARS Brahms simulation has allowed us to model gestures, routes, and field of view, though none of these are yet incorporated in the simulated agents' perception and hence do not affect simulated actions. Research continues to close the loop so the physics model in BrahmsVE feeds back to the simulation while it is running, thus routes will affect how long a movement takes, and fields of view (and hearing) will affect what the agent can perceive. Modeling an agent's perception of gestures and relating them to individual behavior is complex, but is fundamental for relating cognition to social behavior. Figure 7 provides a glimmer of what could be involved.



Figure 7. Unusual posture at end of planning meeting. Square standing distribution suggests a balanced or stable relationship. Individuals move into and hold the encounter in this position. Possibly an important issue is being

reconsidered.

4.4 Distinguishing ways of working together

Another understanding that has resulted from this work is recognizing that people are often working together but not collaborating. For example, the group sometimes sits in the habitat, reading and working on computers without talking, in effect, “working together alone” (Figure 8). They are cooperating in sharing a resource (the facility), but not working on the same project.



Figure 8. The activity of working alone together, an example of cooperating without collaborating. FMARS initial habitation, August 2000

The FMARS investigation, plus related work studying field scientists (Clancey 2004b), has suggested the following distinctions:

- **Coordination:** Sharing a common resource via scheduling or ordering, without requiring changes to how individuals or subgroups behave, e.g., sharing the habitat’s “mess table” during the day. Literally, “co-ordinating,” ordering in time and place to avoid any possible interference with others’ activities.
- **Cooperation:** Sharing a common resource in a way that requires adjusting how individuals or subgroups carry out an activity, e.g., sharing space on the table

during the meeting. Literally, “co-operating,” operating in a way that relates individual actions in time and place. Work flow typically describes how different functional roles cooperate, with one product feeding into another task.

- **Collaboration:** Working on a common project, e.g., most of the planning meeting is devoted to the daily EVA, which will require three or four members of the crew to work together for half of the day or more. Literally, “co-laboring,” conceiving and carrying out a single project. Most generally, this is a triad, two or more agents (or groups) and a group. The relation is in general asymmetric: A and B collaborate on a project originated by A (but might do no work together on B’s project). For example, a geologist may help a biologist do a study in the field, but the biologist doesn’t contribute to the geologist’s investigation (Clancey 2004b).

“Working quietly in the hab” is a cooperative group activity, in which individuals pursue their own agendas. In general, the crew’s schedule is designed to balance collaborations (common projects) with individual agendas stemming from personal needs and interests (e.g., reading a book about the Arctic), disciplinary specialization (e.g., microphotography), and institutional commitments (e.g., writing a column for a news organization). Understanding the relation between individual drivers of behavior and group activities is a fundamental aspect of understanding how cognition relates to social interactions.

4.5 Summary of relation between cognitive modeling and social interaction

To summarize the example and discussion to this point, consider some of the questions posed by Sun (chapter 1) for relating cognitive modeling and social interactions:

- 1) What are the appropriate characteristics of cognitive architectures for modeling both individual cognitive agents and multi-agent interactions?

Experience constructing six work practice models in Brahms suggests that the following Brahms language features are relevant:

- a. Subsumption architecture for conceptualization of activity
- b. Physical layout of facilities modeled explicitly; all behaviors are located
- c. Communication of beliefs (Ask and Tell)
- d. Context-dependent perception (activity-specific detectables)

