Patterns in cooperative cognition

Emily S. Patterson, David D. Woods
Cognitive Systems Engineering Laboratory
Institute for Ergonomics, Ohio State University
210 Baker Systems, 1971 Neil Ave., Columbus, OH 43210
Patterson.150@osu.edu
woods@csel.eng.ohio-state.edu

Nadine B. Sarter
Institute of Aviation – Aviation Research Laboratory
University of Illinois at Urbana-Champaign
nsarter@uiuc.edu

Jennifer Watts-Perotti
Cognitive Systems Engineering Laboratory
perotti@kodak.com

ABSTRACT

In this paper, seven studies of cooperative cognition in complex operational settings conducted by members of the Cognitive Systems Engineering Laboratory (CSEL) are reviewed. These studies were conducted using a variety of methodologies, including naturalistic observations as well as more controlled investigations using scenario-based simulations. Six converging patterns that were observed across these studies are synthesized. These patterns are: 1) breakdowns in coordination that are signaled by surprise, 2) escalations in activities following an unexpected event in a monitored process, 3) investments in shared understandings to facilitate effective communication, 4) local actors adapting original plans created by remote supervisors to cope with unexpected events, 5) calling in additional personnel when unexpected situations arise, and 6) functional distributions of cognitive processes during anomaly response. These patterns further our understanding of the fundamentally cooperative nature of cognition and provide insight for innovative design.

KEYWORDS

anomaly response, common ground, escalation, planning, updating.

1. INTRODUCTION

Traditionally, studies in cognitive psychology have focused on individual cognition, where a subject does not have the ability to use resources that would normally be available in a real-world context, such as other people, machine artifacts, procedures, or experience (Hutchins, 1995). On the other hand, traditional studies of group problem solving are run with a small, homogeneous, co-located group that is given a clearly defined problem to solve. Over the past few decades, there has been a growing movement among many researchers, including members of the Cognitive Systems Engineering Laboratory (CSEL), to extend these perspectives by studying distributed cognition in complex operational settings such as aircraft cockpits, space shuttle mission control centers, and medical operating rooms.

In this paper, the results from seven recent CSEL studies (Table 1) are reviewed and synthesized.

The first series of studies (Sarter and Woods, 1992, 1994, 1997a, 1997b; see Sarter, Woods, and Billings, 1997 for a synthesis) investigated automation surprises in pilot interaction with advanced aircraft automation. This series of studies built upon the recognition by Wiener (1980) that pilots were sometimes surprised by the actions of aircraft automation. Sarter and Woods collected corpuses of cases by asking line pilots to describe situations where the automation surprised them. They then manipulated the factors identified during the first studies in a high-fidelity simulation with experienced pilots, using carefully crafted scenarios to elicit automation surprises.

The second series of studies investigated various aspects of how practitioners coordinate during anomaly response. Johannesen, Cook, and Woods (1994) analyzed real-time coordination and communication in the operating room (direct observation of neuro-anesthesia teams), where practitioners were observed to invest in creating a ‘common ground’ for communications before anomalies occurred so that they could coordinate effectively when they had to respond to an anomaly. Building on this research, targeted observations were conducted in space shuttle mission control at NASA Johnson Space Center during simulated and actual shuttle operations. The results capture the cognitive activities underlying coordination in updating or bringing up to speed incoming personnel (Patterson and Woods, 1997) and how interacting functionally distributed teams are able to meet the demands of and avoid failures in anomaly response (Watts, Woods, and Patterson, 1996, Watts and Woods, 1997). The role of the voice loops technology in supporting this coordination was also investigated (Watts *et al.*, 1996b, see also Woods, 1995, Malin *et al.*, 1991).

