Chapter 9

Patterns in Multi-Threaded Work

The previous chapter considered anomaly response as one of the major forms of work in JCSs (another is modifying plans in progress). The patterns in anomaly response reveal an intermingled set of dynamic processes (Figure 11). As a general form of work, the cognitive task synthesis of anomaly response points out more general classes of demand factors that influence JCSs at work. When studying a JCS, investigators can bring these patterns to bear to facilitate the converging studies in functional synthesis. These general demand factors include tempo, escalation, coupling, premature narrowing, reframing, dilemmas, and over-simplifications.

MANAGING MULTIPLE THREADS IN TIME

Mindset is about attention and its control as evolving situations require shifts in where attention is focused in order to manage work over time, and as there are multiple signals and tasks competing for attention. Sometimes new signals are distractions from the current focus, but other times these signals are critical cues for shifts in focus. This case from Cook et al. (1992) illustrates the complexities of how mindset shifts (see also Woods et al., 1994; Dekker, 2002).

> **Unexpected hypotension in anesthetic management during surgery**
> During a coronary artery bypass graft procedure, an infusion controller device used to control the flow of a sodium nitroprusside (SNP) to the patient delivered a large volume of drug at a time when no drug should have been flowing. Five of these microprocessor-based devices, each controlling the flow of a different drug, were set up in the usual fashion at the beginning of the day, prior to the beginning of the case. The initial part of the case was unremarkable. Elevated systolic blood pressure (>160 torr) at the time of sternotomy prompted the practitioner to begin an infusion of SNP. After starting the infusion at 10 drops per minute, the device began to sound an alarm. The tubing connecting the device to the patient was checked and a stopcock (valve) was found closed. The operator opened the stopcock and restarted the device. Shortly after restart, the device alarmed again. The blood pressure was falling by this time, and the operator turned the device off. Over a short period, hypertension gave way to hypotension (systolic pressure <60 torr). The hypotension was unresponsive to fluid challenge but did respond to

repeated injections of neosynephrine and epinephrine. The patient was placed on bypass rapidly. Later, the container of nitroprusside was found to be empty; a full bag of 50 mg in 250 ml was set up before the case.

The physicians involved in the incident were comparatively experienced device users. Reconstructing the events after the incident led to the conclusion that the device was assembled in a way that would allow free flow of drug. Initially, however, the stopcock blocked drug delivery. The device was started, but the machine did not detect any flow of drug (because the stopcock was closed) and this triggered visual and auditory alarms (low or no flow). When the stopcock was opened, free flow of fluid containing drug began. The controller was restarted, but the machine again detected no drip rate, this time because the flow was a continuous stream and no individual drops were being formed. The controller alarmed again with the same message that had appeared to indicate the earlier no flow condition. Between opening the stopcock and the generation of the error message, sufficient drug was delivered to substantially reduce the blood pressure. The operator saw the reduced blood pressure, concluded that the SNP drip was not required, and pushed the control button marked "off." This powered down the device, but the flow of drug continued. The blood pressure fell even further, prompting a diagnostic search for sources of low blood pressure. The SNP controller was seen to be off. Treatment of the low blood pressure itself commenced and was successful.

The case is interesting because performance was both successful—managing the disturbance to protect the patient's physiology—and unsuccessful—unable to determine that the infusion device was the source of the hypotension (see also Cook et al., 1998).

As many interwoven events were happening, practitioner mindset shifted as attention flowed to some parts of the situation but not to others. Even though the practitioners did not understand the source of the anomaly, they acted quickly to correct the physiologic, systemic threat (safing).

Why didn't the infusion device receive attention to determine if it was the source of the unintended flow of drug? Severe limits on the observability of the device's activity were critical (Moll van Charante et al., 1993): the device display shows only demanded rate, not actual; misleading alarm messages; visual inspection of the device state was blocked (an aluminum shield surrounded the fluid bag, hiding its decreasing volume; the drip chamber was obscured by the machine's sensor; the clamping mechanism was hidden inside its assembly; the complexity of tubing pathways with multiple infusion devices). The belief that the device had been designed to handle all contingencies and failure modes was illusory since the design did not effectively address the potential for misassembly; the device had the potential to fail active as in this case; the sensing mechanism could not detect free flow. This overconfidence, combined with very low observability of the interface, meant that there was little in the design to support resilience of the joint system.

