Multiagent Plan Execution and Work Practice

Modeling plans and practices onboard the ISS

Maarten Sierhuis
RIACS/USRA
NASA Ames Research Center
MS T35B-1, Moffett Field, CA 94035
msierhuis@mail.arc.nasa.gov

Alessandro Acquisti
SIMS/UC Berkeley & RIACS/USRA
102 South Hall
Berkeley, CA 94720
aacquisti@mail.arc.nasa.gov

William J. Clancey
IHMC/UWF
NASA Ames Research Center
MS 269-3, Moffett Field, CA 94035
bclancey@mail.arc.nasa.gov

Abstract

Numerous interdependent and uncertain constraints affect plan execution onboard a space ship. Plans are often invalid as they are being executed in the real world. Human work practices partly develop to deal with these realities. However, practices are difficult to study and represent within traditional planning tools. We discuss how modeling the work practices of the ISS Crew is used to develop a plan execution method that can deal with real world situations onboard the ISS. Brahms—a multiagent activity-based language—is used to model situated action and plan execution of human activities in practice.

1. Introduction

Over the last decade or so, work on planning and execution in agent-based systems has made steady progress using formal models (d'Inverno and Luck 2000) (Grosz et al. 1999) (Wooldridge and Jennings 1995) (Castelfranchi 1995) (Shoham 1993). Recently, these models have been combined in empirical systems, in which executable agent-based systems are developed and tested (Luck 1999) (Tambe and Zhang too appear). The DAI community has studied how the emergence of situation-specific events that are not part of the initial plan affects agent plan execution. (Durfee). It is well known that plans are often invalid at the moment they are being executed, because when plan execution starts, or soon after, the world has changed so much that the plan is already incongruent with reality. This has led to the development of partial global and lazy-skeletal planners in AI (Durfee 1988) (Freed 1998). However, most agent-based planning systems are limited in their ability to deal with real-world constraints on teamwork and situational awareness, and are mostly concerned with the planning and execution of the actions of one agent. Although in some of these systems the context of the plan being executed and the multiagent coordination is taken into consideration, only a limited worldview is taken into account during the execution of the plan, usually only the context and coordination relevant to the overall goal, ignoring other activities and intentions of the agents.

In this paper we describe a planning approach that includes how plans in multiagent systems are executed in practice in a complex, rich world. The world not only includes people and machines that at times interact, but also places, objects and artifacts that can change the world over time independently from the actions of the agents. Beyond that, we are interested in representing how the "daily" activities of people that lay outside the planned actions affect the plan execution. We have developed a multiagent modeling framework—Brahms (Clancey et al. 1998) (Sierhuis 2001)—that allows implementation of agent-based systems that execute multiagent plans, modeled after the practice of "plan execution" by people in real-world environments. Agents developed with the Brahms language are able to more flexibly deal with situation-specific world events that are independent from the plan. We call this *plan execution in practice*.

In (Acquisti et al. 2002) we present an agent-based model of the work practice onboard the International Space Station (ISS). This modeling effort forms the basis for a potential tool to assist NASA planners in their scheduling of the daily activities of the ISS crew and, in broader terms, in their planning of manned space missions. Our research has two functions: 1) to provide an artifact (i.e. a simulation model) that can help us study and understand the way work is done onboard the ISS; and 2) to use this artifact in planning, as well as to provide models of the work practice onboard the ISS to robotic assistants such as the PSA or the Robonaut (Ambrose et al. 2001) (Bradshaw et al. 2000). Thus, robotic assistants may have contextual awareness of the activities onboard the space station, allowing them to coordinate with the crew (Bradshaw et al. in press).

In the remainder of this paper, Section 2 discusses the model of the work onboard the ISS; Section 3 briefly introduces the Brahms programming language; and Section 4 discusses the use of our model for planning purposes and the challenges that arise from this task.

2. Modeling a Day in the Life Onboard the International Space Station

In a typical day, each ISS crewmember divides his or her time between physical exercise, maintenance, experiments, communication with ground personnel, personal time, and bio-needs activities (e.g., rest, eating). Some of these activities are critical for the well being of the crew. Hence, the planned maintenance and research activities must be scheduled around them. At the same time, several interdependent structural constraints must be met to ensure crew safety and productivity: thermal control, power management, communication bandwidth management, and regulation of other systems. These form a network of components that must be accurately timed and orchestrated around crew activities and needs.

