Complex Adaptive Systems —Resilience, Robustness, and Evolvability: Papers from the AAAI Fall Symposium (FS-10-03)

How do Systems Manage Their Adaptive Capacity to Successfully Handle Disruptions? A Resilience Engineering Perspective

Matthieu Branlat and David D. Woods

Cognitive Systems Engineering Laboratory
The Ohio State University
Integrated Systems Engineering
210 Baker Systems
1971 Neil Avenue
Columbus OH 43210
branlat.2@osu.edu, woods.2@osu.edu

Abstract

A large body of research describes the importance of adaptability for systems to be resilient in the face of disruptions. However, adaptive processes can be fallible, either because systems fail to adapt in situations requiring new ways of functioning, or because the adaptations themselves produce undesired consequences. A central question is then: how can systems better manage their capacity to adapt to perturbations, and constitute intelligent adaptive systems? Based on studies conducted in different high-risk domains (healthcare, mission control, military operations, urban firefighting), we have identified three basic patterns of adaptive failures or traps: (1) decompensation – when a system exhausts its capacity to adapt as disturbances and challenges cascade; (2) working at cross-purposes – when sub-systems or roles exhibit behaviors that are locally adaptive but globally maladaptive; (3) getting stuck in outdated behaviors – when a system over-relies on past successes although conditions of operation change. The identification of such basic patterns then suggests ways in which a work organization, as an example of a complex adaptive system, needs to behave in order to see and avoid or recognize and escape the corresponding failures. The paper will present how expert practitioners exhibit such resilient behaviors in high-risk situations, and how adverse events can occur when systems fail to do so. We will also explore how various efforts in research related to complex adaptive systems provide fruitful directions to advance both the necessary theoretical work and the development of concrete solutions for improving systems' resilience.

Introduction

Resilience engineering (Hollnagel, Woods and Leveson, 2006) focuses on work systems as an instance of human systems and complex adaptive systems. In particular, it studies how the performance of systems is affected by disruptions, and seeks to understand how the characteristics of these systems make them more prone to success or to failure in their respective environments (especially in terms of the safety of their components or of the elements they are responsible for). The field borrows concepts and results from a variety of other fields related to the study of complexity, such as control theory, organizational and policy sciences. In return, resilience engineering aims at contributing to the study and improvement of work systems, as well as to the development of knowledge about complex adaptive systems.

A large body of research describes the importance of adaptability and adaptive behavior for systems to be able to face the variability of their environment without loosing control or totally collapsing (Ashby, 1956; Ostrom, 1999; Woods and Hollnagel, 2006; Woods and Wreathall, 2008; Weick and Sutcliffe, 2001). Work systems pursue various goals that require them to control various processes in a given environment (e.g., a hospital treats sick patients, nuclear power plants produce energy from a nuclear reaction, aviation crews fly passengers, firefighters respond to fires). In order to do so, systems rely on models of their tasks and of the environment. Any model is however a

---

necessary simplification that aims at reducing and managing the uncertainty and complexity of the real world for pragmatic purposes (Hollnagel, 2009). As a result, the conditions for operation are underspecified. The predictability associated with the processes under control and the environment is never perfect, even in predominantly closed worlds like manufacturing plants. Any work system therefore needs to be able to adapt to a variety of perturbations. Given the underspecification, these perturbations may be surprising. In addition to the fundamental limitation of the models they rely on, work systems are also confronted with the evolving nature of their environment, as well as their own evolutions (e.g., when new technology, resources or working rules are introduced). Both dynamics constitute sources of new challenges to their normal way of operating, i.e., new – and often unforeseeable – vulnerabilities.

At some level of analysis systems are human systems since it is people who create, operate, and modify that system for human purposes, and since people, not machines, gain or suffer from the operation of that system. One of the central efforts of the resilience engineering field consists of studying and understanding challenges to adaptation in complex socio-technical systems in order to better support human adaptive behavior, in part through technological development (Hollnagel, Woods and Leveson, 2006). From both a performance and safety standpoint, we are interested in how systems adapt their plans in progress, i.e., the way they function when expected or unexpected perturbations occur. To a large extent, such adaptations are managed by the human elements of the work systems, even when they make ample use of advanced technology. An important reason for this is the context-sensitivity issue (Woods and Hollnagel, 2006). Operators fill the gaps by adapting to the real conditions of operations and their dynamics (Cook and Rasmussen, 2005). In particular, we are interested in situations where plan adaptations are not successful, either because systems fail to recognize the need for adaptation, or because the adaptive processes themselves produce undesired consequences (Woods and Shattuck, 2000; Dekker, 2003; Lagadec, 2010). A central question is then: *how can systems better manage their capacity to adapt to perturbations, and constitute intelligent adaptive systems?*

We will present how expert practitioners exhibit such resilient behaviors in high-risk situations, and how adverse events can occur when systems fail to do so. We will also explore how various efforts in research related to complex adaptive systems provide fruitful directions to advance both the necessary theoretical work and the development of solutions for improving systems' resilience.

