The Cognitive Modeling of "Day in the Life" Social Behaviors Using Brahms

William J. Clancey[1], Maarten Sierhuis[2], Bruce Damer[3], and Boris Brodsky[4]

NASA-Ames Research Center, Moffett Field, CA 94035
William.J.Clancey@NASA.gov
http://bill.clancey.name

Introduction: The interaction between cognitive and social processes

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 we examine groups of people living and working together, we are forced to recognize that *whose knowledge is called into play* at what 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 (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 *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 exclusively of a model that deduces actions from goals, behaviors are *primarily* driven by broader patterns of chronological and located activities (akin to scripts).

Rather than saying (only) that knowledge drives behavior, we say that *conceptualization of activities drives behavior*, which includes how knowledge is called into play and applied in practice. Put another way, how problems are discovered and framed, what methods are called into play, and indeed who even cares or has the authority to act, are all constrained by norms, which are conceived and enacted by individuals.

---

[1] NASA/Ames Research Center and Institute for Human and Machine Cognition, UWF, Pensacola, FL
[2] RIACS, NASA/Ames Research Center, Moffett Field, CA
[3] DigitalSpace Corporation, Santa Cruz, CA
[4] Titan Corp, NASA LaRC, Hampton, VA

Of special interest for the cognitive modeler, and which is emphasized in social theory (Lave 1988), is how norms are reinforced and shaped through behavior. Each enacting of a norm (behavior pattern) potentially reinforces it for both 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. We might refer to understanding of norms as an individual's "social knowledge," but we must remain aware that 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.

Background, Theoretical Framework, and Modeling Methodology

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), 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 simultanenously on multiplelevels);
- 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 1997; 1999).
- Simulating collective behavior in a multi-agent simulation with an explicit "geographic model" of places and facilities.

To illustrate these ideas, we present 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 2001b). This Brahms simulation of this mission is broadly described in Clancey (2002b); here we focus on one part, the Brahms simulation of a planning meeting. We will see how people behave during the meeting (e.g., standing at the table during the meeting) exemplifies the nature of norms, and how 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 group decision making interact. We describe the methodology for constructing this model of practice, from video and first-hand observation, and how this modeling approach fundamentally changes how we relate goals, knowledge, and cognitive architecture. Following the methods of Schön (1987), we shift 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. We conclude that a 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.

The idea of modeling a planning meeting is very simple, but required a unique opportunity:

- Crew of six people living 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 (planning meeting, filling water tank, and preparing for EVA)
- The Brahms simulation was integrated with a graphic rendering of agent postures, movements, object manipulations, etc. in a system implemented in Adobe Atmosphere (Brahms Virtual Environment)
- 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).

This effort produced two basic surprises: 1) "Human factors" (e.g., hunger and fatigue), which are typically excluded from a cognitive model, drive activities of eating, resting, using the toilet, and recreation (e.g., playing games or talking at the table). 2) The 2 1/2-d (non-immersive) virtual display, which is typically viewed as just a "visualization tool," is essential for simulating line of sight and movement paths. We believe that 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.

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." They are cooperating, in sharing a resource (the facility), but not working on the same project.

The FMARS investigation, plus related work studying field scientists (Clancey in press b), 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 in press b).

"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.

Figure 1: The activity of working alone together, an example of cooperating without collaborating. FMARS initial habitation, August 2000.

Before examining the Brahms simulation in more detail, a few aspects need to 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).

- 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 as situation-action rules). However, Brahms models are first and foremost the investigator's models, not necessarily patterns articulated by the people being modeled (hence our emphasis that activities in Brahms are models of *conceptualizations*—which are largely non-verbal—not models of knowledge).
- Attitude, emotion, and personality are of fundamental importance, 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. We know that these "human factors" are essential for our application domain of long-duration space missions. We are using the FMARS data and simulation to understand what aspects of personality for example are relevant in understanding the crew's behavior.
- Each person's activities are broadly composed 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. We understand these identities to be dynamically 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 guides, not controllers of human action. (This was perhaps the most misunderstood insight of situated cognition research in the early 1990s, leading to absurd interpretations such as "there is no knowledge in the head") 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).

