Voice Loops as Coordination Aids in Space Shuttle Mission Control

Emily S. Patterson, Jennifer Watts-Perotti[1], David D. Woods
Cognitive Systems Engineering Laboratory
Institute for Ergonomics
The Ohio State University

Computer Supported Cooperative Work, 8, 353—371, 1999.

Abstract

Voice loops, an auditory groupware technology, are essential coordination support tools for experienced practitioners in domains such as air traffic management, aircraft carrier operations and space shuttle mission control. They support synchronous communication on multiple channels among groups of people who are spatially distributed. In this paper, we suggest reasons for why the voice loop system is a successful medium for supporting coordination in space shuttle mission control based on over 130 hours of direct observation. Voice loops allow practitioners to listen in on relevant communications without disrupting their own activities or the activities of others. In addition, the voice loop system is structured around the mission control organization, and therefore directly supports the demands of the domain. By understanding how voice loops meet the particular demands of the mission control environment, insight can be gained for the design of groupware tools to support cooperative activity in other event-driven domains.

Keywords: attention, broadcasting, common ground, coordination, ethnographic study, mission control, mutual awareness, overhearing, voice loops

1. INTRODUCTION

In supervisory control, cognitive activities such as monitoring and anomaly response are often distributed across interdependent sets of practitioners. As Hughes, Randall, and Shapiro (1992) and others have noted, practitioners in these domains must be able to coordinate their efforts on a "moment to moment basis, in response to constantly changing circumstances." Voice loops, a groupware technology which allows synchronous communication among groups of people who are spatially distributed, are used to aid coordination in domains such as air traffic management, aircraft carrier operations (Rochlin, LaPorte, and Roberts, 1987) and space shuttle mission control.

---

[1] Now at Eastman Kodak Company, Rochester, NY

Voice loops are essential coordination support tools for experienced practitioners in space shuttle mission control. Controllers use the voice loops to directly communicate with other personnel in mission control. More importantly, however, controllers use the voice loops to remain aware of the activities of other controllers and mission events in related shuttle subsystems. Controllers continuously monitor approximately four voice loops while directly communicating on a primary loop. By being aware of events when they occur, they can synchronize their activities with other controllers and with the actions of the astronauts. If something that is reported on the loops does not match their expectations, they can direct their attention to that thread of conversation and investigate what the deviations are and how their own activities might be impacted.

This paper analyzes how voice loops facilitate the coordination of physically distributed practitioners in space shuttle mission control based on direct observations of voice loops in use. We describe how mission control is structured and how the different types of voice loops reflect the organizational structure. Then we outline the coordination functions that voice loops support and illustrate these functions in an example. We conclude with a discussion of how this analysis provides insight into the important functions that should be considered in the development of systems intended to support cooperative work in other event-driven domains.

2. METHODS

Our findings are based on direct observations of flight controllers using voice loops to support their activities during space shuttle operations (see Figure 3 for an excerpt of observed voice loop communications). The observations were conducted at the Maintenance Mechanical Arm and Crew Systems (MMACS), Mechanical (Mech), Payloads Officer (Payloads) and Remote Manipulator System (RMS) flight control consoles. Over 130 hours of observation were conducted during portions of four actual missions and 27 flight control simulations. Flight control simulations include a full complement of astronauts and flight controllers supporting each flight control console. The high-fidelity simulations are used to train the controllers to respond to unexpected problems.

In addition to these observations, we interviewed controllers during low-tempo periods in the simulations and missions about how they use voice loops to support their activities. Controllers described formal and informal protocols that govern the usage of voice loops in mission control, which loops they monitor, why they monitor them, and how and when they are expected to speak on the loops. In addition, we interviewed the personnel who manage the voice loop system to learn how the specific voice loop assignments are made for each mission, how the loops are managed during flight operations, and which loops are permanently archived.

3. COOPERATIVE STRUCTURE OF SPACE SHUTTLE MISSION CONTROL

3.1 Hierarchical supervisory control structure

The voice loop system maps onto the supervisory control structure of mission control. NASA's Mission Control Center (MCC) at the Johnson Space Center in Houston, Texas is responsible for managing space shuttle missions from take-off to touchdown. During missions, teams of flight controllers monitor spacecraft systems and activities 24 hours a day, 7 days a week. The head flight controller is the flight director, referred to as "Flight." Flight is ultimately responsible for all decisions related to shuttle operations and so must make decisions that trade off mission goals and safety risks for the various subsystems of the shuttle. Directly supporting the flight director is a team of approximately sixteen flight controllers who are co-located in a single location called the "front room" (Figure 1). These flight controllers have the primary responsibility for monitoring the health and safety of shuttle functions and subsystems. For example, one controller is responsible for the electrical subsystems (EGIL) and another for the payload systems (Payloads). These controllers must have a deep knowledge of their own systems as well as know how their systems are interconnected to other subsystems (e.g., their heater is powered by a particular electrical bus) in order to recognize and respond to anomalies despite noisy data and needing to coordinate with other controllers.