Basic Patterns of How Adaptive Systems Fail

We have identified three basic patterns of adaptive failures or traps based on studies conducted in different high-risk domains such as healthcare, mission control, military operations and urban firefighting (Woods and Branlat, 2011).

(1) Decompensation: Exhausting Capacity to Adapt as Disturbances/Challenges Cascade.

This first pattern is observed in settings implementing some form of supervisory control, in which a supervising role is in charge of controlling a dynamic process through a variety of actors (human, technological, biological). This pattern corresponds to an escalation of demands while the system monitoring the process of interest is not capable of adapting and acting fast enough upon an initial set of disturbances. In a complex system, the amount of coupling between elements or sub-systems creates the conditions for a cascade of disturbances resulting in rapidly increasing demands. Under such circumstances, components are stretched to their performance limits and the system's overall control of the situation collapses abruptly. This loss of control creates conditions that can develop into an incident, i.e., severely impair performance and compromise safety. From the perspective of the supervisor, the critical information is how hard the agents under supervision are working to maintain control on their parts of the process. Difficulties in dealing with escalating demands consist of avoid falling behind the tempo of events by recognizing early enough the trend towards a loss of control, and being able to switch to new modes of functioning in time, such as investing the needed resources (including in the case of unanticipated demands).

(2) Working at Cross-Purposes: Behavior that is Locally Adaptive, but Globally Maladaptive.

This pattern is directly related to the complexity of the environments in which real-world systems exist, and of the systems themselves. Systems' or subsystems' behaviors, which exist in the context of networks of interdependencies (functional, structural, temporal) and cross-scale interactions, have implications at a larger scale than simply at the level of elements producing the behaviors. The pattern is exemplified by the widely discussed *tragedy of the commons* (Hardin, 1968; Ostrom, 1999) in which populations gradually deplete shared physical resources by over-exploiting them in the pursuit of short-term production goals. Since the subsistence of these populations also depends on these resources, this behavior turns out over time to be counter-productive. More generally, this second pattern relates to the difficulty of designing systems that provide both stability, in order to ease the synchronization of distributed actions to reach a common goal, and flexibility, in order to adapt in the face of changing conditions of operation (Branlat, Fern, Voshell and Trent, 2009; Grote, Zala-Mezö and Grommes, 2004). Adaptation in tightly coupled systems indeed increases needs for coordination in order to avoid working at cross-purposes. Critical challenges here correspond to the difficulty in recognizing how goals pursued by different parts of the system might conflict, and how progress at one echelon of the system introduces higher demands at another echelon.

(3) Getting Stuck in Outdated Behaviors: the World Changes but the System Remains Stuck in what were Previously Adaptive Strategies.

This third pattern is at play when a system gets stuck in implementing behaviors that were successful in the past, and fails to recognize that the conditions for their implementation are no longer met. The pattern relates to breakdowns in how systems learn, either from past experience or dynamically, as situations unfold. In the context of work systems, this pattern relates especially to the lack of capacity to revise plans in progress in the face of disruptions or opportunities (Hollnagel and Woods, 2005). Plans might correspond to the implementation of enforced rules of work such as procedures, which are typically appropriate for the very specific context of the envisioned operational world, but often fail to address the complexity of the real world (Woods and Shattuck, 2000). The limitations of plans and the difficulty to revise them can both be related to simplification processes that are needed to handle complex phenomena (Feltovich, Coulson and Spiro, 2001). Such processes become problematic when they correspond to oversimplifications, such as when signs are discarded when they contradict an on-going assessment of a situation at hand. This leads decision makers to overhomogenize conditions for the implementation of behaviors, although they present critical differences. Key difficulties lie in the fact that the signs of the failure of previously accepted strategies might be ambiguous and uncertain.

Successes and Failures to Address the Patterns in Real Work Situations

In this paragraph, we will borrow examples from various resilience-related studies, including our own, to illustrate how the three patterns impact the performance of systems and compromise safety, but also how resilient mechanisms address them. Each of these studies takes place in a high-risk domain in which failures can have dramatic human consequences for the members operating, but also for populations under their responsibility. Following a typical approach in resilience engineering, we will look at critical past incidents to understand how these systems manage and sometimes fail to adapt their plans to perturbations they experience (exceptionally or routinely). Such incidents – which are not necessarily the most dramatic accidents – are chosen because they reveal challenges these systems face, and uncover patterns of brittleness and, often, resilience (brittle points and signs of resilience in the face of high risk). The section is broken down following the description of patterns above for practical purposes. However, cases usually don't correspond to one single pattern; all three patterns are often at play to some degree. These patterns can also constitute by themselves competing goals since trying to address challenges related to a specific pattern can end up contributing to another type of failure (think for example of the first and second patterns, as we will illustrate below).