Simulation Model of Mars Crew Planning Meeting

Over the course of a week, an FMARS participant 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
ParicitpatingInPlanningMeeting
CoveringAgendaItemWeather
CoveringAgendaItemWater
AnnouncingEndOfPlanningMeeting
ConductingEVAPreparation
DonningSuit
DepressurizingInChamber
ConductingEVA
EatingSnack
TakingNap

Many details of 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).

Conventionally, a cognitive model is of an individual person's knowledge and reasoning, organizing around problem solving goals and inferences. An activity model is a kind of cognitive model, but organized around activities (what people do when and where), with conditional-actions (called workframes in Brahms), which specify subactivities or primitive actions (e.g., moving and communicating). A first glance, a Brahms model may appear not to be "cognitive" because there is no mention of goals, and actions are not reasoned about. Brahms uses familiar cognitive constructs, but relates and uses them in a different way:

- Activities are concepts, namely conceptualizations of what an agent is doing, e.g., Participating in a Planning Meeting.
- Activities are activated hierarchically in parallel, e.g., while Participating in a Planning Meeting an agent is also Living in the Mars Habitat, as well as Being a Computer Scientist.
- Workframes for each Activity on the current line of activation remain potentially active, such that interruptions may occur at higher levels to redirect attention (i.e., the relationship is not like a stack of procedure calls, but subsumption of concurrent processes).
- Perception is modeled by "detectables" associated with workframe actions; thus what the agent notices in the environment and how it is interpreted depends on the agent's current activity.
- The conditional part of a workframe matches the agent's beliefs; thus are all actions are dependent on beliefs. A belief is not necessarily an experienced (conscious) proposition (which the agent has or could easily articulate). Though it might be readily recognized and accepted as true if stated by someone else.

- The world is modeled as facts, interpreted by the modeler as what is objectively true; agents may believe any subset or variation of the facts.
- Agents may infer new beliefs through forward reasoning (called “thoughtframes”).
- Agents may receive new beliefs through communication with another agent or by reading them from objects (documents, displays).

The simulation of the planning meeting is a kind of cognitive model, though the present framework contains no examples of goal-directed reasoning. Developing Brahms in this way has been an experiment to determine to what extent purposeful, interactive behavior of a group can be simulated from conditional activity patterns alone.

The present model of the FMARS planning meeting does not attempt to replicate the details of how people plan what they do. Simulating how people plan their actions by reasoning about alternatives would require modeling goal-directed reasoning. As will become clear, we have many more problems to worry about in simulating a planning meeting, so we have chosen to leave out this level of detail. In this respect, 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 an apparently tight 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 dealing with the fire 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.

We have found that focusing on the agents’ changing activities, movements, and interactions with objects provides sufficient challenges, without getting into the discourse structure of the meeting’s conversation. We have omitted what a conventional cognitive model would include and vice versa: We have focused on postures (to understand what constrains them and what they convey), coordination of multiagent activity (e.g., how individual agents transition into a single group activity) and biological motives (e.g., hunger, fatigue). With this perspective, we have uncovered many interesting issues that shed a different light on what cognition accomplishes and how perception and action are related through conceptualization of activity—without needing to incorporate discourse modeling, planning, and goal-directed inference. These choice reflect a stage in Brahms’ development, and not a decision about what is ultimately required for advancing our theories or applications.

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.

Planning Meeting Time Lapse

Using methods developed over several field seasons of studying expeditions (Clancey 2001a), Clancey systematically recorded most of several days using a time-lapse apparatus. A quarter-frame (320x240 pixels) wide-angle view (Figure 2) was captured direct to computer disk every 3 seconds, such that the entire upper deck outside of the staterooms is visible. These frames are manually abstracted in a spreadsheet to show where people are (columns) at different times (rows)[5]. From this, statistics and graphs are generated. Meeting such as the morning planning meeting are often video-recorded in full, so the conversations can also be analyzed.

Figure 2: 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.

[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).

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.

---

[5] 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.

CC works on one project, his paper; SB has many problems with the satellite network, walkie-talkies, power, etc. to resolve.

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.

[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 3) reveals some surprising patterns and provides a basis for characterizing behavior in terms of norms.

