Cogn Tech Work (2012) 14:187–197
DOI 10.1007/s10111-011-0201-8

ORIGINAL ARTICLE

Complexity: learning to muddle through

John M. Flach

Received: 31 October 2011 / Accepted: 10 November 2011 / Published online: 24 December 2011

Abstract The articles in this special issue are placed in the context of the literature of general systems theory. The focus is on the complexity (or requisite variety) of complex work domains and the implications for control. Following the insights of Ashby's law of requisite variety, it is concluded that classical hierarchical or servomechanism-type control systems are inadequate as a basis for dealing with the unanticipated variability endemic to complex work domains. Alternative types of control (e.g., self-organizing systems) and alternative images of cognition are suggested as a theoretical context for modeling performance in complex work domains.

Keywords Complexity · Requisite variety · Cognitive systems engineering · Control · Distributed control

1 Introduction

The theme for this special issue is "adapting to change and uncertainty." I have had the opportunity to read all the articles and have been invited to provide a constructive overview. The articles do a very good job of grounding this problem in the particulars of specific domains, such as medicine. My goal is to complement this with a broader general systems perspective, so that the patterns seen in specific contexts can be appreciated as reflecting more general patterns associated with the dynamics of managing complexity. For those whose primary interest is specific to a domain, this may seem to be a bit esoteric. However, I believe the broader perspective may be useful in helping those concerned with safety in a specific domain to learn from other domains that are also struggling to manage complexity.

Ross Ashby's Law of Requisite Variety (1956, p. 207): "*only variety in R can force down the variety due to D; variety can destroy variety*," is my hook for connecting the dots between these articles and the larger challenge of "coping with complexity" (see Hollnagel, this issue; Rasmussen and Lind 1981):

In doing this, my hope is to contribute toward building a basic theoretical foundation for meeting the challenge of engineering more effective cognitive systems. For example, systems that can help patients to more effectively manage Type 2 diabetes (Klein and Lippa, this issue), systems that can help complex organizations like NASA to coordinate activities so that risks of space flight are minimized (Patterson and Hoffman, this issue), or systems that can improve coordination in the emergency department in ways that increase patient safety (Perry and Wears, this issue).

There are two fundamental constructs that must be appreciated in order to understand the Law of Requisite Variety and its implications for a basic understanding of cognitive systems that, in turn, might be a basis for the design of more effective systems. The first construct is "variety," which is intimately related to the constructs of "complexity," "uncertainty," and "change." The second construct is the idea of "destroy(ing) variety," which is intimately related to the constructs of "regulation," "control," and "adaptation." Thus, the challenge that unites the papers in this issue is the challenge to design systems that effectively control or manage complexity (i.e., that destroy variety). I will begin with the construct of complexity, then

J. M. Flach (✉)
Department of Psychology, Wright State University,
Dayton, OH, USA
e-mail: John.flach@wright.edu

I will consider the construct of control, and finally I will consider the mapping or fit between complexity and control.

2 Complexity

The words "complex" and "complexity" appear multiple times in each of the articles in this special issue. However, as Hollnagel (this issue) notes, despite its wide use, the term "complexity" is "notoriously difficult to define." In common usage, complexity is often used synonymously with "difficulty." For example, a difficult task might typically be described as being more complex than a simple task. This reflects the empirical fact that complexity is typically correlated with difficulty. However, I believe that it is important to keep the constructs of "complexity" and "difficulty" distinct. I will use complexity as an attribute of a problem or work domain (e.g., chess is more complex than tic tac toe) that is independent of any consideration with respect to any specific control solution. I will use difficulty as an emergent property of the fit between the capabilities of a control solution or strategy and the demands of the problem. Thus, for example, with increased capability or skill (e.g., more effective representations or strategies), chess can become easier and easier—though it is no less complex. In other words, the space of possibilities (i.e., complexity) remains the same, but with increased skill, it becomes easier to navigate that space to achieve success (i.e., the complexity becomes easier to manage).