Coping with Increasing Demands

The first pattern described above relates to how systems detect signs that they are running out of capacity to adapt to disturbances and that they need to mobilize new resources.

As a domain in which systems operate close to their maximum capacity on a regular basis, healthcare represents an interesting work environment to investigate. Cook (2006) proposes a description of the task of the *bedmeister* through the account of a minor incident. In Intensive Care Units, the bedmeister is the role in charge of managing the beds, i.e., the physical locations and resources associated (human or equipment), in order to provide beds to the patients in need of intensive care. A critical characteristic of the management of patients in ICUs is that it fundamentally involves planned and unplanned activities. A large number of patients reach the ICU after a scheduled surgical procedure, which represents a largely anticipatable load of work. But an equivalently large amount of patients require ICU treatment because they are patients in unexpected degraded states (e.g., a patient from the ward whose parameters worsened, a patient who just went through the ER). Cook focuses in particular on the notion of *bumpable*, jargon used by medical personnel involved in these tasks to identify patients that could be transferred from the ICU to less critical units, thereby generating additional capacity to accept new patients.

From a wider perspective, Cook and Nemeth (2006) describe how the Israeli health system manages the high and unexpected demands of mass casualty events. Events such as suicide bombings in public places present a high potential for cascading into unmanageable situations: casualties are typically severe and high, requiring injured people to be transported and treated quickly; already busy hospital units face heavy disruptions in planned patient care and other tasks; families and friends search for potential victims, seek information and require psychological support; news media require the latest elements of information; etc. However, the Israeli health system has evolved into a system capable of very resilient management of such events. The system's performance relies on the system's capacity to rapidly mobilize large amounts of resources (from ambulances to social workers), on a general tendency to delegate authority at all levels rather than to centralize decisions (e.g., for the dispatch of ambulances to the scene), and on the successful reprioritization of tasks to handle the emergency before returning to normal.

The same issues of being able to recruit more resources to match the demands of the situation in a timely manner exist in other domains. In urban firefighting, incident commanders (IC) are particularly attentive to the potential for new demands to arise (Branlat *et al*., 2009). As an illustration of critical dynamics, the "all hands" signal is used by ICs when they recognize that all units are operating and need dispatchers to promptly provide additional resources. However, given the time needed for additional resources to be available, "all hands" situations can be very precarious since operations are vulnerable to any additional demands that may occur. Cases show that incidents indeed take place in such conditions, especially when resources are requested at the moment they are actually needed. As with Cook's analysis of bed crunches, urban firefighting illustrates the need for these processes to be more anticipatory than reactive in order to be successful. Examples above concern the system's capacity to recruit more of the same type of resources in the face of an increase in scale from normal conditions. But this capacity also concerns the response to different types of demands than in normal conditions, i.e., the ability to integrate news forms of resources. In their study of space mission control, Patterson and Woods (2001) have observed how the system is designed to rapidly call in and benefit from additional types of expertise when anomalies arise. Such ability is part of the resilient mechanisms that make mission control a successful organization when confronted to unanticipated events.

Working in a Network of Interdependencies

Our study of critical incidents in urban firefighting (Branlat *et al*., 2009) focused primarily on coordination issues. In our work, and following previous research on teamwork (Klein, Feltovich, Bradshaw and Woods, 2005), we define coordination specifically as the implementation of the various mechanisms that allow team members to manage interdependencies between their roles and tasks, and conflicts between their goals. Adaptations and replanning are fundamental processes of the firefighters' activity due to the dynamic and uncertain nature of their work environment. However, firefighting operations are also highly distributed and specialized, thereby relying on interdependent functions to fulfill the overall mission. This tightly coupled system also operates in a shared high-risk environment, a source of additional structural and temporal dependencies. Consider the following case:

> After having initiated standard operating procedures and entered a building from the front door to attack a fire in the cellar, an officer signals a better view of the fire from the rear cellar door. The Incident Commander, seizing this opportunity since access from the inside is proving difficult, orders the first engine company (E1) to redeploy the hose line from the front door to the rear of the building. E1 officer signals the new plan to the first ladder company (L1) operating on the first floor, but this message is not received, probably because of the ambient noise. Hose operations initiated by E1 at the rear of the building push heat and flames toward the front the building. The temperature rises and becomes unbearable on the first floor, inciting L1 to evacuate. While evacuating, L1 officer is knocked back and looses consciousness. Upon exiting, L1 firefighters realize their officer is missing, so reenter the first floor to search for him, followed by the second engine company (E2). While L1 and E2 find the officer and attempt to remove him, E1 redeploys to the first floor, and enters the building from the rear, opposing its line to the firefighters coming from the front, delaying the rescue efforts. The fire is ultimately extinguished and the officer removed but suffers severe burn injuries.