Figure 1. The front room in the Mission Control Center

Each of the flight controllers located in the front room has a support staff that is located in "back rooms." The front room and back room controllers communicate with each other through the voice loop system by activating a voice loop channel through a touch screen and talking into a headset. The back room support staff are more specialized than the front room controllers on specific shuttle subsystems and monitor more detailed information sources. For example, the front room controller responsible for the Maintenance Mechanical Arm and Crew Systems (MMACS) has supporting staff members who specialize in: (1) the mechanical systems (Mech I and II), (2) the photo equipment used by the astronauts (Photo/TV), (3) the escape hatch that would be used in the event of an aborted take-off (Escape), and (4) in-flight maintenance for the tools used by the astronauts (IFM).

3.2 Example of coordination across flight controllers

To illustrate how controllers with different scopes of responsibilities coordinate their activities, consider the case of the unexpectedly high pressure in an Auxiliary Power Unit’s (APU’s) fuel pump during the STS-62 mission. Fuel pumps are used to keep mechanical systems warm in the freezing environment of space. During orbit, the front room mechanical systems controller (MMACS) and his support staff (Mech) noticed that pressure cycles in an Auxiliary Power Unit’s (APU’s) fuel pump inlet pressure were higher than expected. The MMACS controller updated the flight director about the abnormal readings and requested that the crew respond to the anomaly by opening the fuel isolation valves to equalize the pressure in the fuel lines and the fuel tank. The flight director approved the request and then described the plan to the front room communications controller (CAPCOM) who then relayed the requested action to the astronauts.

The results from this intervention as well as four subsequent interventions were unexpected and contradictory. Over several days, the MMACS team coordinated with other front room controllers and various engineers who designed relevant shuttle subsystems to generate nineteen competing hypotheses as possible explanations for the anomaly. The selected best explanation, that water in the fuel line had frozen and then thawed, ended up only minimally disrupting mission plans. Other competing explanations, such as a hydrazine fuel leak that would start a fire upon re-entry into the atmosphere, would have had many more implications for other subsystem controllers.

It is clear in this example that coordination was essential to diagnosing the fault and creating a response plan. The voice loops were important in supporting this coordination. The MMACS and Mech controllers were able to work through detailed diagnoses without moving to the same physical location. The sequence of communications, from the MMACS controller to the flight director to the CAPCOM controller to the astronauts, were less error-prone on the public voice loops system than if the communications had been private because any of the mission controllers could have intervened if there were miscommunications. Other subsystem controllers were able to anticipate questions that they might be asked by listening in on these communications. The MMACS controller did not have to explicitly update controllers with inter-connected subsystems because the other controllers were listening to their public updates to the flight director. In addition, the MMACS controllers were able to listen for unexpected changes in other controllers’ sub-systems that might provide new information for their ongoing diagnostic analyses.

4. THE VOICE LOOP SYSTEM

4.1 Voice loop interface and controls

As previously stated, a voice loop is a real-time auditory channel that connects physically distributed people. A controller who speaks on a loop broadcasts to all controllers who are listening in on that loop. A controller monitoring a loop hears any communication among other controllers connected to that loop.

Multiple voice loops can be monitored at the same time. The multiple loops require a mechanism for controllers to select and modify which loops they are monitoring and which they can speak on. The interface is a map of the available loops. Any controller can choose to monitor any of the available communication channels by directly manipulating this representation of the space of channels.

By formal communication protocols in mission control, flight controllers have privileges to speak on only a subset of the loops they can listen in on. In the voice loop control interface, each channel can be set either to monitor or talk modes. Only one channel at a time can be set to the talk mode, although many channels can be monitored at the same time. In order to talk on a loop set to the talk mode, a controller presses a button on a hand unit or holds down a foot pedal and talks into a headset.

Each controller customizes the set of loops they monitor by manipulating the visual representation of the loops at their console. The controllers can save a configuration of multiple voice loops on ‘pages’ under their identification code. The most commonly used loops are grouped together onto a primary page. The controllers then reorganize and prioritize the loops to fit the particular operational situation going on at that time by changing the configuration of loops that are being monitored and by adjusting the relative volume levels on each loop.

The voice loop interface is generally considered to be easy to use and an appropriate communication tool for a dynamic environment like space shuttle mission control. The fundamental display units are visual representations of each auditory loop, which captures the way controllers think about the system. In addition, if individual loops are analogous to windows in a visual interface, then the pages of sets of loops are analogous to the ‘room’ concept in window management (Henderson and Card, 1986). Controllers are able to customize the interface by putting their most commonly used loops together on a single ‘page.’ Active loops on these pages can be dynamically reconfigured in response to the constantly changing environment. Dynamic allocations of which loops to listen to are done by directly selecting loops to turn off and on. Controllers increase or decrease the salience of particular loops by using loop volume controls to adjust relative loudness.

4.2 Voice loop organization reflects mission control structure

The voice loop system design reflects the cooperative structure in mission control (Figure 2). A primary voice loop, the Flight Director loop, is dedicated to communications between the flight director and the primary controllers in the front room. All controllers continuously monitor the Flight Director loop, but only direct communications between the front room controllers and the flight director are allowed. Because of the importance of this loop, only issues of high significance are discussed on it, and communication is kept clear and concise.