Unlike other space missions, the ISS operates on "a continuous basis, with execution planning, logistics planning, and on-orbit operations occurring simultaneously for long periods of time" ((NASA 1999), p. 1.1-1). Planning for crew expeditions starts months or years ahead. As the expedition begins, just-in-time artifacts are prepared (such as the Onboard Short Term Plan, or OSTP, and Form 24) for execution on the ISS. Any unexpected event or discrepancy between the time allocated for a planned activity and the actual time required in the face of the realities of onboard life has far-reaching impacts on the completion and timeliness of crew activities, and therefore affects efficiency and productivity onboard. Such discrepancies are actually frequent, as the comparison between daily plans and actual ship logs shows. Based on our previous work, we state that in order to develop tools to improve planning and efficiency, we need to study how crew work practices emerge from planned activities and written procedures. Our research tries to understand how well the planned ISS activities and their written procedures fit the reality of onboard life, and more specifically, to determine the work practices that have evolved on the ISS since Expedition 1.

Table 1. Types of activities based on regularity and scheduling.

	Scheduled activity	**Unscheduled activity**
Day-specific activity	Maintenance activities (e.g., Replacement of urine-receptacle in Toilet) Experiments (e.g., LAB PL Status/Monitor)	Emergencies Job-Jar activities Unexpected maintenance or repair activities
Recurrent activity	Physical exercise. Daily Planning Conference. Eating (lunch, dinner, breakfast)	Going to the toilet Sending personal email

To deal with unexpected events and the realities of onboard life, we categorize activities according to the degree to which the activity was scheduled (scheduled vs. unscheduled activities) and the uniqueness or repeatability (day-specific vs. recurrent activities) of the activity (see Table 1). The two activity types, recurrent and day-specific, are represented differently in Brahms:

- *Recurrent activities* are represented at the group level in detailed activity plan-templates called *workframes*. By observing the work practice of actual crews, using crew videos, we abstract the practice that evolved during the mission into behavioral activity descriptions that can either be performed by interpreting a schedule (such as physical exercise), or reactively (such as going to the bathroom).
- *Day-specific* activities are represented in procedures (see section 4). Each agent has a plan-template for executing a prescribed procedure. Changes in the world are handled through reactive behavior and performing a just-in-time replanning activity[1] in which an agent changes the (mental) plan representation.[2]

3. The Brahms Language

Brahms is an agent-oriented language with a well-defined syntax and semantics. A Brahms model can be used to simulate human-machine systems, for what-if experiments, for training, "user models," or driving intelligent assistants and robots (Clancey et al. 1998) (Sierhuis 2001). The run-time component—the Brahms virtual machine—executes a Brahms model as part of a real time system, or as a simulation of agent and object behaviors.

The Brahms architecture is organized around the following representational constructs:

Groups of groups containing
  Agents who are located and have
    Beliefs that lead them to engage in
      Activities specified by
        Workframes
Workframes in turn consist of
  Preconditions of beliefs that lead to
    Actions, consisting of
      Communication Actions
      Movement actions
      Primitive Actions
      Other composite activities
    Consequences of new beliefs and facts
    Thoughtframes that consist of
      Preconditions and
      Consequences

Physical objects are represented as entities whose states change within workframes and thoughtframes; *conceptual objects* represent human conceptualizations (e.g., the idea of an "experiment").

Brahms is based on the theory of situated action (Suchman 1987) (Clancey 1997a). The activity framework,

---

[1] We do not suggest that astronauts perform this activity by executing a computational algorithm similar to artificial intelligence planning systems. We rather represent the astronaut's ability to change the order they decide to perform their activities, based on situational awareness and context.

[2] Agent plans are first constrained by the OSTP document and coordinated with Mission Control during the day.

which describes chronological *behaviors*, may be contrasted with the goal-driven framework in Soar and ACT-R, (Laird et al. 1987) (Anderson and Lebiere 1998), which functionally abstracts behavior in terms of tasks. Brahms offers to the researcher a tool to represent and study human behavior from the perspective of activity theory and "work practice" (Sierhuis 2001) (Clancey in press). A traditional task analysis of work especially leaves out informal logistics, such as how environmental conditions are physically detected (e.g., consider how conventional medical expert systems do not model how physicians perform a patient exam).

Without considering circumstantial factors, analysts cannot accurately model how work and information actually flow, thus they cannot adequately design software agents that help automate human tasks and coordinate with people. For these purposes, we need a model that includes aspects of reasoning found in an information-processing model, plus aspects of geography, agent movement, and physical changes to the environment found in a multi-agent simulation – such as interruptions, coordination, impasses. A model of work practice focuses on informal, circumstantial, and located behaviors by which synchronization occurs (such that the task contributions of humans and machines flow together to accomplish goals) and allows the researcher to capture (at least in part) the distinction in activity theory between motives, activities, and task-specific goals (Clancey 1997b) (Clancey in press).