In this case, an opportunity prompted replanning. Through a missed communication, overall coordination broke down, and companies' knowledge of each others' positions and operations became erroneous, preventing synchronization and resulting in a dangerous situation, and ultimately an incident. The subsequent emergency management was hindered by the failure to recoordinate in a timely manner, before further operations created more threats and challenges to the rescue efforts (all of this occurs within a few minutes). As illustrated by this case, adaptive processes threaten the maintenance of synchronization between tasks, and constitute sources of failures when interdependencies are not sufficiently managed through coordination processes.

Cook's description of the management of bed crunches reveals how intertwined administrative and medical concerns both take advantage and suffer from the complexity of the system in place. Interdependent services and units in the hospitals provide additional opportunities for identifying bumpable patients and freeing resources for sicker patients, but the level of interdependencies is both a source of difficulties and extra work for coordinating tasks, and the origin of additional constraints on the management of the patient's transfer. Borrowing a concept from the nuclear industry, Cook and Rasmussen (2005) describe the system approaching maximum capacity as "going solid", a state in which situations become both hazardous and difficult to control. The difficulty to control comes in part from the increased and costly demands in coordination since the tightly coupled system is operating in abnormal modes for which plans were not designed.

Detecting the Need to Adapt On-Going Strategies

As the common expression from military decision-making says, "No plan survives first contact with the enemy." In other words, plans are critical resources for action (Suchman, 1987), but are also fundamentally limited. The danger is to stick to plans while the conditions for their success are no longer met. In our study of urban firefighting, various cases show the difficulties, and sometimes the failures, in providing such flexibility. In a particular case (developed in greater detail in Woods and Branlat, 2011), companies kept on implementing standard operating procedures for gaining entry into the apartment on fire, while the critical condition of water availability assumed by these procedures was not met due to various difficulties. Although the sequence of activities deployed is relevant under normal conditions, it was counter-productive in this case since it allowed for the situation to degrade: providing oxygen in the apartment fueled the fire in the absence of means to control it. Along with other factors, this contributed to a loss of control of the situation, and, ultimately, a fatal accident.

Military command-and-control is also confronted with the difficulty to provide the flexibility necessary for the success of missions (Shattuck and Woods, 2000). As with firefighting, the field of operations displays characteristics that challenge the pursuit of mission goals, for instance: coordination can be particularly difficult, if not impossible, between distributed and interdependent units; plans and procedures are very likely to be underspecified; and the potential for unanticipated situations is high. Contrary to the common belief, military decision-making is not necessarily purely hierarchical and relies to a large extent on the delegation of authority at the level of local actors to decide upon the needed adjustments or revisions to pre-planned operations. The notion of "commander's intent" refers to how commanders in a supervisory role "impart their presence by communicating intent to coordinate and adapt underspecified plans and procedures" (ibid.), i.e., explain the goals, rationale or constraints of the mission. Local actors confronted with particular local conditions then adapt accordingly based on their understanding of these larger objectives and constraints. The difficulty for the supervisory roles is then to determine the right level of delegation, monitor how operations are progressing and communicate their intent in clear ways when the situation requires.

Our studies of urban firefighting were to a large extent based on the study of investigation reports following dramatic firefighters' accidents (Branlat *et al*., 2009). Organizations such as fire departments that try to learn from experience are confronted with the difficult choice of selecting what investigations to conduct (i.e., how to prioritize their resources to support their learning objectives). As Hollnagel (2008) writes, "the first step, the selection of which events to investigate, is clearly a crucial one. The selection of what to look at and what to learn from will in a very fundamental way determine how safe an organization can become." Often, organizations want to rely on some objective measure of the criticality of the events they face; when a primary concern is safety, it seems natural to consider that the most critical events are those in which the worst accidents occur. However, it is a limitation to consider that the events that have produced the worst outcomes are also the ones that will produce the most useful lessons. Instances where challenging and surprising situations are managed without leading to severe outcomes, often reveal interesting and innovative forms of adaptations. Unfortunately, selecting the fruitful experiences or events is a difficult choice when it has to be made *a priori* (Hollnagel, 2008; Dekker, 2008, chapter 3). Dekker suggests that expert judgment should be used to determine which events are worth investigating (2008, p. 40). Difficult situations encountered are regularly used as bases for discussions or drills in firehouses, and firefighters and officers are encouraged to use the near-miss reporting system to describe and discuss such situations on a voluntary basis. Such processes constitute a selection process based on the expertise of firefighters who participated in the events. In particular, they constitute opportunities to discuss the relevance of procedures in place, needs for adaptation, or successful modifications to the strategies developed on the fly (e.g., innovative ways of managing difficult or unexpected situations). The difficulty is for organizations to transform such learning opportunities that occur at a local level into knowledge generated at a global scale, including into modifications or revisions of practices that have proved successful in various contexts.