The practitioners reported that they turned the device off as soon as the blood pressure fell, at about the same moment that the device alarmed a second time. In their mindset, the device was off, unimportant and outside their focus on the anomaly. This does not mean that the device was examined and found to be uninvolved in the evolving situation. If this had been the case, attention would have

flowed to the device for a time to investigate what it was doing, taking attention away from the other aspects of the situation.

Once the device was "off," it was dismissed from their mindset as attention flowed to bringing blood pressure under control and considering what physiological processes could be leading to the hypotension. The practitioners did not make inferences about how the device was working or not working or how it might be playing a role: they did not attend to the device at all. Once it was turned "off" the device disappeared from practitioner attention, not becoming the focus of attention again until the very end of the sequence. The device was absent from the mindset of the practitioners. This is not to say that the practitioners were idle or inattentive, indeed they engaged in looking elsewhere for sources of low blood pressure and carrying out the activities to manage that anomaly.

But "off" and a blank screen did not mean the device could not deliver drug to the patient—the "on-off" push button powers down the device; powering down the device does not guarantee the tubing is clamped shut, fluid flow is possible.

In hindsight it is easy to say if only this aspect of the situation or that aspect of device design were different the near miss would have been avoided. If the device had been more robustly designed, users might notice when situations could be near the boundaries of the assumptions about the device capabilities and the situations that can arise. If the practitioners had more experience and better feedback about how the device could act differently than intended, they might have examined the device as a (potentially) active influence in the situation. If the situation had been less fast paced, practitioners may have been able to dig out that the device was active though "off". But the point of the case is to shift our focus as researchers and designers to the challenge of supporting skilled control of attention.

TEMPO

Understanding how disturbances cascade represents a part of analyzing how external events pace practitioner activities. Note how periods of event-driven or externally paced activities can be intermingled with periods of more self-paced activity where the practitioners have control of how to invest their resources and efforts over tasks (Hollnagel, 2002). For example, space shuttle operations have high-tempo periods tied to launch, docking, or entry, but in between, depending on the kind of anomaly and the mission plan, orbit is a phase where there is a degree of more self-paced activity for flight controllers and mission engineers.

Discovering how a JCS works depends on understanding the different rhythms that play out at different scales of operation and comprehending how periods of higher tempo, more event-driven operations are intermingled with periods of lower tempo, more self-paced activity. To do this, observation and functional modeling are sensitive to what are leading and lagging factors in these dynamic processes (e.g., triggering events that produce a subsequent cascade of effects).

When we see tempo as a fundamental aspect of JCS at work, we see the fundamental demand on how to keep pace with (or stay ahead of) the changing situation. When we use this question as a way to abstract observations from

multiple natural laboratories, the role of anticipation in expert performance comes to the fore. The mystery for researchers is to understand and support how expertise is tuned to the future, while, paradoxically, the data available is about the past.

The target can be captured in a principle that could be called *Avery's Wish* (Woods, 2002). Avery is an actual high-status and high-skill practitioner who is often asked by vendor engineers and usability specialists what features he would like to see in next generation devices (the questions are framed in terms of features and capabilities of the latest technologies). If the technologists could get better information from users about what features are valuable, they would know how to prioritize feature development and introduction for future products. Finally, tired of being the quarry in the feature game, Avery one day responded, "I yearn for some thing that shows me an image of what the system and the room are going to look like, ten minutes from now."

While this response was useless to the engineers in their search for feature-based requirements, it does in fact capture a fundamental aspect of practice and points to a general direction for exploring what would be useful—innovate new forms of feedback that support anticipation of where the system is headed and what could happen next. Part of the difficulty is how to anticipate future trends without being trapped in predictive models, which may be too brittle, uncertain, and cumbersome to rely on given the variability of real fields of practice and the real consequences of failure.

ESCALATION

There is a fundamental relationship regarding tempo that drives how JCSs adapt (Woods & Patterson, 2000):

> **Escalation Principle**: The concept of escalation concerns a process—how situations move from canonical or textbook to non-routine to exceptional. In that process, escalation captures a relationship—as problems cascade, they produce an escalation of cognitive and coordinative demands that bring out the penalties of poor support for work.

The concept of escalation captures a dynamic relationship between the cascade of effects that follows from an event and the demands for work that escalate in response (Woods, 1994). An event triggers the evolution of multiple interrelated dynamics. The response following an anomaly during ascent of a space shuttle mission (Chapter 8, pp. 69-70, and Figure 7) illustrates the cascade and escalation of demands (the reader can use the case of "Being *Bumpable*" in Chapter 3 to practice charting these cascades themselves).