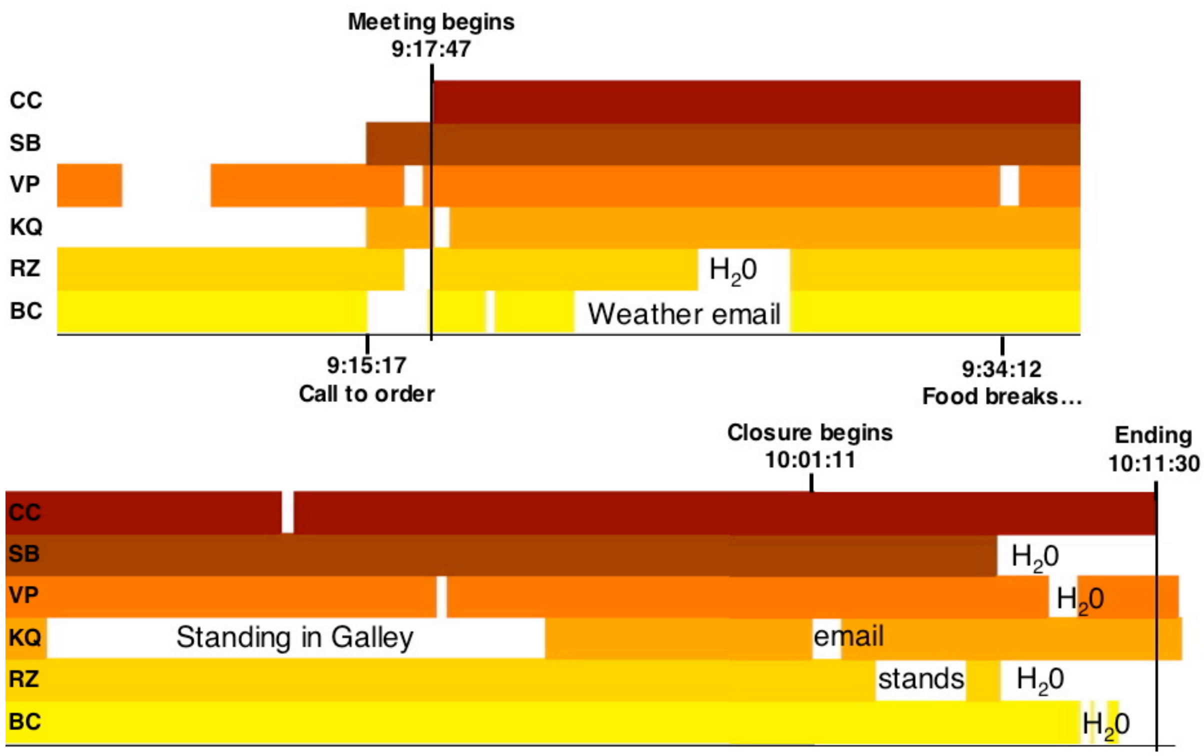

Figure 3: Location of crew members during planning meeting. The timeline is broken into two parts, starting at top left. Color indicates seated at table, otherwise the activity is indicated. See text for analysis.

To understand what we need 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. We would like the graphic simulation, which is a kind of cartoon, to appear plausible. 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 this never occurs. The graph shows what kinds of events are plausible, though may still be unexpected because people do not necessarily reflect on social behavior, even though it may be highly structured. Thus we observe 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, 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 meeting ends, although two people remain at the table.

In short, by modeling how individual agents carry out a group activity as conditional actions organized into activity conceptualizations, we are beginning to unpack how collective (social) behavior results from individual cognition (relating perception, motive, and action). We know a great deal more about social behavior than we have modeled in Brahms. Most importantly, 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). We do not represent this learning (i.e., reinforcement or adaptation of workframes) in Brahms. Furthermore, 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.

Planning Meeting Model Details

To create a model of the planning meeting, we (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 <u>located</u> ***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* <u>*workstation area*</u>*, to his* laptop *(specific subarea TBD), and checks weather report (for ~7 min; show as just* ***sitting*** *facing* laptop*). After* **BC** *is done, he* ***walks back to*** <u>*wardroom table area*</u>*, approaches his* <u>*chair area*</u>*, and* ***sits down*** *on his* chair. *He then communicates the weather data back to* **RZ**.
> *Shortly after* ***BC*** *goes to check the weather,* ***RZ gets up*** *from his* chair*,* ***walks*** *to*