Towards Intelligent Adaptive Systems

Resilience and the Management of Adaptive Capacity

Seeing the patterns at play shows that work systems are capable of avoiding these forms of adaptive traps and manage their capacity to adapt to disruptions. The complexity of system creates new opportunities, but also new challenges, especially in the form of increasing scope and side effects due to wider networks of interdependencies. Unlike more pessimistic approaches (e.g., Perrow, 1984), resilience engineering holds the belief that socio-technical systems are not trapped in complexity: the fundamental human capacity to reflect on their actions and of managing their adaptive capacity represents ways in which work systems can cope with complexity.

From the study of work systems, resilience appears as more than simply bouncing back from disruptions, as it is often defined. Anticipating and avoiding disruptions, rather than simply reacting to them, is actually a much more powerful way to cope with the variability of the real world. In reliable organizations operating in high-risk domains, anticipation doesn't simply mean preparing contingency plans for any imaginable disruption (although this can be useful), but it rather corresponds to the belief that surprising events *will* occur and challenge the system's assumptions about the world (Weick and Sutcliffe, 2001). A resilient system therefore needs to be able to deploy new ways of functioning. Additionally, anticipation also concerns what happens *after* adaptive processes have taken place to cope with a disruption, i.e., what happens when new demands arise as a result of changes in the environment or in the system itself. As a result of these points, we define resilience as fundamentally anticipatory, and related to the capacity to handle the *next* disruption. In short, the resilience of a system corresponds to its adaptive capacity tuned to the future. Woods and Wreathall (2008) claim that any system is capable of two fundamental (and fundamentally different) forms of adaptive behavior. A first form corresponds to the system's built-in adaptive capacity, which allows anticipated perturbations to be handled. Such capacity might come from pre-planned response (including from technological elements or work rules), trained-for perturbations, anticipated resource allocation, etc. However, when surprising events occur, whether the surprise concerns the nature of the disruption or its unusual scale, a second form of adaptive behavior corresponds to the capacity of the system to engage more resources (types of expertise, amount of human resources, time generated by sacrificing production) in order to handle the unusual demands. If all systems show both forms of adaptive capacity to some extent, systems differ widely on the emphasis they put on these capacities. Resilient systems are systems that are careful to provide second order of adaptive capacity.

An important point is suggested by the analysis of urban firefighting and by the account of the bedmeister's task: part of the activity of operators in these systems concerns the creation and maintenance of *margins of maneuver*, i.e., possibilities for adaptation to future demands (Woods and Branlat, 2010; 2011). Practitioners know from experience that new demands and perturbations are likely to arise in these uncertain dynamic domains, although the actual nature and dynamic character of these perturbations is difficult or impossible to predict. A critical task of the firefighting incident commander corresponds to the regulation of adaptive capacity by providing "tactical reserves" (Klaene and Sanders, 2008, p. 127), i.e., an additional capacity to promptly adapt tactics to changing situations. The 'all hands' situations discussed above correspond to situations in which there exists no margin to handle additional demands. Devoting all resources to address demands therefore constitutes a potentially risky decision. In the case of the management of beds in ICUs, hospitals often operate at full capacity. By identifying bumpable patients, the bedmeister re-creates margins of maneuver without which handling unexpected demands would be less efficient and more brittle given the costs and risks associated with "going solid" in this system.

These processes are challenged by the fact that real-world systems exist in a space of fundamental trade-offs identified in the literature on complex adaptive systems: optimality-brittleness, efficiency-thoroughness, acute-chronic. We believe that these fundamental trade-offs are important sources of failure because their management is difficult, or simply because they are often overlooked (i.e., they are not managed for). Doyle and his colleagues (Carlson and Doyle, 2002; Csete and Doyle, 2002; Zhou, Carlson and Doyle, 2005) have identified how successful systems evolve to be "robust yet fragile". Their analyses demonstrate the existence of a trade-off between increasing optimal performance and increasing the capacity to handle unanticipated disturbances. As these two dimensions parallel the first and second order of adaptive capacity (respectively) described by Woods and Wreathall (2008), this result suggests that focusing solely on a system's optimal performance may have a negative impact on its resilience. Hollnagel (2009) describes a similar trade-off in terms of sacrifice decisions to favor either efficiency or thoroughness over the other (experience shows sacrificing thoroughness for efficiency is more common). The quest for optimality or efficiency ends up consuming extra resources that are assimilated to inefficiencies, reducing margins of maneuver and capacity to adapt to unanticipated. In order to manage their adaptive capacity, organizations, as complex adaptive systems, are faced with difficult strategic decisions: Should resources be consumed to address a growing disturbance, or built and sustained to constitute critical reserves? What are the signs or indicators suggesting a shift should be made between these two strategies? How can these strategies be adopted in the context of ambiguous, uncertain and changing environments?