1. There is a cascade of effects in the monitored process.
A fault produces a time series of disturbances along lines of functional and physical coupling in the process (e.g., Abbott, 1990). These disturbances

produce a cascade of multiple changes in the data available about the state of the underlying process, for example, the avalanche of alarms following a fault in process control applications (Reiersen, Marshall, & Baker, 1988).

2. Demands for cognitive activity increase as the problem cascades.
More knowledge potentially needs to be brought to bear. There is more to monitor. There is a changing set of data to integrate into a coherent assessment. Candidate hypotheses need to be generated and evaluated. Assessments may need to be revised as new data come in. Actions to protect the integrity and safety of systems need to be identified, carried out, and monitored for success. Existing plans need to be modified or new plans formulated to cope with the consequences of anomalies. Contingencies need to be considered in this process. All these multiple threads challenge control of attention and require practitioners to juggle more tasks (increasing the risk of workload bottlenecks).

3. Demands for coordination increase as the problem cascades.
As the cognitive activities escalate, the demand for coordination across people, across groups, and across people and machines rises. Knowledge may reside in different people or different parts of the operational system. Specialized knowledge and expertise from other parties may need to be brought into the problem-solving process. Multiple parties may have to synchronize activities aimed at gaining information to aid diagnosis or to protect the monitored process. The trouble in the underlying process requires informing and updating others—those whose scope of responsibility may be affected by the anomaly, those who may be able to support recovery, or those who may be affected by the consequences the anomaly could or does produce.

4. The cascade and escalation is a dynamic process.
A variety of complicating factors can occur, which move situations beyond canonical, textbook forms (Woods et al., 1990). The concept of escalation captures this movement from canonical to non-routine to exceptional. The tempo of operations increases following the recognition of a triggering event and is synchronized around temporal landmarks, particularly those that represent irreversible decision points. The dynamics of escalation vary across situations. First, the cascade of effects may have different time courses. For example, an event may manifest itself immediately or may develop more slowly. Second, the nature of the responses by practitioners affects how the incident progresses—less appropriate or timely actions (or too quick a reaction in some cases) may sharpen difficulties, push the tempo in the future, or create new challenges. Different domains may have different escalation gradients depending on the kinds of complicating factors that occur, the rhythms of the process, and consequences that may follow from poor performance.

5. Interactions with support systems occur relative to demands.
Interactions with support systems (computer based or via other technologies) occur in the context of these escalating demands on knowledge and attention, monitoring and assessment, communication and response. In situations within the envelope of textbook competence, technological systems seem to integrate smoothly into work practices, so smoothly that seemingly little work is required for human roles. However, patterns of distribution and coordination of this work over people and machines grow more complex as situations cascade. Thus, the penalties for poor coordination between people and machines and for poor support for coordination across people emerge as the situation escalates demands.

The difficulties arise because interacting with the technological devices is a source of workload as well as a potential source of support. Interacting with devices or interacting with others through devices creates new burdens that can combine in ways that create bottlenecks. Practitioners are placed in an untenable situation if these new burdens occur at times when practitioners are busiest on critical tasks, if these new attentional demands occur when practitioners are already plagued by multiple voices competing for their attention, or if these new sources of data occur when practitioners are overwhelmed by too many channels spewing out too much competing data.

As active, responsible agents in the field of practice, practitioners adapt to workaround these bottlenecks in many ways—they eliminate or minimize communication and coordination with other agents, they tailor devices to reduce cognitive burdens, they adapt their strategies for carrying out tasks, they abandon some systems or modes when situations become more critical or higher tempo (*JCS-Foundations*, pp. 106-107). Woods et al. (1994) devote a chapter to examples of these workload bottlenecks and the ways that people tailor devices and work strategies to cope with forms of technology-induced complexity (cf., Sarter et al., 1997 for a summary on cockpit automation or Cook & Woods, 1996 for a study on operating room information technology).

Properties of JCSs, such as escalation, illustrate how typical questions about allocating tasks to either the human or the automation are fundamentally misguided and unproductive (*JCS-Foundations*, pp. 121-124). Instead of asking, "Who does what with this in isolation?" CSE asks, "How does the JCS adapt to changing demands?" For example, the advent of technology for unmanned aerial vehicles (UAVs) led most technologists and development managers to ask questions about how many human operators it takes to operate a UAV and to ask for assistance in reducing the ratio of people to UAVs. This is the wrong question and produces a wave of reactive ergonomics work that becomes quickly outdated by the next shift or leap in technology (yet another case of "dustbin human factors"). As a result, the search for design concepts becomes stuck in copying over old interfaces and roles into new situations and capabilities that leaves all stakeholders dissatisfied (in the UAV case, recruiting pilots to tele-operate UAVs).