The remaining studies (Shattuck and Woods, 1997, Dekker and Woods, 1997) examined teams involving local actors and distant supervisors and how they coordinate when the local situation departs from the original plan. Shattuck and Woods (1997) designed scenarios where actual events went beyond pre-planned guidance (either impasses or new opportunities). Commanders (distant supervisors) communicated a plan to subordinates (local actors). In addition to the plan, the commanders also communicated their intent, which was supposed to help the subordinate adapt the plan to handle anomalies. In a simulated situation, both subordinates and commanders were observed as they adapted plans to handle unexpected anomalies based on the commander’s intent. Dekker and Woods (1997) examined how a distant supervisor decided whether or not to intervene and take back authority from another agent in a deteriorating situation -- management by exception in supervisory control. The context was a new envisioned world in air traffic management with a new distribution of authority across controllers, pilots and flight dispatchers.

Table 1: Recent CSEL Studies of Cooperative Cognition

Reference	**Domain**	**Method**	**Theme**
Sarter and Woods (1992, 1994, 1997a, 1997b)	Cockpit aviation	Corpus of cases; Scenario-based Simulation	Automation surprises in human-machine coordination
Patterson and Woods (1997)	Space Shuttle Mission Control	Direct Observations in the Field	Updating during shift changes and the on-call model for intervention
Watts *et al.* (1996a, 1997)	Space Shuttle Mission Control	Direct Observations in the Field	Coordination across functional distributed teams in anomaly response
Watts *et al.* (1996b)	Space Shuttle Mission Control	Direct Observations in the Field	Auditory CSCW technology that mediates the common ground
Johannesen, Cook, and Woods (1994)	Anesthesiology	Direct Observations in the Field	Calibrating the team’s common ground before an anomaly occurs
Shattuck and Woods (1997)	Command and Control	Scenario-based Simulation	Communication of intent from supervisors to local actors
Dekker and Woods (1997)	Air Traffic Management	Scenario-based Simulation	Management by exception as a cooperative architecture

Across these studies from different domains and using different methodologies, several patterns in cooperative cognition have emerged. These patterns are:

- how coordination breakdowns are signaled by an agent that is surprised by the behavior of other agents or the underlying monitored process that is being influenced by another agent,
- how anomalies in a monitored process produce an escalation in cognitive and coordinative demands which bring out the penalties of poor support for coordination,
- how investments in shared understandings before anomalies arise facilitate effective coordination in responding to anomalies, e.g., prevent coordination surprises,
- how to balance flexibility and planning when local actors need to adapt plans created by distant supervisors to cope with unexpected situations,
- how to effectively update and integrate additional personnel called in when unexpected situations arise so that they can coordinate as if they had been present from the beginning of the trouble, and
- how coordination across functionally distributed teams is a robust architecture for meeting the demands and avoiding failure in anomaly response.

2. COORDINATION SURPRISES

In the aviation domain, we have identified and shaped conditions that give rise to what we have referred to as ‘automation surprises’ (Sarter and Woods, 1992; 1994; 1997a; 1997b). These are situations where practitioners are surprised by actions taken (or not taken) by machine agents, such as automation in computerized ‘glass’ cockpits. Automation surprises begin with miscommunication and misassessments between the automation and users which lead to a gap between the user’s understanding of what the automated systems are set up to do, what they are doing, and what they are going to do. As a result, the automation ‘flies’ the aircraft into trouble and the human supervisor is unaware that this is happening/will happen until problems arise.

The evidence shows that the potential for automation surprises is the greatest when three factors converge:

1. the automation can act on its own without immediately preceding directions from its human partner (high autonomy and authority),
2. there are gaps in users’ mental models of how their machine partners work in different situations, and
3. there is weak feedback about the activities and future behavior of the machine agent (low observability).

Parallels to ‘automation surprises’ have been observed in human-human interactions. These observations point to a broader category of cooperative interaction, which we refer to as ‘coordination surprises,’ of which automation surprises are a subset that occur in interactions between practitioners and automation. In coordination surprises, an agent is surprised by the way another agent acts on the distributed cooperative system. A human or machine agent can directly perform an activity that is surprising to an agent. Alternatively, agents can perform activities that affect the underlying monitored process or the coordination of other distributed agents in the system in ways that are not anticipated by the agent who is surprised. The mismatch situation between an agent’s expectations and the situation is believed to result from the convergence of the same three factors: high autonomy and authority, gaps in mental models, and low observability.