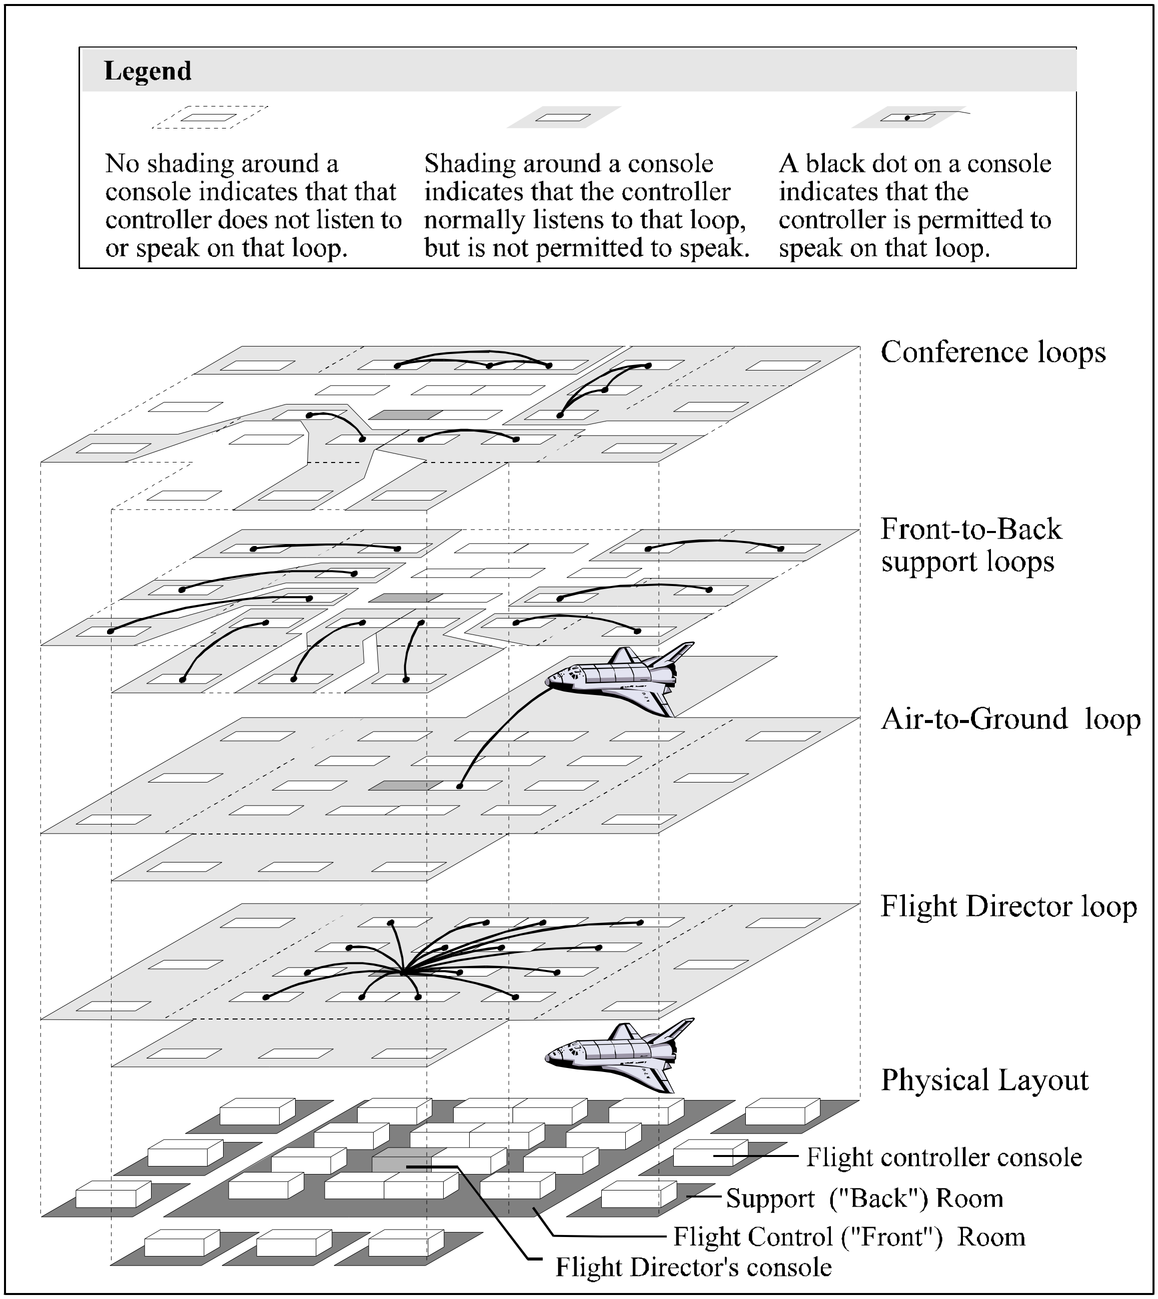

Figure 2. The Voice Loop Structure in Mission Control

Similarly, all controllers monitor the Air-to-Ground loop between mission control and the astronauts, but only one controller, CAPCOM, is authorized to communicate with the astronauts. CAPCOM is an astronaut and physically sits next to the flight director because the flight director makes the final decisions on what should be communicated on this loop. Despite their ability to interact face-to-face, most of the communications between Flight and CAPCOM are done on the flight director loop so that other controllers can listen in.