<u>*water tank area*</u>*,* ***climbs*** *the* water tank ladder*, and* ***checks water level*** *(by looking into the* water tank *– just show him 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 font**):

```
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 sequence of subactivities in the **do** part are all 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 4).

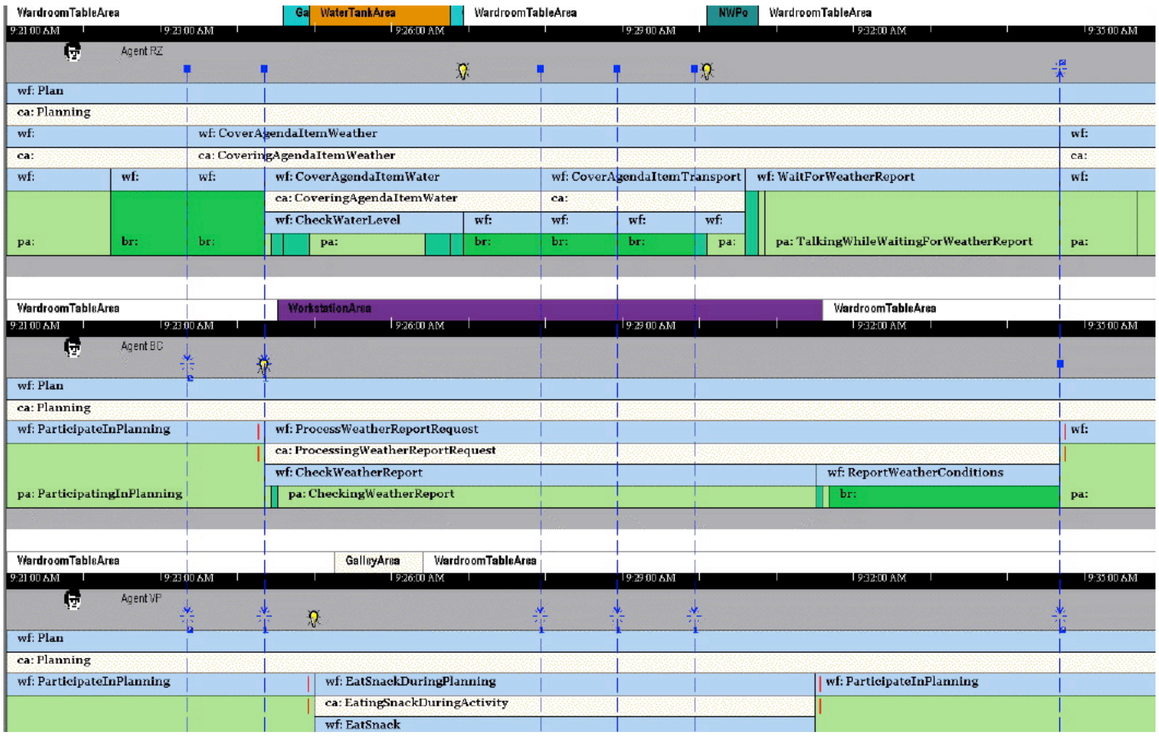

Figure 4: Brahms AgentViewer showing actions, communications, and inferences of agents RZ, BC, and VP during the first part of the planning meeting. At the time RZ does CheckWaterLevel, he is simultaneously engaged in Planning, Covering AgendaItemWeather and Covering AgendaItemWater.

While RZ is checking the water level, BC is checking the weather report. Figure 5 shows this moment graphically using the Brahms Virtual Environment (BrahmsVE).

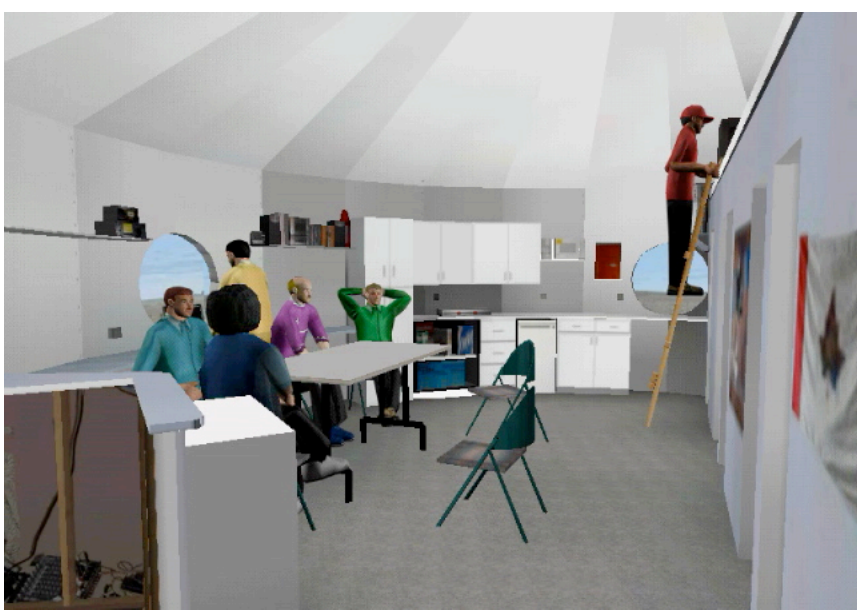

Figure 5: Frame 3:24 from animation showing RZ checking the water level while BC is reading the weather report at his workstation [9:25:19] Developed by Digital Space, Inc.

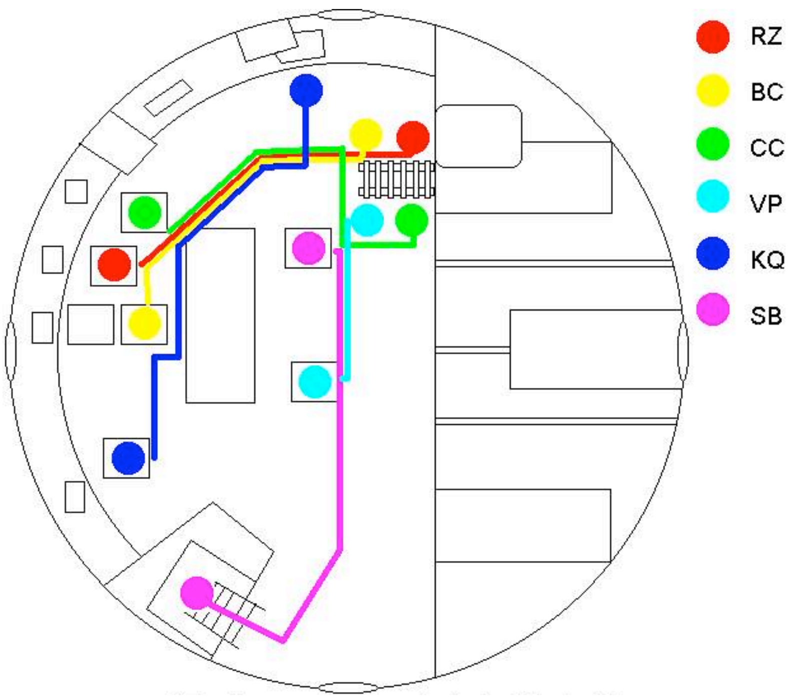

Figure 6: Initial storyboard showing ending of the planning meeting (Digital Space, Inc.)

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 in creating a cartoon or movie (Figure 6). 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.

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. One of the frames is shown in Figure 5.

All together, we simulated three complex FMARS scenarios 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 will all 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 5 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, we can now visualize postures and layout of the planning meeting. For example, we show how RZ sits alone on one long end of the table (Figure 2), which is not visible in the AgentViewer. In effect, we have represented part of the practice of the activity—the details of how people sit and move—in the graphic scripts. Applications of this interface and current work are discussed briefly in the final section.

Modeling Biological Motives and Behaviors vs. Goals

In developing a multiagent model of a day in the life of the FMARS crew, we were led to include biological drivers of behavior, such as fatigue, hunger, and the need to use the bathroom. Such aspects of human behavior are totally ignored by cognitive models, and yet are almost the exclusive concern of the area of psychology and design called *human* *factors* (e.g., Kantowitz and Sorkin 1983). 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 (most notably, one way get up to get something to eat during an FMARS meeting, but would only do this in a business office setting if the food were laid out for the participants in the room itself).

Biological needs are modeled in a simple way; our objectives do 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).

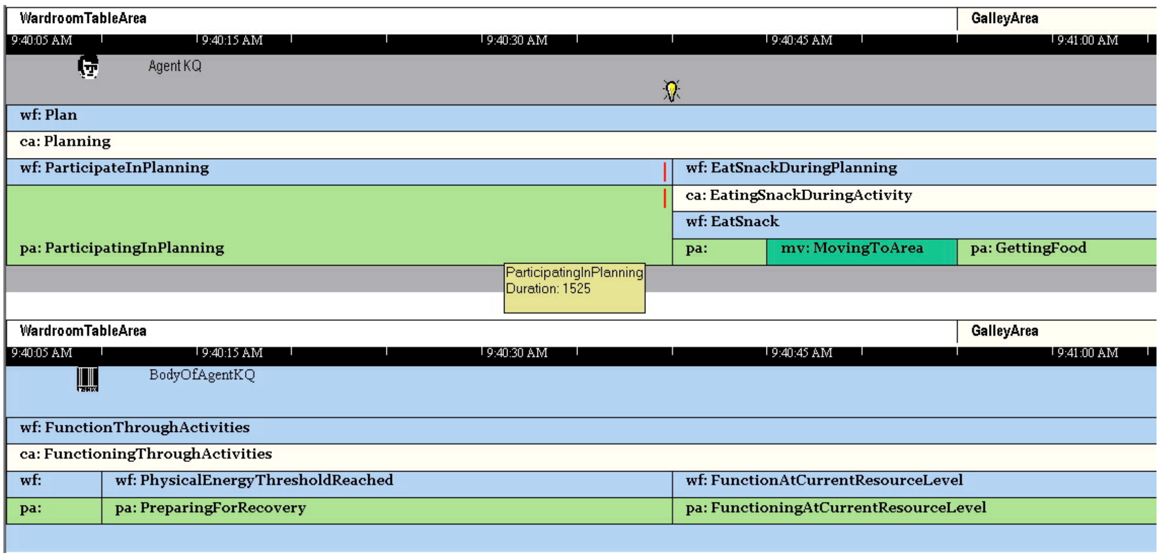

Figure 7: Brahms AgentViewer display of KQ Eating Snack during the planning meeting. This is a workframe within the Planning activity. Eating the snack is itself a composite activity, involving a sequence of conditional actions, including moving and getting food.

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 (Figure 2). 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. We don'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 (a simple example is how 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 action 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. 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). We conclude that it is highly problematic (if not theoretically impossible) to explain by goals actions that have not been deliberately planned. Instead, in a Brahms activity model we represent the context in which the action occurs and attempt to descriptively capture all (gross) movements, sequences, and communications. A goal analysis can always be imposed later, and certainly a task analysis is necessary for designing layouts, procedures, work flow tools, etc.