To illustrate these factors, consider an example of a coordination surprise that was directly observed during the STS-76 space shuttle mission (Patterson, 1997). Prior to the shuttle Atlantis docking with the MIR space station, a NASA mission controller responsible for the mechanical systems on the shuttle (Mech) was surprised by a request by the Russian space agency to close the vent doors prior to docking. Evidence of the surprise includes prior statements made by the controller that he did not believe the action would be requested, a look of surprise when the request was made, and a delay in the timeline because implementing the action took longer than expected. In addition, the observed controller described the event to his replacement in the next shift by saying "In the unlikely event that we do it, I didn't want to be stumbling around…then all of a sudden we're doing this…"

The evolution of the mindsets of the American and Russian space agencies regarding whether or not to close the vent doors prior to docking are shown in Figure 1. Normally the vent doors are left open in space to allow oxygen to escape prior to entry. A hydraulic leak during ascent raised concerns that hydraulic fluid might contaminate the MIR space station. Analyses conducted by both space agencies showed that the amount of leaked hydraulic fluid was negligible so that there was no need to close the vent doors prior to docking. In addition, NASA planned to conduct a space walk during the mission, demonstrating that they were not concerned about the leaked hydraulic fluid contaminating the interior of the shuttle. During various interactions between the American and Russian space agencies, the two organizations presented variations on a stance toward the decision of whether or not to close the vent doors. One day before docking, the Russians announced that they were "90% go" on docking without closing the vent doors. The observed controllers assumed that this was a final decision not to close the vent doors.

Sometime between the conference call and the docking, a representative of the American space agency had a private phone conversation with a representative from the Russian space agency where the decision not to close the vent doors prior to docking was reversed. This reversal in decision was not communicated to the personnel in mission control. This coordination surprise illustrates the three factors that were identified as contributing to mismatch situations. 1) The representative was an agent that had high autonomy and authority in that he could negotiate the closing of the vent doors without being directed to do so by the controllers who would have to carry out the action. 2) In fact, the controllers did not even realize that an American representative could have a private phone conversation where important plans were negotiated, so there was a gap in their mental models of how there might be reversals in decisions following a public statement about a stance toward a decision. 3) In addition, there was also missing feedback in that the representative did not inform the mission controllers of the reversal in the decision.

Time	Mindset of the NASA Mission Controllers		Mindset of the Russian Space Agency
0:05 - 2:45	We need to think about how the hydraulic leak will impact the mission.		Same.
2:45 - 8:00	We need to look into whether or not to close vent doors to reduce MIR contamination.		Same.
8:00 - 17:20	There is no need to close vent doors but we probably will do it to satisfy the Russians.		We should close the vent doors to protect the MIR station.
17:20 - 17:23	We are "100% go" on docking. (conference call)		We are "90% go" on docking. (conference call)
17:20-1:09:00	We will not close the vent doors because the Russians said that it is not necessary.		It is not necessary to close the vent doors.
		PRIVATE CALL	**reversal**
1:09:00 - 1:17:00	**We will not close the vent doors** because the Russians said that it is not necessary.	**mismatch**	**We should close the vent doors** to be safe as long as it is will not make things worse.
1:17:57	Controllers are **surprised** by the Russians' request to close the vent doors.		Russians are **surprised** that it takes the controllers so long to close the vent doors.

Figure 1. Example of a coordination surprise

3. THE ESCALATION PRINCIPLE

During all the field observations in all of the domains, when an unusual or unexpected event was detected, there was an immediate escalation of cognitive and coordinative activities. This pattern can give rise to several problems because of a fundamental relationship: the greater the trouble in the underlying system or the higher the tempo of operations, the greater the information processing activities required to cope with the trouble or pace of activities (Woods *et al.*, 1994). For example, demands for monitoring, attentional control, information, and communication among team members (including human-machine communication) all go up with the unusualness (situations at or beyond margins of normality or beyond textbook situations), tempo and criticality of situations. This means that the burden of interacting with other agents or an interface tends to be concentrated at the very times when the practitioner can least afford new tasks, new memory demands, or diversions of his or her attention away from the job at hand.