Interactions between front room controllers and their support staff are conducted on Front-to-Back support loops. These are the loops that are normally set to the ‘talk’ setting and are often not monitored by other subsystem controllers. Discussions on these loops are much more detailed and less formal than communications on higher-priority loops. For example, unexpected telemetry data values and factors that might account for them would be discussed on these loops.

Conference loops are continuously monitored but lie unused until a situation arises that requires coordination across subsystem controllers. By having dedicated conference loops, groups of subsystem controllers that would need to interact during predictable failure scenarios can be quickly formed without tying up communications on the other loops. When a meeting between controllers who do not have a dedicated conference loop needs to be formed, a front room controller announces on the Flight Director loop for specific controllers to meet on an ad hoc conference loop.

4.3 Monitoring multiple loops in parallel

Each controller typically monitors a minimum of four loops in parallel: the Flight Director loop, the Air-to-Ground loop, the Front-to-Back loop, and a conference loop. These loops are only a small subset of the potentially available loops. For example, 164 voice loops were used during the STS-76 space shuttle mission.

While it may seem difficult to monitor multiple loops in parallel, this ability is essential to controllers' activities and goals. For example, during an observed simulation, the front room mechanical systems controller noticed an abrupt change in the data on her telemetry screens. In order to determine the cause of this change, she monitored and interacted with four loops in parallel. By listening for deviations in standard communications on the Air-to-Ground loop, she could track whether the astronauts were experiencing any abnormal circumstances aboard the shuttle. Listening to the Flight Director loop kept her aware of whether or not other controllers were also seeing strange data patterns. She contacted a related controller on a conference loop to give him a "heads up" that her systems were functioning abnormally, and she discussed the details of the data with her back room staff to determine what might have caused the change in the data. Eventually, she heard the electrical controller inform the flight director on the Flight Director loop that an electrical bus had failed. This failure would account for the unexpected changes in her system data. She contacted the electrical controller on a conference loop to find out whether the bus could be fixed, and then discussed the impact of this failure with her support staff over the Front-to-Back loops.

5. FACILITATING COORDINATION

The purpose of this paper is to suggest reasons for why the voice loop system in space shuttle mission control is a successful coordination aid. The voice loop system is a powerful groupware tool because it allows practitioners to ‘listen in’ on others’ activities while also pursuing their own goals and activities. Perhaps more importantly, listening in on the voice loops does not interfere with ongoing activities of other personnel. In addition, the voice loop technology supports the coordination functions of synchronizing activities, gauging when to interrupt ongoing communications, transferring information about high-level events, and directing other controllers’ attention.

5.1 Listening in

One might think that voice loops are used primarily for direct communication, where one controller uses the voice loops to speak directly to another controller. Although this kind of communication is supported by voice loop technology, the most important function of the voice loops in mission control is that they afford the ability to listen in to events and activities that occur across mission processes. The ability to listen in on discussions on the loops allows controllers to pick up relevant events and activities without disrupting their ongoing work or the communication process between the monitored parties.

Similar kinds of coordination have been observed in the control rooms of other event-driven domains. Luff, Heath, and Greatbach (1992) noticed that underground controllers thought out loud about schedule changes that they made during crisis situations in Line Control Rooms in the London Underground. When the controllers expressed changes out loud, other controllers in the vicinity took note of changes that affected their own schedules without interrupting the busy controller during the crisis situation. Similarly, Rochlin, La Porte and Roberts (1987) noted for voice loop communications in aircraft carrier operations that: “everyone involved...is part of a constant loop of conversation and verification taking place over several different channels at once. At first little of this chatter seems coherent, let alone substantive, to the outside observer...one discovers that seasoned personnel do not ‘listen’ so much as monitor for deviations, reacting to almost anything that does not fit their expectations.”

It is informative to note that we observed that mission controllers who were not assigned to the STS-76 mission preferred to track what was happening during the ascent phase by listening to the voice loops in an empty room rather than watching from the observation deck. They explained that they learned more by listening to the voice loops than by physically observing the mission controllers in the front room.

From a researcher’s standpoint, analyzing what personnel are monitoring what loops and how those configurations change in response to unexpected events gives a great deal of insight into how mission controllers coordinate. For example, when the mechanical systems controller (MMACS) announced a hydraulic leak during the STS-76 mission on the Flight Director loop, many of the other controllers immediately began monitoring the MMACS Front-to-Back loop. By doing so, they could prepare to answer questions that they might be asked. For example, originally the MMACS controller recommended shortening the STS-76 mission from 7 days to 3 days in response to the hydraulic leak. Each controller was asked soon after this announcement to estimate how shortening the mission would impact his or her subsystem. By listening in on the MMACS Front-to-Back loop, the controllers had more time to formulate an estimate of the impact of shortening the mission timeline than if they had waited for the specific announcement to be broadcast on the Flight Director loop.