4. Planning, Execution and Work Practice

In our analysis of the data gathered about life onboard the ISS, we looked for patterns in crew activities and emergence of work practices that are specific to onboard life. These include of course breakfast, lunch and dinner, personal hygiene, exercise, personal time and sleep, but also daily conferences with the ground and ground interventions providing support. We generalized and represented the individual astronaut's daily behavioral patterns as learned and shared activities at the (conceptual) group level. For example, the activity of eating breakfast onboard the ISS is represented at the ISS Crew group-level. This way, we represent that all agents that are a member of the ISS Crew perform this activity. The group structure also allows us to represent differences between social, cultural and other type of communities (for example, the behavioral differences between American and Russian crewmembers, and between male and female crewmembers.

We started by representing one particular day (May 7th, 2001), but soon found that we needed the ability to model any day. Thus, we explored and formalized a plan execution approach that allows the crew agents to perform any daily schedule, while at the same time allowing them to react to situational changes during the execution according to their work practice.

4.1 Daily Schedules

To make our model reusable and applicable to any typical day and scenario on the ISS, we represent daily schedules and procedures as objects that agents can access (i.e., read), have beliefs about, manipulate, and act upon. Daily plans are represented as Form-24[3] objects. Figure 1 shows a part of the morning activities on Form-24 for May 7th, 2001, including the replacement of the urine-receptacle in the toilet starting at 9:50am. Table 2 shows the Brahms source code of the ReplacementUrineReceptable activity, representing what is on Form-24 for May 7th, 2001 for that activity. The Form-24 object is a one-to-one representation of what is on the form, representing the actual information received by the astronauts. The form specifies each activity in sequence, with the time and who is to perform the activity.

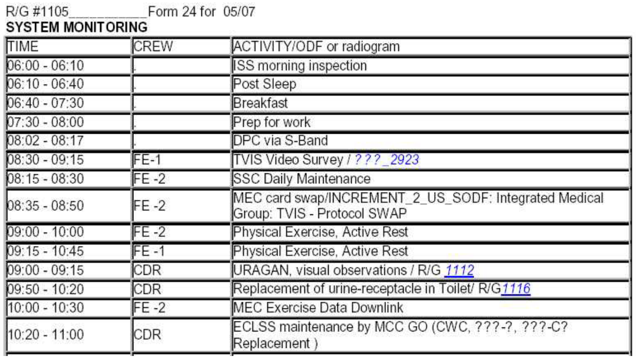

R/G #1105__________Form 24 for 05/07

SYSTEM MONITORING

TIME	CREW	ACTIVITY/ODF or radiogram
06:00 - 06:10	.	ISS morning inspection
06:10 - 06:40	.	Post Sleep
06:40 - 07:30	.	Breakfast
07:30 - 08:00	.	Prep for work
08:02 - 08:17	.	DPC via S-Band
08:30 - 09:15	FE-1	TVIS Video Survey / ? ? ? _2923
08:15 - 08:30	FE -2	SSC Daily Maintenance
08:35 - 08:50	FE -2	MEC card swap/INCREMENT_2_US_SODF: Integrated Medical Group: TVIS - Protocol SWAP
09:00 - 10:00	FE -2	Physical Exercise, Active Rest
09:15 - 10:45	FE -1	Physical Exercise, Active Rest
09:00 - 09:15	CDR	URAGAN, visual observations / R/G 1112
09:50 - 10:20	CDR	Replacement of urine-receptacle in Toilet/ R/G1116
10:00 - 10:30	FE -2	MEC Exercise Data Downlink
10:20 - 11:00	CDR	ECLSS maintenance by MCC GO (CWC, ???-?, ???-C? Replacement)

Figure 1. Form 24 for May 7th, 2001

In the model, the schedule is a *document object*, which a JSC planner officer agent uploads to the station computers. The crew agents access this document object through their laptops. By reading the information in the document (i.e., performing a *communicate activity*) every crew agent receives the information about the schedule in the form of individual beliefs about the activities of the day.