As an illustrative example, during ascent of the space shuttle mission STS-76, there was a hydraulic leak in Auxiliary Power Unit (APU) 3 which triggered escalations in cognitive and coordinative activities (Watts, Woods, and Patterson, 1996, Watts and Woods, 1997). One of the controllers responsible for monitoring the health and safety of the mechanical systems noticed an unexpected drop in hydraulic fluid. The team of controllers immediately calculated the leak rate in order to recommend an action to the astronauts. Because the leak was small enough not to require an immediate abort, the controllers then focused on how to best configure the APUs in order to obtain the most diagnostic information before shutting down the systems while also analyzing if any actions could be taken to protect other interrelated systems. In parallel with these activities, the controllers were constantly updating members of their immediate team, the flight director, and other support controllers who were calling in. In addition, they were giving instructions to be relayed to the astronauts through the CAPCOM controller. Even before the astronauts entered the orbit phase and shut down the APUs, additional support personnel were called in to help assess the impacts of the anomaly on mission plans.

4. TECHNOLOGICALLY-MEDIATED COMMON GROUND

Along with many others, we have observed that practitioners in operational settings rely on a shared understanding or ‘common ground’ when communicating, and that this common ground is mediated by artifacts and technology (Clark and Brennan, 1991; McDaniel, Olson, and Magee, 1996; Kraut, Miller, and Siegel, 1996; Roth, Bennett, and Woods, 1987). In distributed supervisory control, important elements of the common ground include shared understandings about a referent monitored process and the responses of distributed agents in relation to the monitored process (Malin *et al.*, 1991; Johannesen, Cook, and Woods, 1994). In our observations, we have noticed the widespread use of auditory mediating technologies that support a practitioner in maintaining peripheral awareness of ongoing activities of other practitioners in relation to a monitored process without disrupting their ongoing work or the communication process between the monitored parties (e.g., voice loops in space shuttle mission control, Watts *et al.*, 1996b; train controllers in the London Underground, Heath and Luff, 1992; voice loops in air carrier operations, Rochlin *et al.*, 1987). With these technological aids, practitioners are able to ‘listen in’ on the activities of other distributed agents. ‘Listening in’ facilitates coordination in response to events in a monitored process and primes practitioner’s expectations for when other practitioners might need to coordinate with them as well as letting them know what is happening with subsystems of the monitored process that might affect them (Woods, 1995; Watts *et al.*, 1996b).

The common ground is considered to be vitally important in effective coordination during responses to anomalies. Practitioners in several domains are observed to invest in creating a common ground before there is an obvious reason to do so. Space shuttle mission controllers monitor voice loops that are not directly relevant to them in case a problem arises where they would then be required to coordinate with those controllers (Watts *et al.*, 1996b). Medical practitioners proactively update the other medical practitioners involved in an operation before there is any evidence that there might be a problem (Johannesen, Cook, and Woods, 1994). Space shuttle mission controllers, who are not directly scheduled for a mission but who could potentially be called in if a problem arises, are expected to obtain daily updates from staffed controllers (Patterson and Woods, 1997).

5. DISTANT SUPERVISORS AND LOCAL ACTORS

Most of the domains that are studied in the Cognitive Systems Engineering Lab (CSEL) are distributed supervisory control systems. Distributed supervisory control systems are hierarchical and cooperative architectures where remote supervisors work through intelligent local actors to control a process. With this framework, human supervisors, designers, and procedure writers could all be viewed as remote supervisors who implicitly or directly influence local actors. The distant supervisors have a broader scope and a better understanding of the overarching goals for the distributed system. The local actors have privileged access to the monitored process and what is actually happening “on the ground.”

Shattuck and Woods (1997) investigated how local actors adapted when surprises occurred in simulated command and control scenarios and how they used their commander’s statement of intent behind the plan in adapting to unexpected events. At one extreme, practitioners would rotely follow the original plans as described by the supervisor with no regard for the local complicating factors. At the other extreme, practitioners would act completely autonomously, leaving their supervisors ‘out of the loop’ and failing to coordinate with other local actors toward an organizational target. The results demonstrate the need to strike a cooperative balance between remote supervisors and local actors, where local actors have the knowledge and authority that they need to respond to unanticipated local situations in ways that support achieving higher level goals.