The useful direction is to ask how the system of agents adapts to recognize and handle anomalies and opportunities. How is the system of agents able to escalate

work to match cascading effects and tempo variations when situations challenge the envelope of textbook competence? Designing and testing to this target opens whole new territories for innovation such as, in the expanding technology for UAVs, how to help recognize and adapt to disruptions to the plan in progress, how to see if the automation's plan fits the situation that is unfolding, how to support integrating multiple feeds that capture different aspects of the remote situation of interest (Woods et al., 2004).

COUPLING

As one considers the cases and patterns in JCSs at work, note the role of coupling as a form of complexity. Coupling refers to the degree of interconnections between parts in a process to be controlled. Tighter coupling across parts of a process produces more complex disturbance chains following disruptions. Disturbances propagate more quickly and further through the monitored process. Tighter coupling intensifies all of the demands in anomaly response, including factors related to tempo of operations, coordination across parties, switching among multiple threads (control of attention), avoiding workload bottlenecks.[8]

Perhaps the most significant consequence of tighter coupling is the increase in potential side effects of any event that occurs, action that is taken, or plan that is modified. The development of techniques for goal-means analysis (Rasmussen & Lind, 1981; Rasmussen, 1986; Woods & Hollnagel, 1987; Vicente, 1999; Lind, 2003) are important because they provide a means to map couplings across functional and physical levels of description and capture how disrupting events produce cascades of effects (Woods, 1994).[9] By providing information on various forms of coupling, goal-means analyses contribute to tracking what side effects occur and how goals interact in specific situations.

---

[8] Note that organizational theorists and CSE both use the term coupling, but in different senses. CSE restricts coupling to refer to properties of the monitored process that affect how disturbances propagate which in turn creates and intensifies the demands of work. Organizational theorists use the label coupling also to refer to the spread of information and decision making across the management structure.

[9] It is important to avoid confusing means for end in analyses of JCSs at work. The abstraction hierarchy is one analysis mechanism and it is used to capture relationships across goals, functions, and physical mechanisms in a work domain (Rasmussen, 1986; Vicente, 1999). Describing a work domain in the form of an abstraction hierarchy can become an end in itself in cognitive task synthesis, rather than means to map couplings and help decompose the relationships across strategies, demands, and artifacts in work. The criterion to judge any cognitive task synthesis is always: does it provide a functional account how the behavior and strategies are adapted to the goals and constraints of the field of practice. Specific ways to analyze a work domain are useful to the degree they contribute to that end, for example, by helping to map contexts, side effects, interacting goals, and tradeoffs.

Tighter coupling means that disruptions and actions produce multiple effects which challenge coordination and resilience:

• Anomaly recognition and diagnostic search become more difficult since some of the follow-on disturbances can be physically or functionally distant from the original triggering event—"effects at a distance." As disrupting events produce multiple effects, the ability to discriminate red herrings from important "distant" indications will be more difficult (e.g., the scenario in Roth et al., 1992). This exacerbates the already difficult demand to focus in on the significant subset of data as context changes (Woods, 1995a).

• Increased coupling creates more situations where different goals interact and conflict. This means that there are more dilemmas or tradeoffs to be handled and resolved in practice (Woods et al., 1994).

• As coupling goes up, a single action will have multiple effects. Some of these will be the intended or main effects while others will be "side" effects. These side effects must be considered in replanning following disruptions to the plan in progress or assessing the actions of other agents, human or machine. Missing side effects in diagnosis, in planning, and in adapting plans in progress to cope with new events is a common vulnerability that contributes to failure in highly coupled systems (Woods & Shattuck, 2000). For example, when an event disrupts a plan in progress for a tightly coupled process, there are more reverberations to manage during replanning.

• Coupling expands the number of threads or lines of activity and reasoning that can be intertwined during work (directed attention and managing multiple threads). This increases the demand for effective control of attention in order to switch the focus of attention as conditions and priorities change (Woods, 1995b).

• When distant parts are coupled, coordination demands increase as practitioners in one role must know about other parts of the process, know more about the work carried out by other roles that monitor manage those parts, and know more about how activities in their scope of responsibility affects others and visa versa.

One of the innovations in support for JCSs at work (affordances) that has resulted from work in CSE is the need to develop side effect tracking displays for highly coupled systems (e.g., Woods & Hollnagel, 1987; Watts-Perotti & Woods, 1999).