Table 2. Brahms source code of the *ReplacementUrineReceptable* activity (compare to Figure 1)

```
Object form_24_for_May_7_2001_Expedition_2 instance
of DailySchedule {
   initial_beliefs:
     [...]
    (ReplacementUrineReceptable.hour_start = 9);
    (ReplacementUrineReceptable.minute_start = 50);
    (ReplacementUrineReceptable.hour_end = 10);
    (ReplacementUrineReceptable.minute_end = 20);
    (ReplacementUrineReceptable.duration = 1800);    // in
seconds
    (ReplacementUrineReceptable by_whom YuriUsachev);
    (ReplacementUrineReceptable.Cmd_next_activity =
```

---

[3] Form 24 is a Russian form that was created for daily crew schedules onboard MIR. This form is still in use onboard the ISS. The Americans have a more elaborate electronic version of the schedule, called the OSTP.

```
ECLSSMaintenance);
    [...]
```

4.2 Procedures

Every scheduled daily activity has procedures for performing it. We represent the procedures as objects in Brahms. Figure 2 shows the crew's procedure for replacing the urine collection tank; Table 3 shows the Brahms representation. In the model, agents access the procedure objects when needed and read the procedures by performing a communication activity, in which the procedure data are transferred as beliefs to the agent.

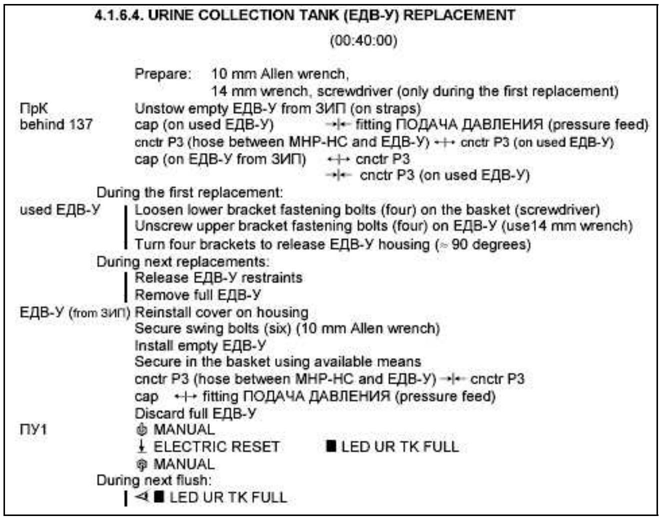

4.1.6.4. URINE COLLECTION TANK (ЕДВ-У) REPLACEMENT
(00:40:00)
Prepare: 10 mm Allen wrench,
14 mm wrench, screwdriver (only during the first replacement)
ПрК Unstow empty ЕДВ-У from ЗИП (on straps)
behind 137 cap (on used ЕДВ-У) →|← fitting ПОДАЧА ДАВЛЕНИЯ (pressure feed)
cnctr P3 (hose between МНР-НС and ЕДВ-У) ←|→ cnctr P3 (on used ЕДВ-У)
cap (on ЕДВ-У from ЗИП) ←|→ cnctr P3
→|← cnctr P3 (on used ЕДВ-У)
During the first replacement:
used ЕДВ-У | Loosen lower bracket fastening bolts (four) on the basket (screwdriver)
Unscrew upper bracket fastening bolts (four) on ЕДВ-У (use14 mm wrench)
Turn four brackets to release ЕДВ-У housing (≈ 90 degrees)
During next replacements:
| Release ЕДВ-У restraints
Remove full ЕДВ-У
ЕДВ-У (from ЗИП) Reinstall cover on housing
Secure swing bolts (six) (10 mm Allen wrench)
Install empty ЕДВ-У
Secure in the basket using available means
cnctr P3 (hose between МНР-НС and ЕДВ-У) →|← cnctr P3
cap ←|→ fitting ПОДАЧА ДАВЛЕНИЯ (pressure feed)
Discard full ЕДВ-У
ПУ1 MANUAL
ELECTRIC RESET ■ LED UR TK FULL
MANUAL
During next flush:
| ◁ ■ LED UR TK FULL

Figure 2. Expedition 2 Procedure for urine collection tank replacement

We have not yet found a complete generic representation for all written procedures as Brahms objects, because the ISS procedures are not written using one standard format. JSC is currently trying to standardize procedures using an XML data definition description (DTD). The modeling effort will of course be simplified if this standardization succeeds.

We represent procedures in Brahms as follows: Each procedure object specifies the *main activity* (e.g., *current.main_activity = ReplaceUrinalReceptacle* in Table 3), thus correlating the procedure to an activity on the Form-24 object. Using this relation, the crew agent knows which procedure to execute for a given activity on the schedule. Next, the procedure object represents the type of activity. A *composite activity* [(*current.main_activity_type = composite_act)]* means that this procedure is decomposed into sub-activities that are further described in the procedure object. A composite sub-activity again has an associated procedure object. In contrast, a *primitive activity* (e.g., the sub-activity *PrepareTools* in Table 3) is described inside the procedure object.