In order to make these strategic decisions, systems need to understand where they are standing in the space represented by the fundamental trade-offs. One attempt to model and represent the system's position in a trade-off space is the dynamic safety model proposed by Rasmussen (Cook and Rasmussen, 2005). This model represents the boundaries of accepted workload, performance and productivity in order to describe how a variety of pressures influence the system in one direction or another. It intends to track how the system's position evolves in time, in the face of changing conditions for operation. Efforts in modeling systemic dynamics can provide useful insight into similar dynamics. For instance, Scheffer and colleagues (2009) have studied how patterns of recovery from disruption can serve as signs of a system's increased difficulty in managing these disruptions. This can constitute an indication of a risk of a future collapse and of the need to adapt behavior (for instance by recruiting more resources). Providing such knowledge first requires fostering learning of how the organization has managed its adaptive capacity in the face of disruptions in past situations. This requires both a culture of organizational learning, as implemented by high-reliability organizations (Weick and Sutcliffe, 2001), and the development of systemic models of incident analysis (e.g., Hollnagel, 2004; Dekker, 2006).

The Need for Polycentric Control

The problem of a system's management of its adaptive capacity can be thought of as a control problem, as suggested by Ashby's law of Requisite Variety (1956). The driving questions are then: How much is the system in control? How can control be amplified? (Hollnagel and Woods, 2005; Woods and Branlat, 2010). Research in large complex systems has identified critical limitations of traditional hierarchical systems in matching the complexity of the real world, and in providing the necessary adaptive capacity (Ostrom, 1999; Woods and Shattuck, 2000; Hollnagel, Woods and Leveson, 2006). The concept of *polycentric control* was developed by Ostrom (1990, 1999) through the study of how real-world complex systems avoid the tragedy of the commons in the management of common pool resources. As such, this conceptual framework is an attempt to understand and avoid an instance of our second maladaptive pattern, created over time by unmanaged vertical interactions between echelons of a system.

The standard view is that the tragedy of the commons is only avoided by allowing central organizations to monitor and regulate these shared resources. Centralized control is believed to act impartially for the common good and to focus on long-term sustainability goals, as opposed to local actors who are not in a position to sacrifice short-term goals as they have too much invested interest. However, several types of challenges or inefficiencies are associated with centralized management: control centers are remote from the scenes of operations, and may have difficulties to acquire precise enough knowledge of local situations (e.g., to make sensible arbitrations between priorities); large exchanges of information are required for decisions to be taken, but channels might be difficult to establish and maintain; control centers create information and decision bottlenecks, especially in the face of cascading events; etc. Ostrom's studies present empirical evidence that successful systems rely to a large extent on local actors to manage shared resources. She emphasizes in particular the fact that local actors, through their greater proximity to the resource than remote experts, develop a much finer knowledge of the resource and its use at their local or regional level. Through this knowledge, they are also more capable of observing trends in the evolution of the resource, therefore more able to adjust their actions to their local environment to meet longer-term sustainability goals. Centralized control, on the other hand, creates the risk of over-homogenizing rules of practice over areas that present different characteristics and situations. These points suggest that the research on polycentric control also resonates directly with our two other patterns of adaptive failure, as it is relevant to the capacity to adapt timely to changing conditions.

More recent approaches to the management of common pool resources (arguably reactionary to the previous one) suggest that control should therefore be entirely decentralized to fully leverage the expertise of local actors. However, Ostrom's research observes the failure of both centralized and decentralized approaches (Andersson and Ostrom, 2008). In our mind, total decentralization of control runs the risk of sub-parts operating in silos, and of falling into other patterns of adaptive failure, especially those that correspond to horizontal interactions (such as in the firefighting case above).

Supporting Polycentric Control

The notion of "adaptive governance" (Ostrom, 1999; Woods and Branlat, 2010) suggests that, in order to efficiently support polycentric decision, levels of autonomy and authority ought to exist and be dynamically balanced across multiple centers pursuing interdependent goals at various echelons of the system. The basic control scheme works as follows: higher echelons of the system are needed to provide a broader perspective (time, space, functional interdependencies) in order to coordinate over emerging trends to meet changing priorities; more local layers need to be empowered to deal with the issues they are directly confronted to; both layers are then in constant interplay as situations evolve and as a result of activities and progress at each center.