The purpose of Brahms is to represent the total system in way that reveals the relation of different levels of analysis (biological, psychological, social), to be contrasted with cognitive models that only represent the reasoning of a single person or multiagent models that focus on functional actions. As in omitting backward chaining, we are not in the least claiming a superior approach, but rather trying to understand what other representations can capture, given a focus on simulating the activities of a typical day, rather than automating a task. For example, Brahms' design was greatly 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 we were modeling fire-fighting in Brahms, we would as a heuristic model the entire day, including where the fire-fighters camp, how meals are prepared, how they are transported to the work site, etc. Thus, a day-in-the-life model complements task analyses, revealing how social activities are interleaved with and constrain how work is actually done. Understanding how individual beliefs, agendas, preferences, ambitions, etc. interact is a significant next step. In particular, the real intersection of cognitive and social analyses move us up a step—from what knowledge is required to accomplish a task, to whose knowledge is called into play. The question for cognitive modeling then becomes, what knowledge and behavior affects who is allowed to participate and by what means?

Conclusions: Lessons about Activity Modeling

We consider here 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 we can learn by reconsidering Newell's social band framework. We conclude with some remarks about applications of multiagent simulations like FMARS to failure analysis.

Use of the virtual environment interface

The most important finding about the graphic interface has been 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 multiagent model. That is, we relegate 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. Oddly enough, 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, as we have seen throughout Brahms modeling, the notion of a primitive activity is fully open, both to the purposes of the model (e.g., is fidelity in modeling the 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 Adobe Atmosphere enables collaborative design and engineering, by which a common virtual world (e.g., rendering 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.

Methodology of constructing a Brahms model

Our experiment of constructing a day-in-the-life FMARS model suggests that we view a Brahms model as a way of stating and organizing information about a work system. A Brahms model is a way of formalizing (expressing, collecting, and organizing) field observations so they can be correlated, shared, and used in software design. 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.

For example, after creating the model, we 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, we could verify whether his summary fits what was observed (including time-lapse data). He distinguishes between an EVA day and a sample analysis day, which would be useful in simulating multiple days (i.e., a typical day includes an EVA, but not everyone goes out every day). CC gives details about his scientific work that were not recorded (e.g., the names of his tools and their parts, and the lab equipment is in sequential order for sample processing), which could be added to our model. He says he performed a procedure 100 times in two weeks; assuming he is correct, the model should reflect this. He tells us that he sent images to a colleague, an activity that was not observed.