- e. Interruption of activities based on priorities and detected conditions
 - f. Model representational objects that agents can read and write (e.g., documents)
- 2) What are the fundamental ways of understanding and modeling multi-agent interactions? How much can they be reduced to individual cognition?
- a. Reductionism is inappropriate; it is better to begin by asking: How can patterns of social interactions emerge from individual cognition and behaviors? What is the nature and role of subconscious perception of interactions by individuals (cf. Figure 7)?
 - b. Ethnography (participant observation) is the fundamental way of understanding and modeling multi-agent interactions: photos, video, time-lapse, activity mapping (person, time, place) (Clancey 2001; 2004a, b; Jordan 1992; Scribner & Sachs 1991; Wynn 1991) .
 - c. As a heuristic, model at least a day in the life of the group (24 hours); move to multiple days as soon as practical; especially, consider the rhythm of a week (Clancey 2002a).
 - d. Model both group and individual activities; consider how the methods for accomplishing goals are adapted in cooperative activities; recognize that not all group activities are collaborative.
- 3) What additional representations...are needed in cognitive modeling of multi-agent interactions?
- a. Activities (including motives, goals, operations) described by Activity Theory (Leont'ev 1979)
 - b. Biological needs (fatigue, hunger, toilet, cold) affect choice of activity and manner of carrying it out.
 - c. Perception of posture, attitude, tone of voice, etc. affect relationships (not included in Brahms).
 - d. Perception of space, line of sight, voices (e.g., to determine paths, what is visible, what can be heard)
 - e. Facilities (e.g., lack of proper heating at FMARS, available work space) influence personal experience and attitude towards cooperation.

- 4) How can we best characterize and model social relations, structures, and organizations in relation to individual cognition?
 - a. See 2d
 - b. In a multi-agent simulation, social structure can be modeled in terms of the activities of groups to which agents belong.
 - c. Model roles (e.g., meteorologist) and identities (e.g., graduate student) as inherited group behaviors.
 - d. Model behaviors descriptively: What individuals do when and where for how long—do not focus on goals and tasks.
 - e. Model the broad activity chronology of a day and refine to tasks to the level required for the application of the model
 - f. Focus on how group activities begin, the norms for how they are carried out, and how they are brought to a closing
 - g. Attempt to model belief change as much as possible in terms of communication, perception, and forward-chaining; goal-directed inference occurs during planning activities (e.g., deciding what to do next)—observe why and how often it occurs
 - h. Do a statistical analysis of where people are located and what they are doing throughout a day
 - i. Observe reminders and peripheral attending (how individuals keep each other synchronized); group and individual tolerance for delays
 - j. Consider how the group decides whose knowledge will be called into play and how individual methods of working are facilitated, blended, or inhibited by the group's schedule, other goals, or conflicting modes of operation (e.g., when one is driving in a caravan during an EVA it may be impossible to stop and take photographs)
 - k. Recognize that some social patterns (e.g., paths left by ATVs) may be perceived and direct individual behavior; others may be only tacitly conceived and yet be influencing individual behavior (e.g., how people arrange themselves and interact, Figures 1, 2, 7, and 8).

This outline resembles more a list of examples than a comprehensive perspective and goes beyond what is incorporated in the planning model. Thus at least from the perspective of this project it represents the edge of scientific understanding.

4.6 Relation to Newell's Social Band Framework

One way of appraising progress is to compare the FMARS planning model to Newell's (1990) discussion of the "social band" in *Unified Theories of Cognition*. Newell's position was comprehensive and contains many sound pieces of advice: "models of the individual as intelligent agents interacting with ... real worlds would seem essential" (p. 493). The aspect of his analysis that appears perhaps most foreign are the "system levels" called "bands." By analogy to physical computer systems, the bands are defined in terms of time scales, with the social band having "time units" of days to months (p. 152). In contrast, simulating the most simple norms, such as standing at a table during a meeting, involves momentary dynamics of perceiving and moving within a conceptualization of the conscious person ("what I'm doing now," Clancey 1999).

Possibly Newell viewed "social" as just meaning *direct, physical interaction* with others: "As the time scale increases from days to weeks to months, the systems involved become social. They comprise multiple individuals in interaction. Humans do not lead solitary lives at the level of days and above" (p. 154). The idea that all human activity is *socially conceived* (in terms of the norms of roles, methods, purpose), so solitary activity is always inherently social, was apparently not part of Newell's notion of social or his notion of knowledge. He viewed knowledge as "socially conditioned" (p. 490) as opposed to being formulated in social terms ("who am I being now?" Clancey 1997a).

Anderson (2002) makes a similar conclusion: "Newell thought that issues of cognitive architecture became relatively unimportant at the Rational Band and were completely irrelevant at the Social Band" (p. 3-4). Indeed, Anderson disagrees with Newell: "fine-grained temporal factors at the Biological Band do influence higher-level outcomes" (p. 4). But Anderson's analysis focuses on the mechanism of "unit-task" learning, rather than the individual's *conceptualization* of motivation and value (Lave & Wenger 1991)—social factors that explain why learning is occurring at all.