Within multi-layered networked systems, the capacity becomes essential for particular roles or centers to adjust their behavior by taking into account interdependencies with other roles, activities, and events (e.g., Kulathumani *et al.* 2008). Understanding and managing cross-scale interactions (vertical or horizontal) allows avoiding undesired side-effect in tightly coupled systems. In human systems, such capacity resides essentially in the standard practices in place (which organize interdependent work based on models of the tasks) or people's expertise (who hold the knowledge of how their activities relate to others), both powerful but limited mechanisms, as we have detailed earlier. Amplifying control then consists of developing tools that help reveal/track relevant interdependencies and help anticipate how projected actions will propagate (resonate) across interdependencies relative to goals (Woods and Branlat, 2010). Hollnagel (2004) has proposed FRAM (Functional Resonance Accident Method) as an attempt to represent functional interdependencies and how variability resonates across the system, constituting sources of failure. A recent study following the last financial crisis has shown how such representations could provide explanatory value for understanding complex systemic failures (Sundström and Hollnagel, 2011), and equivalent principles would be valuable as situations unfold. In the context of computer-supported work, research on distributed decision-making further shows the impact of representations on the overall quality of decision. Studies in the domain of air traffic management have shown the value of a variety of representations to support distributed decision-making in this domain characterized by high stake, scope and interdependencies (Smith, McCoy and Layton, 1997; Smith, Spencer and Billings, 2007). Different forms of representations are needed to support the various forms of cooperative activity, from independent work to collaborative problem-solving, and support the recognition of the need or opportunity to transition between more or less costly forms of cooperation.

Conclusion

We have presented a resilience engineering perspective on how work systems, particular instances of complex adaptive systems, sometimes fail to deal successfully with the variability of their environment. We have identified three basic maladaptive patterns based on studies of real cases. These patterns suggest directions of development so that work organization, as an example of a complex adaptive system, can see and avoid or recognize and escape the corresponding failures. Study of resilient or brittle behaviors can serve as inspiration or point to limitations in order to indentify fruitful directions to support work systems in managing their adaptive capacity. More generally, the dynamics and principles identified in the study of work environments as complex adaptive systems should resonate with concepts and themes from other fields related to complexity.

References

Andersson, K., & Ostrom, E. (2008). Analyzing decentralized resource regimes from a polycentric perspective. *Policy Sciences, 41(1)*, 71-93.

Ashby, W. R. (1956). *An Introduction to Cybernetics*. London, UK: Chapman & Hall.

Branlat, M., Fern, L., Voshell, M., & Trent, S. (2009). Understanding Coordination Challenges in Urban Firefighting: A Study of Critical Incident Reports. *In Proceedings of the 52nd Annual Meeting of the Human Factors and Ergonomics Society*. *San Antonio, TX, Oct. 2009*.

Carlson, J. M., & Doyle, J. (2002). Complexity and robustness. *In Proceedings of the National Academy of Sciences of the United States of America* (Vol. 99, pp. 2538-2545).

Cook, R. I. (2006). Being bumpable: consequences of resource saturation and near-saturation for cognitive demands on ICU practitioners. In D. D. Woods & E. Hollnagel (Eds.), *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. (pp. 23–35). Boca Raton, FL: Taylor & Francis/CRC Press.

Cook, R. I., & Nemeth, C. (2006). Taking Things in One's Stride: Cognitive Features of Two Resilient Performances. In E. Hollnagel, D. D. Woods, & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 205-221). Adelshot, UK: Ashgate.

Cook, R. I., & Rasmussen, J. (2005). “Going solid”: a model of system dynamics and consequences for patient safety. *Quality and Safety in Health Care, 14(2)*, 130-134.

Csete, M. E., & Doyle, J. C. (2002). Reverse Engineering of Biological Complexity. *Science, 295(5560)*, 1664-1669.

Dekker, S. (2003). Failure to adapt or adaptations that fail: contrasting models on procedures and safety. *Applied Ergonomics, 34(3)*, 233-238.

Dekker, S. (2006). *The Field Guide to Understanding Human Error*. Adelshot, UK: Ashgate.

Dekker, S. (2008). *Just Culture: Balancing Safety and Accountability*. Adelshot, UK: Ashgate.

Feltovich, P. J., Coulson, R. L., & Spiro, R. J. (2001). Learners' (mis)understanding of important and difficult concepts: a challenge to smart machines in education. In K. D. Forbus & P. J. Feltovich (Eds.), *Smart machines in education: the coming revolution in educational technology* (pp. 349-375). Cambridge, MA: MIT Press.

Grote, G., Zala-Mezö, E., & Grommes, P. (2004). The Effects of Different Forms of Co-ordination on Coping with Workload. In R. Dietrich & T. M. Childress (Eds.), *Group Interaction in High Risk Environments* (pp. 39-54). Adelshot, UK: Ashgate.

Hardin, G. (1968). The tragedy of the commons. *Science, 162(3859),* 1243-1248.