The idea of modeling "a day in the life" is a starting point. The day we have 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 over-arching 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, we could provide the statistics to Cockell for his own report. Related work by Clancey during NASA's Haughton-Mars Project during 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 an average of every minutes, while they estimated ten). 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 the analysis of the Haughton-Mars expedition in 1999 (Clancey 2001a), which showed 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 observing 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, the fundamental problem in constructing a model of social behavior is knowing what everyone is doing at all times. A Brahms 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, we believe at this time 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.

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 we know 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. We have a great deal to learn about this phenomenon, particularly how individuals are rationalizing their actions and where 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 agent's perception and hence do not affect simulated actions. We are working on closing 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 a complex project for the future, but is fundamental for relating cognition to social behavior. Figure 8 provides a glimmer of what could be involved.

Figure 8. Unusual posture at end of planning meeting. Square standing distribution suggests a balanced or stable relationship. Individuals apparently subconsciously move into and hold the encounter in this position. Possibly an important issue is being reconsidered.

Summary of relation between cognitive modeling and social interaction

To summarize the example and discussion to this point, we consider questions posed by Ron Sun for relating cognitive modeling and social interactions:

1) What are the suitable characteristics of cognitive architectures for modeling both individual cognitive agents and multi-agent interactions?
    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 8)?
   b. Ethnography (participant observation) is the fundamental way of understanding and modeling multi-agent interactions: photos, video, time-lapse, activity mapping (person, time, place)
   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
   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) How does individual cognition change and adapt in relation to social and physical interactions?
   a. Biological needs (fatigue, hunger, toilet, cold) affect choice of activity and manner of carrying it out.
   b. Perception of posture, attitude, tone of voice, etc. affect relationships (not included in Brahms)
   c. 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 structures and organizations in relation to individual cognition?
   a. See 2d
   b. In a multiagent 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 we arrange ourselves and interact, Figures 1, 2, and 8).

In constructing this outline, we are aware that it resembles more a list of examples than a comprehensive perspective, and already goes beyond what we have incorporated in the planning model. Thus at least from the perspective of our efforts it represents the edge of what we have considered and understand.

Relation to Newell’s Social Band Framework

One way of appraising our progress is to compare what we have done to Newell’s (1991) 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 is his framework based on systems levels he calls “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, we find that 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,” cf. Clancey 1999).