Dekker and Woods (1997) observed another aspect of coordination between local actors and distant supervisors. In an envisioned new form of air traffic management, authority to control flight paths will be distributed mostly between pilots and company dispatchers. Air traffic controllers will monitor the aircraft and intervene only to preserve safety -- a management by exception cooperative architecture (Billings, 1996). Dekker and Woods investigated how this architecture creates a dilemma for the distant supervisor about whether and when to take back partial or complete control over what flight path some subset of aircraft will fly. If the supervisor intervenes early (when it is easy to understand the situation and act constructively), it will tend to be seen as over-intervention. If the supervisor intervenes late, it will be very difficult to act constructively given the tempo of the situation and the workload involved in assessing and directing multiple aircraft. This particular case raises a variety of questions about what shared models, shared information, and common or overlapping fields of view are needed to support dynamically adjusting the distribution of authority to match changing constraints.

6. UPDATING AND ON-CALL ARCHITECTURES

Under pressure to use expertise more efficiently, many of the observed domains are fundamentally changing their supervisory control architecture. When situations are routine, fewer staff and less experienced staff monitor and adjust the underlying process. Only when anomalies occur and situations depart from routine are more personnel, more expert personnel, and more specialized personnel brought in to cope with the anomalous situation. Personnel at work during nominal operations need to recognize off-nominal situations, to call on appropriate expertise, and to call in additional resources. This ‘on-call’ architecture for supervisory control has the potential to effectively utilize expert practitioners by using them only when expert knowledge and more intense analyses and planning are required in order to respond to a problem. This means the distributed supervisory control system needs to be able to bring increasing expertise to bear smoothly and coordinate multiple cognitive activities as situations escalate in difficulty and tempo.

When practitioners are called in, they must be updated so that they know and function as if they had been present during the previous activities (Johannesen, Cook, and Woods, 1994; Patterson and Woods, 1997). They must learn what events have taken place and how plans have been revised as a result of these events. They also need to know what analyses have been performed. In many domains, commitments to irreversible decisions are delayed, so it is also important to know the team's stance toward critical decisions in order to influence the choice if the opportunity arises. Practitioners' expectations for monitoring system parameter values depend on past events and current system configurations (e.g., a system is leaking so fluid levels are expected to be lower). Knowing this information helps them to anticipate what to do in the case of contingencies (e.g., open the relief valve if pressure begins to rise in the suspect system). They must also learn the status of communications with other agents in the system, such as who has been updated, what written documents need to be distributed, who has requested permission for changes to plans, and who is involved in finalizing the commitment to a course of action. Observations in space shuttle mission control during shift changes revealed that practitioners employ strategies that rely heavily on prior knowledge and shared understandings to have quick and effective updates. For this reason, the organization requires controllers who are assigned to provide support in the event of an anomaly to receive periodic updates before any problems arise.

7. COORDINATION ACROSS FUNCTIONALLY DISTRIBUTED TEAMS

In many domains, the interdependent cognitive processes of anomaly response are distributed across a set of functionally distinct but cooperative teams who possess distinct but overlapping expertise and perspectives. For example, space shuttle ground support consists of many sub-communities such as operations and engineering with common overall goals. However, each sub-team has distinct responsibilities, resources, and authority, which lead them to approach these overall goals from different perspectives.

Watts *et al.* (1996a, 1997) studied the coordination across these functionally distinct teams during actual anomalies in shuttle operations. They described a pattern of cooperative advocacy that seemed to provide a robust mechanism to cope with the demands and potential pitfalls of anomaly response. Each sub-team developed their own assessment and response strategy of the anomaly and its consequences for the mission -- a within-team perspective -- and then shared their perspectives in a series of coordinative meetings. Preparing for a possible critique and actually confronting another group's perspective on the situation revealed inaccuracies, gaps, uncertainties, and conflicts. The process of sharing each sub-team's assessment stimulated individuals to call to mind other possibilities, constraints, and side effects.