Newell claimed that “the group’s behavior is explainable and predictable by its use of knowledge in service of its goals” (p. 154). This is by definition true when one constructs a model that refers to conditional actions as “knowledge” and describes all behavior as deriving from goals. However, as shown in this chapter, other kinds of models are possible. More generally, a group’s behavior is explainable and predictable by 1) interacting normative behaviors of individuals (e.g., when the planning meeting begins depends on how long they delay after the commander’s call to order) and 2) habitual patterns of “how we do things,” which are not all scheduled or reasoned about in plans (e.g., sharing hot water during breakfast, allowing people to stand during the middle of a meeting).

Referring to all human behaviors as determined by goals and knowledge seems inappropriate when a day in the life of a group such as the FMARS crew is considered. The task-goal-knowledge analysis applies best when people are working on specific tasks, focusing on using laboratory equipment, downloading and analyzing EVA science data, or preparing a meal. Put another way, at the time scale of 10 seconds or more—Newell’s “Intendedly Rational Band” (p. 150)—behavior is both deliberately reasoned about *and* habitually patterned by previous interactions. Although one may ignore biological impulses during intendedly rational activities (e.g., continuing to read a fascinating book chapter despite having the urge to use the bathroom), all behaviors are always in a social context, that is, they are *conceived* with respect to social norms, relations, and values. People frame their activity in terms of their socially constructed identities; this determines what they do, when, where, and how, including what problems they discover or tackle and what methods they use to resolve them (for elaboration, see Clancey 2002b).

In summary, the heuristic of modeling a day in the life of a group living and working together reveals an interaction of biological, task-oriented cognitive, and social influences that cannot be separated into temporal bands. Social behavior is not only occurring (or rolling up) over longer time scales as Newell posits, in the manner of individual actions accumulating into a social history or a person being forced to interact with others (e.g., going to the store to buy milk). The “bands” in Newell’s analysis are not isolated systems

in practice. Different emergent aspects of the scene (biological, task-goal oriented, and collective) causally influence each other:

Biology and culture interpenetrate in an inextricable manner.... Individuals are not real and primary, with collectivities...merely constructed from their accumulated properties. Cultures make individuals too; neither comes first, neither is more basic.... Thus, we cannot factor a complex social situation into so much biology on one side, and so much culture on the other. (Gould 1987, p. 153)

Cognition—whether the person is physically alone or in a group—is immersed in norms and emergent physiological, physical, and cooperative constraints (Wynn 1991)⁴.

4.7 Application to failure analysis

Because NASA's failure analysis reports (e.g., CAIB 2003) consistently emphasize social problems, it is worth considering how a Brahm's activity analysis might be useful in understanding or identifying organizational and cultural problems in a highly structured task setting. One approach is to represent how people are actually conceiving of a given activity in broad terms. For example, as MER scientists are working at JPL during a Mars rover mission, do they conceive of their activity as geologists exploring Mars or see the mission through the eyes of the "flight control" team operating a rover? How do these conceptions interact as concerns in practice and influence the quality of the outcome from scientific and engineering perspectives? Notice how this analysis is different from a task model that frames the problem in one way (e.g., controlling the rover) or uses a multi-tasking or linear architecture (e.g., first I solve the geology planning problem, then I solve the flight control sequencing problem). In practice, these tasks are not strictly partitioned into different roles, nor when they are separated organizationally can the constraints be strictly ordered. An activity analysis asks how a given individual might be blending alternative ways of perceiving, interpreting, and acting, such that they experience conflicts in their judgment (e.g., as a geologist, I'd first take a look over the top of this crater I'm standing in and possibly return, but the mission success criteria imply that the

⁴ For a more detailed discussion, see the chapter "Dialectic Mechanism" in Clancey (1997b) as well as the discussion of Maturana's "structural coupling" (p. 89).

rover's path must omit loops). In a task analysis, these are just "conflicting goals."