5.2 Group synchronization

Anticipation has been found to be important for effective team coordination in high-tempo, dynamic situations (Decortis, de Keyser, Cacciabue, and Volta, 1991). When flight controllers are aware of upcoming mission events and activities, they can anticipate problems in their systems and prepare for future actions that will be required. Anticipation is important because it allows controllers to synchronize their communications and actions over time. For example, if a failure occurs in a subsystem, the flight director will ask related subsystem controllers about the impacts of that failure on their systems. When controllers hear about the failure on the Flight Director's loop, they can anticipate related questions from the flight director and prepare to answer them without delay. Controllers can also anticipate actions that will be required of them. For example, an anomaly in one subsystem might require diagnostic tests in another system. When the controller hears about the anomaly on the voice loops, he can anticipate the requirement of these tests, and prepare to conduct them when they are requested.

One way that voice loops aid synchronization is by affording the ability to track the tempo of mission processes. Since shuttle systems are interconnected, a failure in one subsystem may cause a cascade of disturbances throughout related systems. This cascade of disturbances causes an escalation of cognitive activities as controllers respond to these disturbances (Woods and Patterson, in press). For example, if an event like a complex anomaly occurs in a shuttle subsystem, the event triggers diagnostic activity in all related subsystem teams. This activity generates more communication across teams over the voice loops. Therefore, it is possible for controllers to track the cascade of disturbances in shuttle systems by tracking the escalation of activities that occur in response to these disturbances. This general indication of activity <u>tempo</u> allows controllers to synchronize their processes and activities with rest of the flight control team.

5.3 Gauging interruptibility

In contrast to a system of direct communications where only invited parties are involved in a conversation, voice loops allow controllers to listen to communications without announcing their virtual presence. This ability allows controllers to better gauge the relevance of their communications in relation to what is happening on a loop before interrupting. Controllers are then better able to time their communications, either by speeding up or postponing communications in relation to spurts of activity or by waiting for a pause in the communications to interject.

During our observations, we noticed that controllers gauge the interruptibility of practitioners outside their immediate team before communicating with them. When a controller needed to communicate with another controller working on a different subsystem, the controller would first listen to that controller’s Front-to-Back support loop. By listening to that loop, he could estimate the controller’s current workload to judge how interruptible the controller would be in terms of the criticality of the issues that she is addressing. Using this strategy reduces the number of unnecessary interruptions and allows controllers to judge the priority of their item against the ongoing work context. This reduces the chances that a controller will be forced to direct her attention away from current tasks in order to receive information about a lower priority item.

When the contacting controller misjudges the interruptibility of the person whom they are trying to contact, the receiving controller has the option to postpone communications by replying with "Standby." The standard protocol for initiating communications on the loops is to name the person that you wish to speak to in order to get his or her attention and then identify yourself (e.g., "Flight, MMACS"). The person who is called should then respond with either "Standby" or "Go ahead." With this protocol, the controllers can flexibly negotiate when to start communications on the loops. Compare these interactions with those conducted on a single-channel communication system like a standard telephone, where requesting people to standby would then tie up the only communication channel for both parties. Similarly, in face-to-face communications, waiting for interactions when the person to be contacted is busy generally means that the person's resources are tied up while waiting for the interaction.
In extreme situations, controllers may interrupt regardless of what other communications are going on. In crisis situations, controllers will use the Flight Director loop to broadcast critical information after declaring "Break! Break!" The "Break! Break!" communication is an explicit protocol that is rarely used and would instantly gain the attention of mission controllers.

5.4 Integrating levels of information

A common problem in information design is the display of raw, instantaneous sensor values rather than integrated information about the monitored process. For example, the mission control display screens provide continuously updated telemetry data for system parameters like temperature and pressure. The controllers must integrate this information to determine the system's global status and behavior by comparing the displayed data with memorized nominal ranges for specific contexts. The voice loops, however, allow controllers to take advantage of the data integration performed by other controllers. Instead of passing raw data about related systems from one controller to another, voice loops allow controllers to pass integrated, event-level information between controllers monitoring interconnected systems.

The voice loop system does more than simply make raw data available to controllers who would not otherwise have access to the data. For example, consider a hypothetical groupware system where all mission controllers could look at any of the data screens for any other controller. This system would effectively allow other controllers to "listen in" at a much lower level of information abstraction without allowing them to take advantage of processing by the controllers with primary responsibility for those subsystems.

The ability to pass event-level information is essential to controllers' tasks and goals. As mentioned previously, anomalies occurring in one system often cause a cascade of disturbances through other systems. Therefore, in order for controllers to meet the goal of anticipating and preparing for potential problems in their subsystems, they must be aware of events that are

happening in related subsystems. The voice loops facilitate this awareness by passing event-level information between controllers of related subsystems.
Communications on different voice loops provide information at different levels of abstraction and aggregation. Controllers can shift among these levels by selectively attending to specific loops, while still peripherally monitoring the other loops. The Flight Director loop functions as an overview of events and activities. Communications on the Flight Director loop consist of a brief summary of events related to a subsystem, combined with a recommendation for action. This overview loop can be monitored in parallel with Front-to-Back loops, where events are discussed at greater levels of detail. For example, if there is a hydraulic leak, controllers must infer this event from unusually low quantities of hydraulic fluid and pressure. On the Front-to-Back loops, teams of controllers would discuss detailed information such as thermal factors that might affect the hydraulic fluid readings. However, instead of describing these detailed factors to the flight director on the Flight Director loop, the controller responsible for the hydraulic systems (MMACS) communicates diagnostic results or status (e.g., there is a hydraulic leak) and implications for modifying mission plans (e.g., shortening mission duration).