For example, during the observed anomalous space shuttle mission STS-76, there was a hydraulic leak in an Auxiliary Power Unit (Figure 2). To avoid the somewhat unlikely but high-risk scenario of losing capability in another of the three APUs, the operational community (MOD) wanted to avoid using an APU to test the flight control systems, as would normally be done a day before entry. In seemingly direct opposition, the engineering community (MER) wanted to use an APU for the flight control system test in the same way as would be done on a nominal mission in order to gain more information about the affected systems.

This opposition in stance resulted from the two community's differing scopes of responsibility. The operational community is directly responsible for safety during a particular mission. They believed that any risk associated with gathering additional information should be justified in terms of how the information might affect operational decisions on that flight. Since it was unlikely that additional information would affect the configuration of the systems for entry and there was some risk associated with using an APU, they were opposed to the test. The engineering community, on the other hand, is responsible for the safety of the shuttle system design across all the missions. They suggested that since the risk of using the APU was low, it was a reasonable decision to test the flight control systems. In this way, they argued that they could ensure that there were no hidden problems with the flight control system as well as gain information that might be valuable in redesigning the APU.

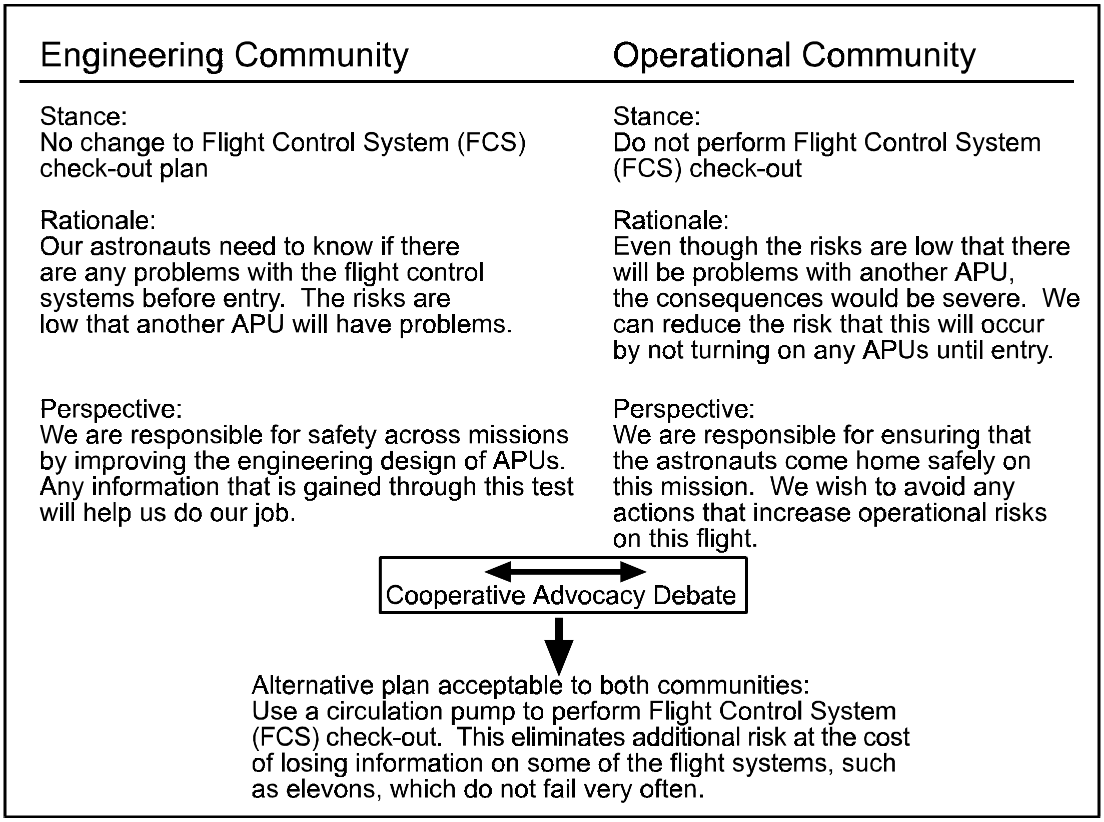

Engineering Community

Stance:
No change to Flight Control System (FCS) check-out plan

Rationale:
Our astronauts need to know if there are any problems with the flight control systems before entry. The risks are low that another APU will have problems.