Hollnagel, E. (2004). *Barriers and accident prevention*. Adelshot, UK: Ashgate.

Hollnagel, E. (2008). Investigation as an Impediment to Learning. In E. Hollnagel, C. P. Nemeth, & S. W. A. Dekker (Eds.), *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure* (pp. 259-268). Adelshot, UK: Ashgate.

Hollnagel, E. (2009). *The ETTO Principle: Efficiency-Thoroughness Trade-Off - Why Things That Go Right Sometimes Go Wrong*. Farnham, UK: Ashgate.

Hollnagel, E., & Woods, D. D. (2005). *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. Boca Raton, FL: Taylor & Francis/CRC Press.

Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). R*esilience Engineering: Concepts and Precepts*. Adelshot, UK: Ashgate.

Klaene, B. J., & Sanders, R. E. (2008). *Structural Firefighting: Strategies and Tactics* (2nd ed.). Sudbury, MA: Jones & Bartlett Publishers.

Klein, G. A., Feltovich, P. J., Bradshaw, J. M., & Woods, D. D. (2005). Common ground and coordination in joint activity. In W. B. Rouse & K. R. Boff (Eds.), *Organizational simulation* (pp. 139-184). Hoboken, N.J.: Wiley-Interscience.

Kulathumani, V., Sridharan, M., Ramnath, R., & Arora, A. (2008) Weave: an architecture for tailoring urban sensing applications across multiple sensor fabrics. *In Proceedings of the international workshop on mobile device and urban sensing (MODUS), April 2008*.

Lagadec, P. (2010). La Question des Plans: nécessité tactique, risques stratégiques [Plans in question: tactical necessity, strategic risks]. *Préventique-Sécurité, 110*, 25-28.

Ostrom, E. (1990). *Governing the commons: the evolution of institutions for collective action*. New York, NY: Cambridge University Press.

Ostrom, E. (1999). Coping with Tragedies of the Commons. *Annual Reviews in Political Science, 2(1)*, 493-535.

Patterson, E. S., & Woods, D. D. (2001). Shift Changes, Updates, and the On-Call Architecture in Space Shuttle Mission Control. *Computer Supported Cooperative Work (CSCW), 10(3)*, 317-346.

Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books.

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., *et al*. (2009). Early-warning signals for critical transitions. *Nature, 461(7260)*, 53-59.

Shattuck, L. G., & Woods, D. D. (2000). Communication of intent in military command and control systems. In C. McCann & R. Pigeau (Eds.), *The Human in Command: Exploring the Modern Military Experience* (pp. 279-291). New York, NY: Plenum Publishers.

Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine communication*. Cambridge, UK: Cambridge University Press.

Smith, P. J., McCoy, C. E., & Layton, C. (1997). Brittleness in the design of cooperative problem-solving systems: The effects on user performance. *IEEE Transactions on Systems, Man and Cybernetics, Part A, 27(3)*, 360-371.

Smith, P., Spencer, A., & Billings, C. (2007). Strategies for designing distributed systems: case studies in the design of an air traffic management system. *Cognition, Technology & Work, 9(1)*, 39-49.

Sundström, G. A., & Hollnagel, E. (2011). The importance of functional interdependencies in financial services systems. In E. Hollnagel, J. Pariès, D. D. Woods, & J. Wreathall (Eds.), *Resilience Engineering in Practice*. Farnham, UK: Ashgate.

Weick, K. E., & Sutcliffe, K. M. (2001). *Managing the Unexpected: Assuring High Performance in an Age of Complexity*. San Francisco, CA: Jossey-Bass.

Woods, D. D., & Branlat, M. (2010). Hollnagel's test: being 'in control' of highly interdependent multi-layered networked systems. *Cognition, Technology & Work, 12(2)*, 95-101.

Woods, D. D., & Branlat, M. (2011). Basic Patterns in How Adaptive Systems Fail. In E. Hollnagel, J. Pariès, D. D. Woods, & J. Wreathall (Eds.), *Resilience Engineering in Practice*. Farnham, UK: Ashgate.

Woods, D. D., & Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton, FL: Taylor & Francis/CRC Press.

Woods, D. D., & Shattuck, L. G. (2000). Distant Supervision–Local Action Given the Potential for Surprise. *Cognition, Technology & Work, 2(4)*, 242-245.

Woods, D. D., & Wreathall, J. (2008). Stress-strain Plot as a Basis for Addressing System Resilience. In E. Hollnagel, C. P. Nemeth, & S. W. A. Dekker (Eds.), *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure* (pp. 143-158). Adelshot, UK: Ashgate.

Zhou, T., Carlson, J., & Doyle, J. (2005). Evolutionary dynamics and highly optimized tolerance. *Journal of Theoretical Biology, 236(4)*, 438-447.