One purpose of a social simulation of work practices is to understand how "intendedly rational" behaviors fail to accomplish goals within broader time scales because behavior derives from norms and emotions, and not just local reasoning about technical matters. An example appears in the *Columbia Accident Investigation Board Report* (2003), involving a management meeting that reviewed and accepted a faulty damage analysis. People based decisions on previous interpretations of similar problems and scheduling constraints for subsequent launches. A social analysis is required to explain why knowledge and concerns of individuals and subteams were not brought to bear. In this case the norms of management prevented specialists from getting data they needed to support their tentative damage analyses, creating a Catch-22 situation.

The Columbia disaster highlights how the group's roles, schedules, and even representational practices (e.g., PowerPoint bullets, Tufte 2003, pp. 7-11) determine the salience of events—how to evaluate a situation, what effects are important, and hence what constitutes a problem and how or to what extent it is resolved. The FMARS models shows how cognitive modeling might apply to real-world applications by developing a multi-agent simulation, with multiple groups interacting over a day or more. Just as conventional task analysis works backwards from goals to knowledge, an activity-based analysis works backwards from the quality of the work product (e.g., ways in which it fails) to the representations (e.g., presentations at meetings), interactive patterns (e.g., how time is allocated during a meeting), and norms of authority that influence who may speak to whom about what, when and where. Modeling these relations and effects in Brahms in a general way is an open research problem.

How were people during the Columbia management meetings conceiving of their activity? Planning for the next launch or trying to return the crew safely? Were they conceiving the meeting as managing the agenda (i.e., controlling who participates and how) or trying to ferret out and understand anomalies? Of special interest to the Columbia analysis are informal (not role or task-defined) communications by which people assist or influence each other, a consideration naturally revealed when a modeler focuses on describing behaviors instead of only goals and inferences. In other words, communication of information is not necessarily traceable to missing or wrong technical

knowledge, but instead will point to misconceptions about practice, a presumption about how the work is supposed to be done, including especially lines of authority and when and how people are allowed to influence the group's work. Thus modeling how people conceive of their activity, which is always pervaded by social relations, is essential for explaining human behavior. This is a very different kind of cognitive model than emphasized heretofore in understanding expertise and problem solving ability.

5. Summary

Simulating an FMARS planning meeting in Brahms produced several surprises:

- “Off-task” activities of eating, resting, using the toilet, and recreation (e.g., playing games or talking at the table) must be included in a work practice simulation because they causally affect the duration, timing, and methods by which tasks are accomplished.
- Characteristics and experiences of people often studied by human factors specialists (e.g., hunger and fatigue), which are typically excluded from cognitive models, must be included a work practice simulation because they determine when off-task activities occur.
- Everyday behaviors, such as getting something to eat, are carried out according to norms, but improvised in a way (e.g., standing while eating during a meeting) that exercise the open nature of norms, while possibly accomplishing many goals simultaneously. Such behaviors appear to blend rituals or habits with both premeditated intentions and emergent affects (e.g., calling attention to oneself and hence being better able to influence the decisions being made).
- The non-immersive virtual display of BrahmsVE, which was at first considered to be only a “visualization tool,” provides a means of simulating line of sight and movement paths—information that is essential for simulating what agents can detect in the environment and how long movement between two points requires.

The heuristics of modeling a full “day in the life” of the habitat and simulating all agent movements and use of tools were crucial for making these discoveries.

The modeling experiment shed a different light on what cognition accomplishes and how perception and action are related through conceptualization of activity. For example,

a conventional cognitive model of a planning meeting would focus on the discourse structure of the meeting's conversation. Such fine-grained explanations of topic relationships, based on the semantics of what is being presented, explained and decided, might be improved by including what the FMARS model focused upon: postures (e.g., which may convey boredom or disagreement to participants), transitional activities (e.g., how individual agent behaviors become coordinated into a coherent group activity), and biological motives (e.g., fatigue, which may affect the meeting's agenda).