PREMATURE NARROWING

When we step back and examine studies on anomaly response, one theme that recurs is the danger of premature narrowing. Through careful observation of the development of expertise and coordinative mechanisms in a JCS, one begins to notice the presence of strategies for avoiding this basic vulnerability (Cook et al., 2000). Findings have highlighted the danger of becoming stuck in one assessment

and being unable to revise the assessment even as new evidence comes in or situations change (Woods et al., 1987). Studies of hypothesis generation in diagnostic reasoning found that the key is broadening the set of possible explanations to be considered (Gettys et al., 1987). Studies of professionals in information analysis found that premature narrowing was a basic vulnerability as analysts moved to new computer-based systems in order to cope with massive increases in data availability and were challenged by new tasks outside their home base of experience (e.g., Patterson et al., 2001; Elm et al., 2005).

All processes for anomaly response and information analysis must, within some horizon, funnel-in on key sources, on the basic unexpected finding to-be-explained, on the storyline that explains the unfolding events and evidence gathered, and on the critical action or plan revision that needs to be undertaken to accomplish goals in changing situations. The danger is a premature narrowing that misses or discounts evidence that would lead to revision. Experienced practitioners develop "broadening" checks that they combine with the normal funneling processes to reduce the risk of premature narrowing or closure. In effective performance, a JCS adjusts the sequence of funneling-in plus broadening checks to converge in a timely manner while remaining sensitive to the need to revise previous assessments. Using broadening checks is a balance between the need to be sensitive to the potential for misassessment and the need to accomplish work within time and resource bounds inherent in evolving situations or imposed to meet organizational goals. Notice that funneling-in plus broadening is another example of Neisser's perceptual cycle (*JCS-Foundations*, p. 20) and demonstrates how this concept is fundamental to analysis and synthesis of JCSs at work.

Criteria that assess broadening versus premature narrowing also are key in studies of collaboration. Studies of human-automation coordination find that poor collaborative architectures narrowed the range of data considered and hypotheses explored (Layton et al., 1994). Studies of human collaboration find that diversity across participants improves problem-solving performance (Smith et al., 1997; Hong & Page, 2002). Studies of error detection and correction point to the need for collaborative cross checks across the multiple practitioners involved in providing care (Patterson et al., 2004; or Fischer & Orasanu, 2000, for aviation). Thus, one part of the value of effective collaborative interconnections lies in how they can broaden focus, reduce mis-assessments, and support revision.

REFRAMING

The dynamic character of anomaly response has revealed the importance of revision. When anomaly response breaks down it is often associated with an inability to revise plans and assessments as new evidence arrives and as situations change. Failures to revise, as a basic vulnerability in JCSs, are a form of under-adaptation.

In the perceptual cycle (*JCS-Foundations*, p. 20), data noticed about the world trigger frames (including purposes) that account for the data and guide the search for additional data. At the same time, the current framing orients the observer to the

world, changing what counts as data (e.g., the discussion of unexpected events). Both activities occur in parallel—data generating frames, and frames defining what counts as data. The revision of assessments is more than a simple adjustment to the current assessment, but rather the more difficult process of reframing (Klein et al., in press). For example, a review of research on problem detection found that failures of problem detection are not so much failures to detect an early indicator, but rather they are failures to re-conceive or redefine the situation (Klein et al., 2004).

Fixating on a situation assessment, thus missing or discounting new evidence is another example of the need for and difficulty of reframing. The basic defining characteristic of fixations is that the immediate problem-solving context has biased the problem-solver in some direction—framing. The data on successful and unsuccessful revision of erroneous situation assessments shows that reframing usually takes a person with a fresh point of view on the situation (Woods et al., 1987). For example, after the Three Mile Island accident (Kemeny, 1979), the nuclear industry took actions to avoid failures to revise assessments by adding a new role to the control room team with a different background and viewpoint (called the Shift Technical Advisor) and giving the operators new kinds of representation of the behavior of the plant relative to goals (new displays and procedures organized around safety functions).

As in this example, one can aid reframing by changing how different people in the system coordinate their roles to try to ensure a fresh point of view, i.e., one that is unbiased by the immediate context, or to ensure effective cross checks (Patterson et al., 2004). One can aid reframing by providing new forms of feedback on the behavior of the monitored process that captures events, future directions (what could happen next), views that integrate data around models of the work domain (Vicente, 1999), multiple views that provide contrasting perspectives on what data mean, and displays on goal achievement (Woods, 1995a).