A primitive activity in a procedure has standard attributes: what resources are needed to perform the activity (*Screwdriver1*, *10mmAllenWrench*, *14mmWrench*), in what location the agent has to be to start executing the procedure (*ToiletArea*), and what the next sub-activity of the procedure is, after the current activity is finished (*UnstowEmptyContainer*).

Using the sub-activity descriptions in the procedure object, a crew agent can execute the scheduled activity on the Form-24. How the agents do this is explained in the next section.

Table 3. Brahms representation of urine collection procedure (compare to Figure 3)

```
object UrineCollectionTankReplacementProcedure
instanceof Procedure {
   initial_beliefs:
   (current.main_activity = ReplaceUrinalReceptacle);
   (current.main_activity_type = composite_act);
   (current has_sub_activity PrepareTools);
   (current has_sub_activity UnstowEmptyContainer);
   (current has_sub_activity CapOnUsedUrineContainer);
     [...]
   (ReplaceUrinalReceptacle first_sub_activity
PrepareTools);
   (PrepareTools.type_act = primitive_act);
   (PrepareTools object_needed Screwdriver1);
   (PrepareTools object_needed 10mmAllenWrench);
   (PrepareTools object_needed 14mmWrench);
   (PrepareTools.where_performed = ToiletArea);
   (PrepareTools next_sub_activity
UnstowEmptyContainer);
   [...]
} // UrineCollectionTankReplacementProcedure
```

4.3 Plan Execution

Plan execution in the model is generalized using an *activity plan template* (see Figure 3). The template represents when and how to execute an OSTP activity from the daily schedule. When the (current) agent knows (i.e., has a belief that matches a condition of the template) what OSTP activity he is supposed to do at that moment, and he knows what procedure describes how to perform the activity, then the agent concludes that he is doing the activity and immediately starts to execute the generic *doOSTPActivity* activity.

Figure 3 shows the template represented as a Brahms workframe in the ISSCrew group. Every ISS crewmember agent is a member of the ISSCrew group and thus inherits this workframe. To understand how this workframe is applied, the reader must be familiar with how Brahms' belief-driven engine operates. The engine schedules an agent's current workframe to be executed based on an efficient Reasoning State Network (RSN) of workframes, preconditions and activities. At any moment during execution, an agent has only one current workframe being executed, and a list of unavailable, available, interrupted, and impassed workframes. The engine schedules the current workframe at every engine clock-tick from the highest priority activities in the set of current, available, and interrupted workframes (Sierhuis 2001).

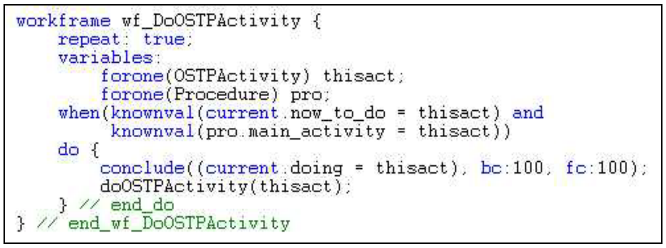

```
workframe wf_DoOSTPActivity {
    repeat: true;
    variables:
        forone(OSTPActivity) thisact;
        forone(Procedure) pro;
    when(knownval(current.now_to_do = thisact) and
         knownval(pro.main_activity = thisact))
    do {
        conclude((current.doing = thisact), bc:100, fc:100);
        doOSTPActivity(thisact);
    } // end_do
} // end_wf_DoOSTPActivity
```

Figure 3. ISS Crew Agent Activity Plan Template

When an agent executes the *doOSTPActity*, the agent executes a *composite activity*. Figure 4 represents the composite activity *doOSTPActity*.[4] The model states that when an agent knows what scheduled OSTP activity to perform, the agent first needs to retrieve the procedure, and read it to determine how to execute the OSTP activity, using the *wf_retrieveProcedure* plan template. After the agent has done this, the agent executes every procedure step, using the *wf_executeProcedureActivity* plan template, based on the description of the sub-activities on the procedure object (e.g. Figure 2). To decide what next activity in the procedure to execute, the agent uses a set of production rules (e.g., thoughtframe *tf_decideNextSubActivity*). These thoughtframes are executed when the end-condition for the current procedure activity is met.

4.4 Integrating Work Practice with Planning

A question arises from the above representation: How adequately does Brahms model work practice if it is naturally improvised and involves learning? If the question is taken to mean "can human activities be predicted in detail", the answer is no. Brahms models represent patterns, norms of behavior; Brahms agents do not mimic human flexibility in detail. It is not possible using this framework to represent the multitude of factors that affect human behavior, even in a relatively controlled environment such as the ISS or a manned space mission.