Possibly Newell viewed “social” as just meaning 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) 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).

Furthermore, 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 other kinds of models are possible, as we have shown. In particular, 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) simply 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 actions as determined by goals and knowledge seems bizarre when actual behaviors are considered in a day in the life of a group, such as the FMARS crew. The task-goal-knowledge analysis applies best when we view people at work, focusing on using laboratory equipment, downloading and analyzing EVA science data, or even preparing a meal. Put another way, at the time scale of 10 seconds or more—the "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 paper despite having the urge to use the bathroom), all behaviors are always in a social context. People frame their activity in terms of their socially constructed identities, and this determines what we do, when, where, and how, including what problems we discover or tackle and what methods we use to resolve them (Clancey 2002b).

Obviously, much more needs to be said in relating multiagent models like the FMARS Brahms model to Newell's analysis and other related work. However, for the present purposes it should be apparent that 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 occurring or rolling up over longer time scales, 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). Cognition—whether the person is physically alone or in a group—is immersed in norms and emergent physiological, physical, and cooperative constraints like a fish in water (Wynn 1991).

Application to failure analysis

The FMARS model has greatly clarified our understanding of the relation of task models to activity models. This topic is discussed at length in Clancey (2002b), including an extension of activity theory that takes into account aspects of perceptual-motor coupling and process dynamics (e.g., the goal of sustaining an activity, as opposed to arriving at a particular place or state). Properly relating task analysis to activity analysis is essential for using Brahms models in applications such as instruction and work systems design, especially failure analysis.

Of particular interest in applications is representing how people are actually conceiving of an 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 as a "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" architecture, rather than dynamically blending alternative coherent and sometimes conflicting ways of perceiving, interpreting, and acting.

In contrast with a conventional task analysis, focusing on goals and knowledge, an activity analysis emphasizes who is participating, where, when, and where (the context), and thus how problems are articulated, framed, and approached. In effect, from the perspective of attempting to evaluate and design and effective work systems (facilities, roles, schedules, procedures, tools, etc.), we are in some sense still focusing on knowledge, but shifting from considering only *what knowledge* is required to *whose knowledge* comes into play. This is an ironic twist, because it doesn't deny the central importance of cognition in the quality of the work product, but shifts from an idealized functional view of a problem to be solved, to a practical, realistic view of how people act within the physical and conceptual spaces of the group—insisting on a personal agenda or remaining silent as circumstances develop around them.