Perspective:
We are responsible for safety across missions by improving the engineering design of APUs. Any information that is gained through this test will help us do our job.

Operational Community

Stance:
Do not perform Flight Control System (FCS) check-out

Rationale:
Even though the risks are low that there will be problems with another APU, the consequences would be severe. We can reduce the risk that this will occur by not turning on any APUs until entry.

Perspective:
We are responsible for ensuring that the astronauts come home safely on this mission. We wish to avoid any actions that increase operational risks on this flight.

Alternative plan acceptable to both communities:
Use a circulation pump to perform Flight Control System (FCS) check-out. This eliminates additional risk at the cost of losing information on some of the flight systems, such as elevons, which do not fail very often.

Figure 2. Example of functionally distinct teams advocating a stance

While debating these stances, it became clear that an alternative plan was acceptable to both communities. Instead of using an APU to check out the flight control systems, a circulation pump could be used. This would eliminate the risk of losing capability in either of the remaining two APUs but also meet the need to test the flight control systems. There was a loss of some information on the flight control systems that could only be gained by running an APU, but this plan met the most critical needs of the goals advocated by both communities. It seems unlikely that this plan would have been formed without having functionally distinct teams openly critiquing alternative options.

Cooperative advocacy as a strategy for coordinative anomaly response assumes that there are multiple sub-teams or people, each diverse in perspective, background, methods, or goals (Watts and Woods, 1997). There is a movement as each makes an investment to develop their own assessment or stance and then shares that result with the other sub-teams or people in a shared event or environment. The mixing of the separate analyses spawns revision, cross checks for detecting and recovering from erroneous assessments, and cues that call to mind new possibilities. Cooperative advocacy seems like an interesting possible method for structuring coordinative activity when there is a wide scope of factors that could be relevant, when all things cannot be thought through in advance, and when there are uncertainties and potentially serious consequences as a result of misassessments or mis-plans.

8. CONCLUSION

This paper describes patterns in cooperative cognition that were observed in several different operational settings and investigated through different field research methodologies. These patterns further our understanding of how cognition is fundamentally cooperative. This understanding could provide insight into how the design of machine artifacts could aid practitioners in various ways. Systems could be designed to minimize coordination breakdowns, such as by avoiding coordination surprises or supporting bringing more knowledge to bear as situations escalate. Alternatively, patterns of existing cooperative strategies for coping with domain complexities could be augmented and extended. For example, technologies could be provided to better mediate building and maintaining a common ground, better balance the goals of distant supervisors and with the need for local actors to adapt to unique circumstances, augment updates given to personnel who are called in to respond to an unexpected situation, and support interactions between functionally distinct teams during anomaly response.

9. ACKNOWLEDGMENTS

This paper represents contributions from many current and prior members of the Cognitive Systems Engineering Laboratory, including Lawrence Shattuck, Leila Johannesen, Sidney Dekker, James Corban, Scott Potter, Matthew Holloway and Marie Walters, as well as many collaborators, including Richard Cook, Jane Malin, Charles Billings, Emilie Roth and Philip Smith. Two of the authors were supported under National Science Foundation Graduate Research Fellowships. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

10. REFERENCES

Billings, C. E. (1996). *Aviation Automation: The search for a human-centered approach.* Hillsdale, NJ: Erlbaum.

Clark, H.H. & Brennan, S.E. (1991). Grounding in communication. In L. Resnick, J.M. Levine, and S.D. Teasley (Eds.) *Perspectives on Socially Shared Cognition.* Washington, D.C.: American Psychological Association, 127-149.

Dekker S. & Woods D. D. (1997). Management by exception in future air traffic control: An empirical study of coordination in an envisioned distributed environment. Manuscript submitted for publication.