5.5 Integrating information about mission events and practitioner activities

Research has shown that practitioners in process control domains must not only keep track of the monitored process, but they must also track the activities of other agents who are affecting that monitored process (Johannesen, Cook, and Woods, 1994; Patterson and Woods, 1997). For example, if the electrical controller performs a test of an electrical bus, his actions will affect the subsystems that receive power from that bus. Therefore, controllers of subsystems that depend on electrical power must be aware of the electrical controller's assessments and plans. They must use this information to alter their expectations for nominal telemetry values and possible contingency plans.

Voice loops provide an elegant view into the processes and activities of other practitioners. By monitoring discussions on the loops, controllers can pick up communications about activities performed by other controllers that will eventually impact their subsystems. The ability to track activities over the voice loops allows controllers to pick up information about the status of processes and activities without interrupting the controllers involved. If an anomaly in the mechanical systems occurs, the flight director can monitor the Front-to-Back loops of the mechanical console to track the status of the diagnosis or response process. This way, he can track the mechanical controllers' reasoning processes and progress without interrupting their activities by asking for an update.

5.6 Directing attention

In event-driven supervisory control domains, it is critical for practitioners to dynamically shift their focus of attention to re-orient to newly relevant communications, events, or activities. The voice loops are a powerful tool for redirecting controllers’ attention because it allows them to remain peripherally aware of others’ activities while still focusing on their current goals and activities. Background communications on the loops are occurring constantly. Sometimes these communications are noise, i.e., they are not relevant to a particular flight controller. But in other contexts, any of these communications could serve as a signal that they should interrupt ongoing lines of thought and re-orient their attention.

A basic challenge for any cognitive agent at work is where to focus attention next in a changing world. Laboratory-based research in attention and perception has revealed that the object, event, goal, or line of thought that we focus on depends on the interaction of two sets of activity (Yantis, 1993; Folk, Remington, and Johnston, 1992; Ward, 1997). The first is the set of knowledge-directed, endogenous processes that depend on the observer's current knowledge, goals, and expectations about the task at hand. The second set of processes is stimulus-driven, exogenous processes, where attributes of the stimulus world elicit attentional capture of shifts of the observer's focus. These two sets of processes combine in a perceptual cycle (Neisser, 1976), where unique events in the environment shift the focus of attention, call to mind knowledge, and trigger new lines of thought. The activated knowledge, expectations, or goals in turn guide further exploration and action.

Broadcast messages are an effective means of redirecting the attention of multiple practitioners to a particular item. When a broadcast is made on the Flight Director loop, relevant controllers focus their attention on the message while other controllers do not respond to the broadcast. Broadcasts are more robust than direct communications because they eliminate the opportunity for a controller to unintentionally leave someone "out of the loop." On the other hand, broadcasts have an associated cost in that they increase the amount of auditory activity that must be attended to and therefore should be limited to important announcements.

The ability to remain aware of what is happening without utilizing limited focal attentional resources is generally accepted as an important function in Computer-Supported Cooperative Work (Dourish and Bellotti, 1992, Woods, 1995; Gaver, 1997). Different labels have been used to refer to this characteristic of representations and media for collaboration. For example, Woods (1995) used the term "preattentive reference." The label preattentive was chosen because most models of attentional control include some type of preattentive processing to support the general cognitive function or competence in natural perceptual fields that observers/listeners monitor for new events or activities peripherally without disrupting or diverting focal attention (e.g., LaBerge, 1995). It is preattentive <u>reference</u> to capture the fact that we are concerned with virtual environments where external representations and support tools mediate contact with the world. Preattentive reference then refers to how the design of external representations, support tools, and communication media affects practitioners' abilities to remain peripherally aware of other events and others' activities, either supporting this cognitive function or undermining it.

6. AN ILLUSTRATIVE EXAMPLE

We have annotated a transcript from the observed STS-76 mission (Figure 3) to illustrate the coordinative functions that the voice loops support. This two-minute excerpt details what the back room Mech controller, who supports the front room mechanical systems controller (MMACS), was listening to during a segment of the high-tempo ascent phase. During this time, only three of his voice loops are active, although he was monitoring several others that were not being talked on in that specific time interval. Communications on the same row occurred at the same time.