The model describes situated behaviors (referring to time, location, detected/perceptual properties of objects, group beliefs, and communications), but cannot replicate the flexibility of human behavior in all its complexity, which involves breaking patterns, and thus establishing new practices. In addition, because Brahms does not model human reasoning and learning, the simulation depends heavily on initial conditions such as the models attributed to the agents about procedures and—what we have observed to be—their work practice. In the absence of ground intervention, the simulated agents would not be able to find new solutions to unexpected problems. Fortunately, this is not an important problem in modeling the ISS because the practice is to seek detailed advice from the ground. For a Mars surface simulation, where time delay prevents such conversations, we would have a greater need to model how agents learn from available resources. Nevertheless, the model can capture *routine adaptations*. The analysis of the data and the comparison between planned activities and daily logs highlight frequent, and up to a point, regular discrepancies between the plan and the practice (cf. Table 4). The discrepancies we refer to are not only those caused by imprecise timing of new activities, or triggered by unforeseeable error and mismatches with systems or procedures.[5] Rather, as Table 4 shows, we also consider more substantial discrepancies involving a deliberate (though possibly not planned-in-advance) behavior of the crew.

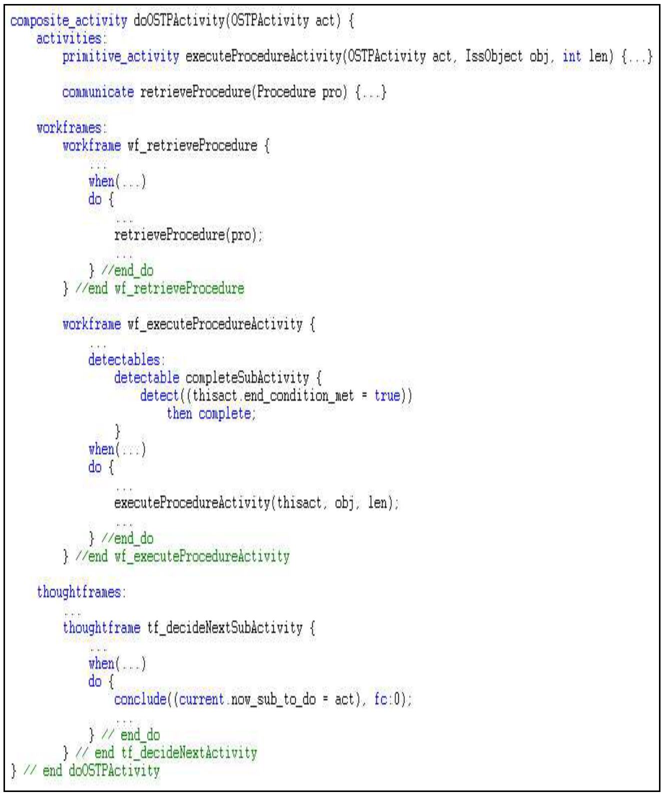

```
composite_activity doOSTPActivity(OSTPActivity act) {
    activities:
        primitive_activity executeProcedureActivity(OSTPActivity act, IssObject obj, int len) {...}

        communicate retrieveProcedure(Procedure pro) {...}

    workframes:
        workframe wf_retrieveProcedure {
            ...
            when(...)
            do {
                ...
                retrieveProcedure(pro);
                ...
            } //end_do
        } //end wf_retrieveProcedure

        workframe wf_executeProcedureActivity {
            ...
            detectables:
                detectable completeSubActivity {
                    detect((thisact.end_condition_met = true))
                        then complete;
                }
            when(...)
            do {
                ...
                executeProcedureActivity(thisact, obj, len);
                ...
            } //end_do
        } //end wf_executeProcedureActivity

    thoughtframes:
        ...
        thoughtframe tf_decideNextSubActivity {
            ...
            when(...)
            do {
                conclude((current.now_sub_to_do = act), fc:0);
                ...
            } // end_do
        } // end tf_decideNextActivity
} // end doOSTPActivity
```

Figure 4. *doOSTPActivity* activity

A traditional planning approach typically does not take into consideration some of the items highlighted in Table 4 or the concatenated circumstantial effects caused by the highlighted discrepancies. In contrast to typical planning approaches, by virtue of representing behaviors and not just abstracted tasks, the Brahms simulation is capable of showing how the practice of onboard activities often

---

[4] The hierarchy of active activities is handled through a variation of the subsumption mechanism, allowing agents to be in multiple subsumed activities at the same time Brooks, R., A. (1991). "Intelligence without representation." *Artificial Intelligence*, 47, 139-159. Thus workframes and thoughtrames on conceptually higher levels may change the agent's behavior. In particular, the activity of an agent interpreting a procedure should be contrasted with procedure invocation in a computer program.