DILEMMAS

… to see "very well that it was necessary to perish in order not to perish; and to expose oneself to dangers of all kinds, in order to avoid all dangers.
Jesuit Relations (1656-57)

The rising complexities of practice create or exacerbate competing demands such as conflicting goals. Multiple, simultaneously active goals are the rule rather than the exception for virtually all domains in which expertise is involved. Practitioners must cope with the presence of multiple goals, shifting between them, weighing them, choosing to pursue some rather than others, abandoning one, embracing another. Many of the goals encountered in practice are implicit. Goals often conflict. Sometimes these conflicts are easily resolved in favor of one or another goal, sometimes they are not. Sometimes the conflicts are direct and irreducible, for example, when achieving one goal necessarily precludes achieving another one. But there are also intermediate situations, where several goals may be

partially satisfied simultaneously. An adequate analysis of JCSs at work requires explicit description of the interacting goals, how they create tradeoffs and dilemmas, and how these are resolved in practice (Woods et al., 1994).

Some dilemmas are embedded in the very nature of technical work in that field of practice. These include technical dilemmas that arise from demands inherent in the process to be managed and controlled. For example, some situations demand that anesthesiologists attempt to keep blood pressure both high and low to meet different goals on the patient's cardiac system. For an anesthetized cardiac patient, a high blood pressure works to push blood through the coronary arteries and improve oxygen supply to the heart muscle. On the other hand, because increased blood pressure adds to cardiac work, a low blood pressure is desirable to reduce cardiac work. The appropriate blood pressure target adopted by the anesthesiologist depends in part on the practitioner's strategy, the nature of the patient, the kind of surgical procedure, circumstances within the case that may change (e.g., the risk of major bleeding), and the negotiations between different people on the operating room team (e.g., the surgeon who would like the blood pressure kept low to limit the blood loss at the surgical site).

Organizations constrain and pressure practice in ways that create or intensify dilemmas. For example, consider how increased pressure for efficiency creates more same-day surgeries, which may exacerbate dilemmas. In the daily routine, the goal of discovering important medical conditions with anesthetic implications before the day of surgery may drive the practitioner to seek more information about the patient. Each hint of a potentially problematic condition provides an incentive for further tests that incur costs (e.g., the dollar cost of the tests, the lost opportunity cost when a surgical procedure is cancelled and the operating room goes unused for that time, the collaborative costs of disgruntled surgeons). The goal of minimizing costs, in contrast, provides an incentive for the use of same-day surgery even though this constrains preoperative evaluations. Thus, the medical practitioners face a dilemma (Woods, 2006a). Some practitioners in some situations might not follow up hints about some aspect of the patient's history because to do so would impact the usual practices relative to throughput and economic goals. This is especially true since following up the hint will only rarely lead to important information, though the delay surely will interrupt the workflow and incur costs. Other practitioners will adopt a defensive stance and order tests for minor indications even though the yield is low, to be on the safe side. This generates increased costs and incurs the wrath of their surgical colleagues for the delays thus generated.

Organizational factors at the blunt end of systems shape the world in which practitioners work by influencing the means available for resolving dilemmas. In aviation, an aircraft is de-iced and then enters the queue for takeoff. After the aircraft has been de-iced, the effectiveness of the de-icing agent degrades with time. Delays in the queue may raise the risk of ice accumulation. However, leaving the queue to go back to an area where the plane can be de-iced again will cause additional delays, and in addition, the aircraft will have to re-enter the takeoff queue again. Thus, the organization of activities (where de-icing occurs relative to where queuing occurs in the system) can create conflicts that the practitioners must resolve because they are responsible at the sharp end of the system. Unfortunately,

there have been aircraft crashes where, in hindsight, crews accepted delays of too great a duration and ice did contribute to a failed takeoff (Abbott, 1996). In general, balancing tradeoffs created by conflicts of safety versus production goals can prove difficult, as turned out to be the case in the Columbia space shuttle accident (CAIB, 2003; Woods, 2005b).

Some tradeoffs emerge in the nature of demands placed on the work of JCSs. In anomaly response, for example, we have seen how there is a trade-off with respect to when to commit to a course of action. Practitioners have to decide whether to take corrective action early in the course of an incident with limited information, to delay the response and wait for more data to come in, to search for additional findings, or to ponder additional alternative hypotheses (see Figure 11). In control of attention, there is the trade-off between being too easily interrupted by new signals or events and being too focused on the current priority, which affects the balance between the risk of vagabonding from thread to thread incoherently and the risk of failing to revise an assessment as the situation changes (Woods, 1995b).