During this segment, the back room mechanical systems controller (Mech) notices an unexpected decrease in hydraulic fluid. The Mech controller speaks on the Front-to-Back loop in order to direct the attention of the front room controller who is his supervisor (MMACS) to this system by mentioning that the hydraulic quantity is decreasing. The MMACS and Mech controllers diagnose the drop in hydraulic quantity as a hydraulic leak because it is clearly the best explanation for the data. They independently estimate the leak rate and then compare their estimates on the Front-to-Back loop. They decide that the leak rate is approximately between 0.3-0.6%/min. The front room mechanical systems controller (MMACS) then updates the flight director on the Flight Director loop about the hydraulic leak. This update also informs all of the mission control personnel as they are monitoring this loop. Note that the update on the Flight Director loop is a higher-level description than on the Front-to-Back loop. In addition, the description of the event is accompanied by the implied statement that there are no immediate changes to mission plans because the leak rate is less than the commonly known decision cut-off of 1%/min. to abort the ascent. The flight director then anticipates that the MMACS controller will recommend that the astronauts close an isolation valve (TVC ISO) in an attempt to isolate the leak after the main engines have been cut off and they have attained a stable altitude (post-MECO).

In parallel with discussions about the hydraulic leak, the astronauts are performing nominal sequences of events that are being reported on the Air-to-Ground loop. The Flight Dynamics Officer (FDO), who is responsible for planning orbiter maneuvers, is calling out the commands which are then relayed to the astronauts through the dedicated communication controller (CAPCOM). An astronaut on the shuttle (Atlantis) then reads back the commands in order to ensure that they were heard correctly. These commands form temporal landmarks around which the ground controllers synchronize their activities. For example, the Main Engine Cut-Off (MECO) is a critical landmark for the MMACS controller. Any action in response to the hydraulic leak, such as closing the isolation valve, needs to be taken before the main engines are cut off.

Figure 3. Voice Loop Excerpt From the STS-76 Mission

Before and after the discussion between MMACS and Flight about the hydraulic leak, the controller responsible for the booster rocket engines (Booster) is updating the flight director on the status of his subsystem. This update is particularly important because the previous mission had been aborted due to problems with the booster engine sensors. Many of the other controllers are carefully monitoring the status of the booster engines because they are anticipating another sensor failure. When Booster reports that the performance was nominal, the other controllers can stop focusing their attention on updates from Booster and can divert their attention away from how to implement contingency plans in the event of booster engine problems. Note that Booster's update is judged to be a lower priority than the discussions about the hydraulic leak. Rather than interrupting their ongoing communications, Booster waits for those discussions to pause before updating Flight.

7. DISCUSSION

In summary, the voice loops are a successful medium for supporting cooperative activity. They support critical coordination functions for practitioners in event-driven, supervisory control domains. Support for these functions can become criteria for Computer-Supported Cooperative Work tools in event-driven, supervisory control domains in general. Voice loops allow practitioners to synchronize their activities, gauge when to interrupt ongoing communications, transfer information about high-level events, and direct their attention to newly relevant activities or data. These are critical tasks for distributed supervisory controllers in dynamic environments that can be supported and augmented with groupware technology.

The communication medium created by voice loops is distinct from other media. It allows for communications among spatially distributed people, unlike face-to-face interactions. It is an auditory medium, so does not use the overloaded visual channel such as video-based or electronic mail systems would. The ability to listen to multiple loops in parallel extends telephone-based interactions and other single channel media, such as radio-based systems. The ability to dynamically adjust volume levels on the different loops enables differentiation of the loop communications and directing attention to particular threads of activity. It supports real-time interactions, unlike electronic mail systems, although it also allows interactions to be archived and replayed on request, unlike face-to-face interactions.

Overall, we attribute the wide acceptance of the voice loop technology in mission control mainly to two important factors. The first is that the controllers are able to pick up relevant information without disrupting either their own activities or the activities of others. The voice loops provide a medium for communication that allows controllers to remain peripherally aware of communications between physically distributed but functionally interdependent practitioners without requiring focal attention. In addition, controllers can listen in on communications without disrupting or even alerting the participants in the communications. In this way, the burden of interaction rests with the people who benefit from the information, which has previously been found to be important in the acceptance of groupware technology (Markus and Connolly, 1990).

The second factor is that the voice loop system is designed around the mission control organizational structure. Members within a team have dedicated Front-to-Back loops for detailed discussions about specific subsystems. Controllers responsible for subsystems that are inter-connected have dedicated conference loops to coordinate their efforts and allow for easy communication. The Flight Director loop allows highly observable interactions between subsystem controllers and the flight director. These interactions have the dual functions of providing input to a central decision maker as well as broadcasting information to all of mission control. The Air-to-Ground loop serves the function of efficient and effective communication with the astronauts who ultimately are the ones who implement the recommended actions. Ad hoc groups can form on loops in response to unexpected anomalies and the loops allow flexible reprioritization by adjusting relative loudness and the set of loops that are monitored.

Compare this context-sensitive design structure with a single global loop design where practitioners would take turns speaking (e.g., Thunderwire; Hindus, Ackerman, Mainwaring, and Starr, 1996). With a single-channel design, communications would take much more time because practitioners would have to wait their turn. There would be no way to dynamically select what to listen to, communications could not be conducted in parallel, and there would essentially be only one level of discussion that would have to be heard by everyone on the loop.