[5] In this regard, Expedition 2 reported a substantial improvement with respect to Expedition 1 in the accuracy of the predicted duration of scheduled activities and in the feasibility of the planned daily workload.

diverges, both in timing and execution, from the originally scheduled activities and procedures. Distances and movements, noises, tools location, work practice, and so forth are considered. Hence, delays caused by crew movement constraints, the search for tools and other items, and the inability to share resources or access to electronic procedures can be discovered from the simulation. For example, we model that the work practice of the astronauts is to move from one module to the other to communicate face-to-face, rather than using the internal audio system.

Table 4- Discrepancies between plan and practice[6]

Cause of Discrepancy	Example	
	Plan	***Practice***
Procedures not easily accessible or not clear	During emergency, refer to procedure	During emergency, rely on training and memory
Noise level on internal audio system	Use internal audio system to communicate between modules	Move from module to module to communicate with crew members
Personal preferences	Do medical tests as scheduled	Do medical test in the morning
Shared resources not always available	Upload on laptop computer medical/physical data after experiment/exercise	Upload data rarely
Personal habits	Read procedure	Read electronic procedure (from laptop), or read printed procedure
Inventory system not always reliable	Use tools indicated in procedure	Tools must be found and time can be lost in this operation
Inventory system not always reliable	Use bar-code reader for inventory	Rarely use bar-code reader

From a modeling perspective, the primary challenge we have addressed is how to represent the schedule *and* the practice of following a schedule. Work practice might diverge from a procedure in several ways:

- Work practices specific to certain procedures might emerge. That is, the crew has learned how to perform the work, and what they do is not documented (yet) in the procedure. By observations of videos, debriefs, and interviews, we can note that certain activities are regularly executed in a particular way, possibly not the exact way described in the procedure.
- Other events not described in the procedures might occur: interruptions, delays (such as the time spent looking for tools), and so on. These events could be defined as work practice not specific to any particular activity (e.g., a preferred location that a crew member uses for storing the tools needed for a task).
- Errors and failures might arise. These might have well-define statistical properties that we can observe from data and insert in the model.
- In addition to this, we are currently working on the meta-level of "*just-in-time replanning activity*", where the simulated agent, noting that a certain activity scheduled for a certain time and length is taking more (or less) than the time allocated for it, has to decide in real time the next step in his execution plan, which might include contacting mission controllers for advice.

Table 5. Thoughtframe representing learning that deviates from the procedure

```
thoughtframe tf_ WorkPracticeReplaceUrinalReceptacle {
   when(knownval(current.doing = ReplaceUrinalReceptacle))
   do {
    conclude((PrepareTools object_needed 14mmWrench is
false), fc:0);
    conclude((PrepareTools object_needed 12mmWrench is
true), fc:0);
    }
}
```

To summarize, work practice can be represented on top of the procedures in one of the ways described in Table 6. To offer a hypothetical example, while the Urinal Container Replacement procedure might imply that a certain step is always followed by another step, the agent/astronauts might have a work practice (represented as a thoughtframe in Table 5) that overrides the existing beliefs (for example, about which tools are needed) that the agent received from reading or memorizing the procedure.

In short, while formally modeling work practice might sound like a contradiction in terms, the combination of data-based work practice study and an agent-based simulation approach can provide a powerful tool to incorporate *some* useful aspects of the daily life of the crew onboard a space station.

5. Conclusions

On the basis of our ongoing modeling effort of the ISS crew work practice, we have discussed the use of a Brahms model for planning and scheduling. The Brahms model of a day in the life of the ISS crew is not hard-coded, in the sense that the model does not represent a single specific

---

[6] Sources: ISS Ship logs; Expedition 2 debriefs; interviews with ISS training specialists.

day. Instead, we can simulate any typical day by providing a different daily plan as input to the model. The combination of a work practice based analysis of the crew activities, and an agent-based approach to their representation offers powerful instruments, both for studying, and then influencing, human activities in manned space missions. Consequentially, our ongoing efforts to model emergency scenarios might be useful to predict ISS crew behaviors and their outcomes. In our continued research we are also exploring the use of the ISS model as part of an environment for teamwork between ISS crews and onboard software assistants and robotic systems, and as a short term planning and scheduling tool for mission planners.