Heath, C., & Luff, P. (1992). Collaboration and control: Crisis management and multimedia technology in London underground line control rooms, *Computer-Supported Cooperative Work* (1), 69-94.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Johannesen, L. J., Cook, R. I., & Woods, D. D. (1994). *Grounding explanations in evolving diagnostic situations* (CSEL Report 1994-TR-03). The Ohio State University, Cognitive Systems Engineering Laboratory.

Kraut, R. E., Miller, M. D., & Siegel, J. (1996). Collaboration in performance of physical tasks: effects on outcomes and communication. *CSCW '96 Proceedings*, Boston, MA., 57-66.

Malin, J., Schreckenghost, D., Woods, D. D., Potter, S., Johannesen, L. J., Holloway, M., & Forbus, K. (1991). *Making Intelligent Systems Team Players*. NASA Technical Report 104738, Johnson Space Center, Houston TX.

McDaniel, S.E., Olson, G.M., & Magee, J.C. (1996). Identifying and analyzing multiple threads in computer-mediated and face-to-face conversations. *CSCW '96 Proceedings*, Boston, MA., 39-47.

Patterson, E. S. (1997). *Coordination across shift boundaries in space shuttle mission control.* (CSEL Report 1997-TR-01). The Ohio State University, Cognitive Systems Engineering Laboratory.

Patterson, E.S., & Woods, D.D. (1997). Shift changes, updates, and the on-call model in space shuttle mission control. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*. Albuquerque, NM: 243-247.

Rochlin, G. I., La Porte, T. R. & Roberts, K. H. (1987). The self-designing high-reliability organization, Aircraft carrier flight operations at sea. *Naval War College Review,* 76-90.

Roth, E. M., Bennett, K., & Woods, D. D. (1987). Human interaction with an 'intelligent' machine. *International Journal of Man-Machine Studies,* 27:479-525.

Sarter, N. B., & Woods, D. D. (1997a). "Teamplay with a Powerful and Independent Agent": A Corpus of Operational Experiences and Automation Surprises on the Airbus A-320. *Human Factors*, 39(4), 553-569.

Sarter, N. B., & Woods, D. D. (1997b). Mode Errors of Omission and Commission: Observed Breakdowns in Pilot-Automation Coordination in a Full Mission Simulation Study. Manuscript submitted for publication.

Sarter, N. B., & Woods, D. D. (1994). Pilot Interaction with Cockpit Automation II: An Experimental Study of Pilot's Model and Awareness of the Flight Management System. *International Journal of Aviation Psychology*, 4:1-28.

Sarter, N. B., & Woods, D. D. (1992). Pilot Interaction with Cockpit Automation I: Operational Experiences with the Flight Management System. *International Journal of Aviation Psychology,* 2:303--321.

Sarter, N. B., Woods, D. D., & Billings, C. (1997). Automation Surprises. In G. Salvendy, editor, *Handbook of Human Factors/Ergonomics*, second edition, Wiley, New York.

Shattuck, L. G., & Woods, D. D. (1997). Communication of intent in distributed supervisory control systems. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*. Albuquerque, NM, 259-268.

Watts-Perotti, J, & Woods, D. D. (1997). *A cognitive analysis of functionally distributed anomaly response in space shuttle mission control*. (CSEL Report 1997-TR-02). The Ohio State University, Cognitive Systems Engineering Laboratory.

Watts, J., Woods, D. D., & Patterson, E. S. (1996a). Functionally distributed coordination during anomaly response in space shuttle mission control. *Human Interaction with Complex Systems '96,* Dayton, OH.

Watts, J., Woods, D. D., Corban, J., Patterson, E. S., Kerr, R., & Hicks, L. (1996b). Voice loops as cooperative aids in space shuttle mission control. *CSCW '96 Proceedings,* Boston, MA., 48-56.

Wiener, E.L., & Curry, R.E. (1980). Flight-deck automation: promises and problems. *Ergonomics,* 23(10), 995-1011.

Woods, D. D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics,* 38(11), 2371-2393.

Woods, D. D., Johannesen, L. J., Cook, R. I., & Sarter, N. B. (1994). *Behind Human Error: Cognitive Systems, Computers, and Hindsight*. Dayton, OH: CSERIAC.