Practitioners also trade off between following standard routines or adapting the routine to handle the particular situation they face in order to meet the intent behind the plan or procedure (cf., Shattuck & Woods, 2000). Do the standard rules apply to this particular situation after a disrupting event or when some additional complicating factor is present? This is the problem of coordinating the distant but global perspective of supervisors/management with the local and up-to-date perspective of the operator on the scene. Mis-balancing the trade-off risks the failures of under-adaptation—continuing to apply a plan that doesn't fit the situation at hand—or over-adaptation—ad hoc adaptations to disruptions that fail to take into account the broader goals and constraints. Supervisors and the larger organizational context must determine the latitude or flexibility they will give actors to adapt plans and procedures to local situations, given the potential for surprise in that field of activity. Supervision that establishes centralized control inhibits local actors' adaptations to variability, increasing the vulnerability to under-adaptation failures. At the other extreme is supervision that provides local actors with complete autonomy. In the latter case, the goals and constraints important in the remote supervisors' scope are disconnected from the activity and decision making of local actors. As a result, the response across multiple local actors may not be coordinated and synchronized properly, increasing vulnerability to over-adaptation failures. The systems concept of resilience (Hollnagel et al., 2006) suggests that organizations look for evidence of gaps between distant plans and the factors that challenge those plans. A large gap places demands on practitioners to adapt in order to resolve the dilemmas and conflicts.

Dilemmas also arise in advisory interactions across agents (Roth et al., 1987). A machine expert recommends a particular diagnosis or action, but what if your own evaluation is different? What constitutes enough evidence that the machine is wrong to justify disregarding the machine expert's evaluation and proceeding on your own evaluation of the situation?

The dilemmas may be resolved through conscious effort by specific teams, or practitioners can also simply apply standard routines without deliberating on the nature of the conflict. In either case, they may follow strategies that are robust (but

still do not guarantee a successful outcome), strategies that are brittle (work well under some conditions but are vulnerable given other circumstances), or strategies that are very vulnerable to breakdown. Uncovering the dilemmas that occur in practice and how they are resolved through practice will reveal how behavior and strategies are adapted to constraints (and this provides a good tip for designing problems/scenarios to use in studies intended to discover how a JCS works).

Value of mapping goal—means relationships in the work domain lies in how the analysis reveals dilemmas and characterizes how multiple factors come together to actualize conflicts in particular situations.

OVER-SIMPLIFICATIONS

The cases discussed to this point also reveal a larger pattern about coping with complexity. When confronted by complexity, people have a tendency to simplify (JCS-Foundations, p. 82). This is a locally adaptive response wherever the complexity originates. The stories of adaptation in JCSs reveal simplifications as a coping response when a new artifact is misfit to demands such as control of attention, when an increase in coupling produces additional disturbance chains cascading from a disrupting event, or when management creates double binds for practitioners wrestling in specific situations with conflicts between simultaneously important but inconsistent goals.

The problem of modern systems is a rise in complexity that results from successful adaptation to the pressure to be "faster, better, cheaper"—or the Law of Stretched Systems (see NASA, 2000; Woods, 2005b; Woods, 2006a). Driven by demands for new performance levels and by pressures to reduce resources; plus, fuelled by new capabilities from the expanding powers of new technology for connectivity, for data pick-up/collection/ transmission, and for extending our presence into remote environments, fields of practice face new complexities. In addition, the rising tide of complexities to meet the pressure to be "faster, better, cheaper" undermines the viability of previous simplifications adopted to manage the demands and tradeoffs in work. CSE arose as a response to this growth in complexity (*JCS-Foundations*, pp. 3-5).

Feltovich and his colleagues have studied how people understand complex concepts (often highly trained practitioners such as cardiologists), and have found a set of over-simplification tendencies at work (cf. also Rasmussen, 1986, on simplifications as shortcuts in decision making). Table 4, taken from Feltovich, Spiro & Coulson (1997), captures some of the basic dimensions along which simplification and over-simplification can occur. First, note that these tactics no longer work well in the face of the complexities of modern work systems under "faster, better, cheaper" pressures. They become over-simplifications ill-suited to handle the situations, risks and uncertainties of modern systems, e.g., how the increased coupling or interactions between parts produce more complex cascades of effects to be tracked and controlled. Second, close examination of Table 4 reveals that simplifications along those dimensions are all essentially **narrowing** heuristics. When the work of JCSs is based on these heuristics, the system is vulnerable or

exposed to failure by missing side effects of an action given the interconnections across parts, through poor synchronization of multiple lines of activity, by getting stuck on seeing one factor as the cause of outcomes when results actually emerge from the interactions and contributions of multiple factors.