Table 6. Approach to Planning Work Practice

1. Gather information about procedures, work practices, and statistics on failures/errors/delays
2. Represent procedures in Brahms
3. Insert procedures in Brahms model, and let agents/astronauts read them
4. Model agents' execution of procedures. Deviations will occur:
    a. Procedure specific deviations (e.g., work practice specific to a certain activity) – *override procedure?*
    b. General crew work practice (e.g., work practice that emerge independently from a specific activity: leaving or looking for tools in a certain location, interacting with crew members passing by, etc.) – *mix with procedure?*
    c. Independent events (human errors, machine failures for which we have statistics – see 1) – *mix with procedure and requires replanning?*
    d. Just in time replanning
5. Execute several runs of model, examine results, find regularities and determine sensitivity
6. On the base of 1,2, and 5, make library of "abstract" procedures and library of "actual" statistical procedures, to be used by planners in coordination with Brahms.

Acknowledgements

This work is being funded by the NASA Intelligent Systems Program in the area of Human-Centered Computing. We acknowledge the many contributions of other members of related projects: J. Bradshaw, J. Brodsky, R. van Hoof , R. Jeffers, D. Prescott, M. Shafto, N. Suri, and A. Uszok.

References

Acquisti, A., Sierhuis, M., Clancey, W. J., and Bradshaw, J. M. (2002). "Agent Based Modeling of Collaboration and Work Practices Onboard the International Space Station." *11th Computer-Generated Forces and Behavior Representation Conference*, Orlando, Fl.

Ambrose, R., Culbert, C., and Rehnmark, F (2001). "An Experimental Investigation of Dexterous Robots Using Eva Tools and Interfaces." *AIAA*.

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

Bradshaw, J. M., M.Sierhuis, Y.Gawdiak, Jeffers, R., Suri, N., and Greaves, M. (in press). "Adjustable Autonomy and Teamwork for the Personal Satellite Assistant." Agent Autonomy, H. Hexmoor and R. Flacone, eds., Kluwer.

Bradshaw, J. M., Sierhuis, M., Gawdiak, Y., Thomas, H., Greaves, M., and Clancey, W. J. (2001). "Human-Centered Design for the Personal Satellite Assistant." *International Conference on Human-Computer Interaction in Aeronautics 2000*, Toulouse, France.

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

Castelfranchi, C. (1995). "Guaranties for autonomy in cognitive agent architecture." Intelligent Agents: Theories, Architectures, and Languages, M. Wooldridge and N. R. Jennings, eds., Springer-Verlag, Heidelberg, Germany, 56-70.

Clancey, W., J. (1997a). *Situated Cognition: On Human Knowledge and Computer Representations*, Cambridge University Press.

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

Clancey, W. J. (in press). "Simulating Activities: Relating Motives, Deliberation, and Attentive Coordination." *Cognitive Systems Review*, special issues on Situated Cognition.

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

d'Inverno, M., and Luck, M. (2000). "Formal Agent Development: Framework to System." Formal Approaches to Agent-Based Systems, J. L. Rash, C. A. Rouff, W. Truszkoswski, D. Gordon, and M. G. Hinchey, eds., Springer-Verlag, Greenbelt, MD, 133-147.

Durfee, E. H. (1999). "Distributed Problem Solving and Planning." Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, G. Weiss, ed., The MIT Press, Cambridge, MA, 121-164.

Durfee, E. H. (1988). *Coordination of Distributed Problem Solvers*, Kluwer Academic Press, Boston.

Freed, M. (1998). "Simulating Human Behavior in Complex, Dynamic Environments," Ph.D. Thesis, Northwestern University.

Grosz, B. J., L. Hunsberger, and Kraus, S. (1999). "Planning and Acting Together." *AI Magazine*, 20(4), 23-34.

Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). "Soar: An architecture for general intelligence." *Artificial Intelligence*, 33, 1-64.

Luck, M. (1999). "From definition to deployment: What next for agent-based systems?" *The Knowledge Engineering Review*, 14(2), 119-124.

NASA. (1999). ""International Space Station - Operations and Planning Training Manual" ISS OPS & PL TM 21109, TD 9711, Rev. A." Johnson Space Center, Houston, TX.

Shoham, Y. (1993). "Agent-oriented programming." *Artificial Intelligence*, 60(1), 51-92.

Sierhuis, M. (2001). "Modeling and Simulating Work Practice; Brahms: A multiagent modeling and simulation language for work system analysis and design," Ph.D. thesis, University of Amsterdam, Amsterdam, The Netherlands.

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

Tambe, M., and Zhang, W. (too appear). "Towards flexible teamwork in persistent teams: extended report." *Journal of Autonomous Agents and Multi-agent Systems, special issue on "Best of ICMAS 98"*.

Wooldridge, M., and Jennings, N. R. (1995). "Intelligent Agents: Theory and Practice." *Knowledge Engineering Review*.