Another contrasting design concept would be to allow controllers to create any loops that they might want at any point in time. This extreme flexibility would create unnecessary burdens for the practitioners. It would force the controllers to figure out for themselves all the people that they might want to talk to and negotiate who should be on each of the loops. The need for loops that are used infrequently, such as conference loops, might not be recognized until a problem occurs. It would then be too difficult during the high-tempo response period to create the loops. Instead, communications that would be appropriate for a conference loop would probably be conducted on other loops, such as the Flight Director or Front-to-Back loops, disrupting those communications and tying up important communication pathways. In addition, loops would be created in ways that would be idiosyncratic to the particular teams rather than standardized. Without standardization, controllers would have to memorize the setup of specific loops in order to know who listens to them. In other words, the voice loop structure in mission control seems to provide a balanced level of flexibility that allows flight controllers to adapt to circumstances without creating new workload demands to reconfigure their tool set at the very point where the tempo and criticality of operations are increasing (Woods, 1993).

In this paper we have elucidated why voice loops are viewed as successful coordinative aids in space shuttle mission control. By understanding how voice loops meet the particular demands of the mission control environment, insight can be gained for the design of groupware tools to support cooperative activity in other event-driven domains.

8. ACKNOWLEDGMENTS

Research support was provided by NASA Johnson Space Center under Grant NAG9-390, Dr. Jane Malin technical monitor. We are particularly grateful to the flight controllers who shared their expertise with us and allowed us to observe their operations during training simulations and actual missions. Additional support was provided by two National Science Foundation (NSF) Graduate Fellowships. Any opinions, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

9. REFERENCES

Bentley, R., Hughes, J. A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D., and Sommerville, I. (1992). Ethnographically-informed systems design for air traffic control. In ACM (Ed.), *Proceedings of the CSCW '92 Computer Supported Cooperative Work*. Toronto, Canada.

Clark, H., and Brennan, S. (1991). Grounding in Communication. In L. Resnick, J. Levine, and S. Teasley (Eds.), *Socially Shared Cognition.* Washington, DC: American Psychological Association.

Decortis, F., de Keyser, V., Cacciabue, P.C., & Volta, G. (1991). The temporal dimension of man-machine interaction. In George R. S. Weir and James L. Alty (Eds.) *Human-Computer Interaction and Complex Systems*. San Diego, CA: Academic Press Inc.

Dourish, P., and Bellotti, V. (1992). Awareness and coordination in shared workspaces. In ACM (Ed.), *Proceedings of the CSCW '92 Computer supported Cooperative Work.* Toronto, Canada.

Folk, C. L., Remington, R. W. and Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology, Human Perception and Performance*, 18, 1030-1044.

Gaver, W.W. (1997). Auditory Interfaces. In M. Helander, T.K. Landauer, and P. Prabhu (Eds.) *Handbook of Human-Computer Interaction*, Second Edition, 1003-1041.

Henderson, D. and Card, S. (1986). Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics, 5(3)*, 211-241.

Hindus, D., Ackerman, M.S., Mainwaring, S., and Starr, B. (1996). Thunderwire: a field study of an audio-only media space. In ACM (Ed.), *Proceedings of CSCW '96 Computer-Supported Cooperative Work*. Boston, MA.

Hughes, J. A., Randall, D. and Shapiro, D. (1992). Faltering from ethnography to design. In ACM (Ed.), *Proceedings of CSCW'92 Computer-Supported Cooperative Work.* Toronto, Canada.

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

LaBerge, D. (1995). *Attentional Processing: The Brain's Art of Mindfulness*. Cambridge, MA: Harvard University Press.

Luff, P., Heath, C., and Greatbatch, D. (1992). Tasks-in-interaction: Paper and screen based documentation in collaborative activity. In ACM (Ed.), *Proceedings of CSCW'92 Computer-Supported Cooperative Work.* Toronto, Canada.

Markus, M.L., and Connolly, T.(1990). Why CSCW Applications Fail: Problems in the Adoption of Interdependent Work Tools. In ACM (Ed.), *Proceedings of CSCW '90 Computer-Supported Cooperative Work*, 371-380.

Murray, C. and Cox, C. B. 1989, *Apollo, The Race to the Moon*. New York: Simon & Schuster.

Neisser, U. (1976). *Cognition and Reality*. San Francisco: W. H. Freeman and Company.

Patterson, E.S., and 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. and Roberts, K. H. (1987). The self-designing high-reliability organization, Aircraft carrier flight operations at sea. *Naval War College Review*, Autumn, 76-90.

Ward, L. (1997). Involuntary Listening Aids Hearing. *Psychological Science*, 8(2), 112-118.

Woods, D.D. (1993). Price of flexibility in intelligent interfaces. *Knowledge-based systems, 6(4),* 189-196.

Woods, D. D. (1994). Cognitive Demands and Activities in Dynamic Fault Management, Abduction and Disturbance Management, in N. Stanton (ed.)., *Human Factors of Alarm Design.* New York: Taylor & Francis.

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

Woods, D.D., & Patterson, E.S. (in press). How Unexpected Events Produce An Escalation Of Cognitive And Coordinative Demands. P.A. Hancock and P.A. Desmond (Eds.). *Stress Workload and Fatigue*. Hillsdale NJ: Lawrence Erlbaum.

Yantis, S. (1993). Stimulus-drive attention capture. *Current Directions in Psychological Science*, 2(5), 156-161.