Table 4. Over-simplifications (from Feltovich, Spiro & Coulson, 1997)

1. Discreteness/continuity. Do processes proceed in discernible steps, or are they unbreakable continua? Are attributes adequately describable by a small number of categories (e.g., dichotomous classifications like large/small), or is it necessary to recognize and utilize entire continuous dimensions (e.g., the full dimension of size) or large numbers of categorical distinctions?

2. Static/dynamic. Are the important aspects of a situation captured by a fixed "snapshot," or are the critical characteristics captured only by the changes from frame to frame? Are phenomena static and scalar or do they possess dynamic vectorial characteristics?

3. Sequentiality/simultaneity. Are processes occurring one at a time, or are multiple processes happening at the same time?

4. Mechanism/organism. Are effects traceable to simple and direct causal agents, or are they the product of more system-wide functions. Can important and accurate understandings be gained by understanding just parts of the system, or must the entire system be understood for even the parts to be understood well?

5. Separability/interactiveness. Do processes occur independently or with only weak interaction, or is there strong interaction and interdependence?

6. Universality/conditionality. Do principles hold in much the same way (without the need for substantial modification) across different situations, or is there great context-sensitivity in their applicability?

7. Homogeneity/heterogeneity. Are components or explanatory schemes uniform (or similar) across a system—or are they diverse?

8. Regularity/irregularity. Is a domain characterized by a high degree of routinizability across cases, or do cases differ considerably from each other even when commonly called by the same name? Are there strong elements of symmetry and repeatable patterns in concepts and phenomena, or is there a prevalence of asymmetry and absence of consistent pattern?

9. Linearity/nonlinearity. Are functional relationships linear or nonlinear (i.e., are relationships between input and output variables proportional or non-proportional)? Can a single line of explanation convey a concept or account for a phenomenon, or are multiple and overlapping lines of explanation required for adequate coverage?

10. Surface/deep. Are important elements for understanding and for guiding action delineated and apparent on the surface of a situation, or are they more covert, relational, abstracted?

11. Single/multiple. Do elements in a situation afford single (or just a few) interpretations, functional uses, categorizations, and so on, or do they afford many? Are multiple representations required (multiple schemas, analogies, case precedents, etc.)?

The functional synthesis of anomaly response points to factors that make it necessary to avoid the simplifications in Table 4. For example, as tempo and coupling increase, the danger of failures to revise increases. Thus, we begin to see a pattern where a routine, plan or algorithm is deployed correctly, but in the wrong situation—i.e., the actual the situation demands a different response (Mitroff's 1974 error of the third kind). This means design needs to provide new kinds of artifacts that promote observability of dynamic, multi-factor processes, new kinds of artifacts for coordination that support escalation of knowledge and expertise as a situation cascades, new forms of broadening checks and cross-checks that enhance revision of assessments and plans as situations change or disruptions occur (cf., *JCS-Foundations*, pp. 87-91).

These over-simplifications captured by Feltovich's research do not apply just to work at the sharp end of systems (Figure 2). Designers and managers at the blunt end of systems are quite vulnerable to over-simplification tendencies, as well, as they make investment and development decisions that affect the future of operations (Woods & Dekker, 2000). Over-simplification tendencies have dominated the investigation of how systems fail (the red herring of "human error" and linear causality in accident and root cause analysis (Hollnagel, 1993; Woods et al., 1994, chapter 6; Dekker, 2002; Hollnagel, 1998; Hollnagel et al., 2006).

As we begin to consider stories of work about miscoordination between automation and people in Chapter 10, note that a second story line is going on in parallel. The "second" story (Cook et al., 1998; Woods & Tinapple, 1999) is a story of how the blunt end is trapped in over-simplifications about how new technological powers affect the future of work and in over-simplifications about how people and machines coordinate as joint cognitive systems (see Roesler et al., 2001; Woods et al., 2004; Feltovich, Hoffman et al., 2004).

To summarize: Patterns in Multi-Threaded Work

Multi-threaded work challenges simplification heuristics. These heuristics become over-simplifications that expose the JCS to new vulnerabilities for failure. Describing, modeling and supporting multi-threaded work requires new ways to represent dynamic balances as in control of attention, new ways to manage cascades of effects from disrupting events, and new ways to re-conceptualize the problem to be solved or the key goal to be achieved.