In some respects, behaviors emphasized in the FMARS model might be viewed as noise in a conventional discourse model. For example, it might appear humorous to ask a cognitive modeler, "What if the person is hungry and doesn't want to continue talking?" Cognitive simulations often assume that people are motivated (i.e., the goals of the task at hand are not in question) and that work occurs in a controlled setting. The FMARS simulation emphasizes that the context includes people's activities, which have both broad and narrow forms that influence what goals and methods are established, how they are adapted—affecting the quality of the resulting work.

This chapter has focused on what can be learned from the use of a virtual environment interface, the methodology of constructing a Brahm's model of practice, how individual behaviors reflect and reinforce group dynamics, the relation of cognitive modeling and social interaction, and what can be learned by reconsidering Newell's social band framework. The examples throughout illustrate many aspects of behavior that protocol analysis would not consider because they are visual relationships (e.g., how people stand when talking), off-task (i.e., would not be included in an experimental setting that presents a task to a subject), and conceptualizations that are not articulated in common experience or sought in task-oriented studies (e.g., understanding of norms, how participation is negotiated).

The observational methodology used in the FMARS study includes both systematic (e.g., time-lapse video) and informal (e.g., field notes) records. By design, the recording is intended to record and learn more than can practically be analyzed, and thus (perhaps) include information that is only later found to be useful (as illustrated by the analysis of the July 13, 2001 planning meeting). Clancey (in preparation a, b) shows how time lapse,

diaries, and surveys can be systematically recorded and analyzed to produce information about productivity and work system design problems.

The focus of the FMARS simulation is to provide a proof of concept that the simulation can fit what actually occurs. The main criteria used were the episodes visible on the time lapse (e.g., movement of crew members during the planning meeting for different reasons) and the duration of events. As discussed in considering CC's report of his crew experiences, to more thoroughly verify the model would require simulating at least a week, which is well beyond what modeling resources have permitted. The present model includes three episodes identified as recurrent and involving distinctive combinations of attention and interpersonal interaction (the planning meeting, refilling the water tank, and preparing for an EVA, Clancey 2002a). From the perspective of practical design and ongoing Mars analog investigations, the most important scientific product of such research is identifying new issues to systematically study (e.g., the frequency of interruptions, Clancey in preparation b).

Although cognition is sometimes considered narrowly as relating goals, inference, and actions, cognitive science (as represented by the journal and society of that name) more broadly includes perception, the nature of conceptualization, social interaction during learning, and many other topics. This chapter focuses on relating collective (social) behavior to individual cognition (involving perception, motive, and action) by emphasizing that individual behaviors are conceptually coordinated with respect to an understanding of norms. Such an investigation touches upon the nature of culture, as embodied in individuals (Lave 1988), and realized in episodes that exercise, extend, test, and interpret other people's conception of how to behave. These normative behaviors include: What topics should be discussed when, by whom, and using what tools?

Thus, the analysis presented is part of a much larger project that might examine the decisions made during the planning meeting, and tie them to interpretations of the group's role structure, competing motivations, and so on. This analysis would again be primarily episodic until many such meetings had been analyzed and statistically related. The FMARS 2001 rotation studied here ended after a week, and the group never lived or worked together again. Developing a full-fledged theory of such social interactions may therefore require a series of related studies in other contexts.

Finally, the FMARS modeling experiment illustrates what mechanisms other than backward chaining of goals capture, given a focus on simulating the activities of a typical day, rather than automating a task. The project revealed the relation of different levels of analysis (biological, psychological, social). A contrast can be drawn with multi-agent models that focus on functional actions. For example, Brahms' design was inspired by the Phoenix system (Cohen et al. 1989), which showed how an environment model of a fire-fighting setting interacted with a hierarchical communication and command structure. If modeling fire-fighting in Brahms using the same approach used for FMARS, one would model the entire day, including where the fire-fighters camp, how meals are prepared, how they are transported to the work site, etc. This day-in-the-life model would complement Cohen's multi-agent task analysis, revealing how social activities are interleaved with and constrain how work is actually done.