For NASA, the importance of a work practice analysis is highlighted in the *Columbia Accident Investigation Board Report* (2003), specifically in the actions of individuals within the key management meeting that considered and accepted a faulty tile damage analysis (basing decisions on previous interpretations of similar problems and scheduling constraints for subsequent launches). Perhaps one could say that 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 emotional pressures, and not just local reasoning about technical matters. Indeed, a social analysis is required to explain why existing information and knowledge was not brought to bear.

The Columbia disaster highlights how the group's roles, schedules, and even representational practices (e.g., PowerPoint bullets) 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 multiagent simulation, with multiple groups interacting over a day or more. Just as conventional task analysis works backwards from goals to knowledge, an activity-based critique 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 at what, when and where.

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 of 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 we have heretofore emphasized in understanding expertise and problem solving ability.

Acknowledgements

Our colleagues including especially Paul Feltovich, Chin Seah, Dave Rasmussen, and Mike Shafto, as well as other members of the Work Systems Design & 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. 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. Digital Space, Inc. has been funded through SBIR and STTR NASA contracts.

References

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., Sachs, P., Sierhuis, M., and van Hoof, R. 1998. Brahms: Simulating practice for work systems design. *Int. J. Human-Computer Studies*, 49: 831-865.

Clancey, W. J. 1999. Conceptual Coordination: How the Mind Orders Experience in Time., Hillsdale, NJ: Lawrence Erlbaum.

Clancey, W. J. 2001a. Field science ethnography: Methods for systematic observation on an arctic expedition. *Field Methods*, 13(3):223-243. August.

Clancey, W.J. 2001b. Simulating “Mars on Earth”—A Report from FMARS Phase 2. In F. Crossman and R. Zubrin (Eds.) *On to Mars: Colonizing a New World*, Apogee Books, 2002

Clancey, W. J. 2002a. A Closed Mars Analog Simulation: the Approach of Crew 5 at the Mars Desert Research Station. Mars Society Annual Conference, Denver.

Clancey, W. J. 2002b. Simulating activities: Relating motives, deliberation, and attentive coordination. *Cognitive Systems Research*, 3(3) 471-499.

Clancey, W. J. in press a. Automating CapCom: Pragmatic Operations and Technology Research for Human Exploration of Mars. In C. Cockell (ed.), *Martian Expedition Planning*. AAS Publishers.

Clancey, W. J. in press b. Roles for Agent Assistants in Field Science: Understanding Personal Projects and Collaboration. *IEEE Transactions on Systems, Man and Cybernetics*. Issue on Human-Robot Interaction.

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, August. Available at: 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. Exobiological Protocol and Laboratory for the Human Exploration of Mars—Lessons from a Polar Impact Crater, *Journal of the British Interplanetary Society*, 56 (3-4): 74-86 , Mar-Apr.

Damer, B. 1997. *Avatars: Exploring and building virtual worlds on the internet*. Berkeley: Peachpit Press.

Jordan, B. 1992. Technology and Social Interaction: Notes on the Achievement of Authoritative Knowledge in Complex Settings. IRL Technical Report No. IRL92-0027. 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: Cambridge University Press.

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.

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,* 17(5):32-41. Special issue on Human-Centered Computing at NASA, September/October.

Sierhuis, M., W.J. Clancey, Seah, C., and 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.

Suchman, L. A. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge: Cambridge Press.

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.* Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 45-64.