Understanding the nature and influence of individual emotions, agendas, preferences, ambitions, etc. is a significant next step. Thus, the intersection of cognitive and social analyses broadens the research perspective—from what knowledge is required to accomplish a task, to why certain people are participating at all. How do leaders in high-risk situations manage fear and temerity in assigning individuals to tasks? To allowing someone to present a contrary view and plan to the group? A question for cognitive modeling then becomes, what knowledge and motives affect who is allowed to participate and in what manner?

Acknowledgments

Our colleagues, especially Paul Feltovich, Chin Seah, Dave Rasmussen, and Mike Shafto, as well as other members of the Work Systems Design and Evaluation group in Computational Sciences at NASA Ames, have made important contributions to this research. We also thank the FMARS Rotation #2 July 2001 crew for providing explanations of their work and completing surveys. The FMARS study would not have been possible without the Mars analog concept, support, and facilities of the Mars Society led by Robert Zubrin. Field support and research ideas were also provided by Pascal Lee and the Haughton-Mars Project (1998 to 2003). This work has been supported in part by NASA's Computing, Communications, and Information Technology Program,

Intelligent Systems subprogram, Human-Centered Computing element, managed by Mike Shafto at NASA Ames. DigitalSpace Corporation has been funded through SBIR and STTR NASA contracts (see www.digitalspace.com/reports/index.html).

References

- Acquisti, A., Sierhuis, M., Clancey, W. J., Bradshaw, J. M. (2002). Agent-Based Modeling of Collaboration and Work Practices Onboard the International Space Station. Proceedings of the 11th Conference on Computer-Generated Forces and Behavior Representation, Orlando, FL, May, pp. 181-188.
- Anderson, J. R. (2002). Spanning seven orders of magnitude: a challenge for cognitive modeling. *Cognitive Science*, 26(1), 85-112.
- Clancey, W. J. (1993). Situated action: A neuropsychological interpretation (Response to Vera and Simon). *Cognitive Science*, 17(1), 8-116.
- Clancey, W. J. (1994). Comment on diSessa. *Cognition and Instruction*, 12(2), 97-102.
- Clancey, W. J. (1995). A boy scout, Toto, and a bird. In L. Steels and R. Brooks (Eds.) *The "Artificial Life" Route to "Artificial Intelligence": Building Situated Embodied Agents* (pp. 227-236). New Haven: Lawrence Erlbaum.
- Clancey, W. J. (1997a). The conceptual nature of knowledge, situations, and activity. In P. Feltovich, K. Ford, & R. Hoffman (Eds.), *Human and Machine Expertise in Context*, (pp. 247–291). Menlo Park, CA: The AAAI Press.
- Clancey, W. J. (1997b). *Situated Cognition: On Human Knowledge and Computer Representations*. NY: Cambridge University Press.
- Clancey, W. J. (1999). *Conceptual Coordination: How the Mind Orders Experience in Time*. Hillsdale, NJ: Lawrence Erlbaum.
- Clancey, W. J. (2001). Field science ethnography: Methods for systematic observation on an arctic expedition. *Field Methods*, 13(3), 223-243. August.
- Clancey, W.J. (2002a). Simulating "Mars on Earth"—A Report from FMARS Phase 2. In F. Crossman and R. Zubrin (Eds.) *On to Mars: Colonizing a New World* (CD-ROM). Apogee Books.
- Clancey, W. J. (2002b). Simulating activities: Relating motives, deliberation, and attentive coordination. *Cognitive Systems Research*, 3(3), 471-499.

- Clancey, W. J. (2004a). Automating CapCom: Pragmatic Operations and Technology Research for Human Exploration of Mars. In C. Cockell (Ed.) *Martian Expedition Planning, Vol. 107, AAS Science and Technology Series*, pp. 411-430.
- Clancey, W. J. (2004b). Roles for Agent Assistants in Field Science: Understanding Personal Projects and Collaboration. *IEEE Transactions on Systems, Man and Cybernetics*, Part C: Applications and Reviews, Special Issue on Human-Robot Interaction, 34(2) 125-137.
- Clancey, W. J. (in preparation a). Observation in natural settings. To appear in K. A. Ericsson, N. Charness, P. Feltovich, & R. Hoffman, *Cambridge Handbook on Expertise and Expert Performance*, "Methods for Studying the Structure of Expertise." New York: Cambridge University Press.
- Clancey, W. J. (in preparation b). Participant observation of a Mars surface habitat mission simulation. To appear in *Habitation*.
- Clancey, W. J., Sachs, P., Sierhuis, M., and van Hoof, R. (1998). Brahms: Simulating practice for work systems design. *Int. J. Human-Computer Studies*, 49, 831-865.
- Cohen, P. R., Greenberg, M. L., Hart, D. M., & Howe, A. E. (1989). Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3), 34-48.
- Columbia Accident Investigation Board. (2003). *CAIB Report, Volume 1*. NASA. (Online), August. <http://www.caib.us/news/report/volume1/default.html>
- Cockell, C. S., Lim, D.S.S., Braham, S., Lee, P., and Clancey, W. J (2003). *Journal of the British Interplanetary Society*, 56(3-4), 74-86.
- Damer, B. (1997). *Avatars: Exploring and Building Virtual Worlds on the Internet*. Berkeley: Peachpit Press.
- Damer, B. (2004). Final Report, SBIR I: BrahmsVE: Platform for Design and Test of Large Scale Multi-agent Human-Centric Mission Concepts, DigitalSpace Documents. (Online), July. www.digitalspace.com/reports/sbir04-phase1
- Gould, S. J. (1987). *An Urchin in the Storm: Essays about Books and Ideas*. New York: W. W. Norton.

- Jordan, B. (1992). Technology and social interaction: Notes on the achievement of authoritative knowledge in complex settings. IRL Technical Report No. 92-0027. Palo Alto, CA: Institute for Research on Learning.
- Kantowitz, B. H., and Sorkin, R. D. (1983). *Human Factors: Understanding people-System Relationships*. New York: John Wiley.
- Lave, J. (1988). *Cognition in practice*. Cambridge, UK: Cambridge University Press.
- Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. New York: Cambridge University Press.
- Leont'ev A. N. (1979). The problem of activity in psychology. In Wertsch, J. V. (Ed.), *The Concept of Activity in Soviet Psychology* (pp. 37-71). Armonk, NY: M. E. Sharpe.
- Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Schön, D. (1987). *Educating the Reflective Practitioner*. San Francisco: Jossey-Bass Publishers.
- Scribner, S., and Sachs, P. (1991). Knowledge acquisition at work. *IEEE Brief*. No. 2. New York: Institute on Education and the Economy, Teachers College, Columbia University. December.
- Seah, C., Sierhuis, M., & Clancey W. (2005). Multi-agent Modeling and Simulation Approach for Design and Analysis of MER Mission Operations. *Proceedings of 2005 International Conference on Human-Computer Interface Advances for Modeling and Simulation (SIMCHI'05)*. January, pp. 73-78.
- Sierhuis, M. (2001). *Modeling and Simulating Work Practice*. Ph.D. thesis, Social Science and Informatics (SWI), University of Amsterdam, The Netherlands, ISBN 90-6464-849-2.
- Sierhuis, M., and Clancey, W. J. (2002). Modeling and simulating work practice: A method for work systems design. *IEEE Intelligent Systems*, special issue on Human-Centered Computing at NASA, 17(5), 32-41.
- Sierhuis, M., W. J. Clancey, Seah, C., Trimble, J., and Sims, M. H. (2003). Modeling and simulation for mission operations work systems design. *Journal of Management Information Systems*, 19(4), 85-128.

Clancey, et al: Cognitive modeling of social behaviors

Spradley, J. P. (1980). *Participant Observation*. Fort Worth: Harcourt Brace College Publishers.

Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine communication*. Cambridge, UK: Cambridge University Press.

Tufte, E. R. (2003). *The cognitive style of PowerPoint*. Cheshire, CT: Graphics Press LLC.

Wenger, E. (1998). *Communities of Practice: Learning, Meaning, and Identity*. New York: Cambridge University Press.

Wynn, E. (1991). Taking practice seriously. In J. Greenbaum and M. Kyng (Eds.), *Design at work: Cooperative Design of Computer Systems* (pp. 45-64). Hillsdale, NJ: Lawrence Erlbaum.