The more formal definition of complexity that Hollnagel refers to is more convenient to use relative to my preference to focus exclusively on attributes of the problem. In this context, the term complexity refers to the number of possibilities in the problem space—the greater the number of possibilities, the greater the complexity of the problem. As Hollnagel notes, this is the sense of complexity that was at the center of early work on control (Wiener 1948) and communication (Shannon and Weaver 1963). Shannon's information statistic was developed as a way to formally index complexity (i.e., the space of possibilities). In Shannon's system, information is not a function of any specific message that was sent, but rather it is a function of the total set of messages that *might* have been sent. Using Shannon's information statistics, complexity increases in proportion to the doubling of the number of possibilities. Also, there is an explicit connection between complexity and uncertainty. That is, the greater the number of possibilities, the greater the uncertainty there will be about which of these possibilities will be realized at any particular time. Thus, coping with complexity is synonymous with coping with uncertainty! Or, in other words, a system is complex if its future is uncertain.

To dig deeper into the question of complexity or uncertainty, it is reasonable to ask, what makes a system complex or what makes the future uncertain. Figure 1 suggests two dimensions that contribute to system complexity. First, there is the dimensionality of the problem space. This reflects the number of variables, parameters, degrees of freedom, or states that contribute to shaping the field of possibilities. In general, as the dimensionality increases, the complexity will also increase. Second, there is the nature of the interdependence between the dimensions. When the interdependence is low, progress through the state space (e.g., toward achieving success) will be a simple function of the component variables (e.g., either a linear causal chain or a summation of the main effects). When the interdependence is high, progress through the state space will depend on interactions among the dimensions, such that the behavior of any specific variable might change as a function of the behavior or state of other variables. In general, as the interdependence becomes higher, the complexity or uncertainty will increase.

Note that Fig. 1 is very similar to a diagram introduced by Perrow (1984) as the interaction/coupling chart. However, in Perrow's diagram, "complex" was an anchor for the interaction axis, whereas Fig. 1 is designed so that complexity is a property of the space, with complexity increasing as one moves up and to the right in the enclosed space. I have included several points in the complexity space illustrated in Fig. 1 from Perrow's diagram to illustrate the difference between the two ways of thinking about problem spaces or work domains (e.g., Military adventures,

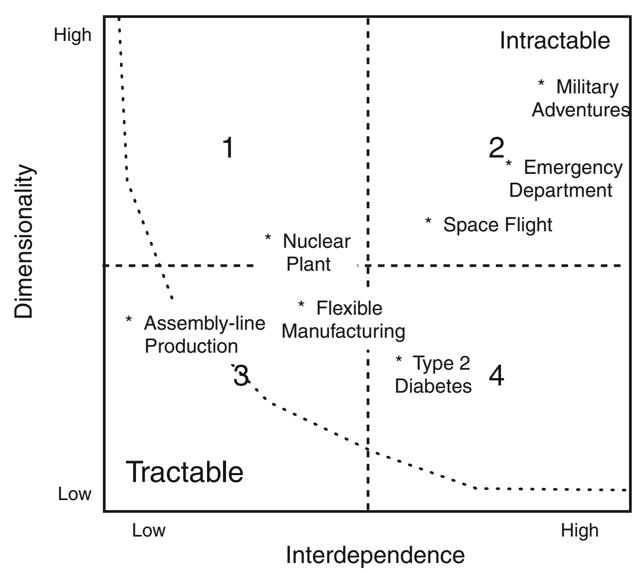

Fig. 1 This diagram illustrates problem complexity as a function of the dimensionality (e.g., number of variables or degrees of freedom) and the interdependence among the variables [from additive (linear) to interactive (nonlinear)]. Complexity is greatest in the top right corner of this space and decreases toward the *bottom left corner*

Nuclear Plant, and Assembly line production). In a following subsection on dynamics, I will have more to say about the relation between Fig. 1 and Perrow's interaction/complexity space and the coupling of behavior over time.

2.1 Dimensionality

The dimensionality of a problem reflects the number of variables that are involved in specifying the "state" of the world. For example, this could be the different variables associated with total blood cholesterol (HDL, LDL, triglycerides). The "state" of the cholesterol would be a function of the values associated with each of the three variables. The total field of possibilities would be defined by the "space" created by the joint values of the three variables. Thus, for example, the field of possibilities for cholesterol could be visualized as a three-dimensional space with each axis of the space associated with one of the three variables.

An important aspect of the dimensionality of a system is whether the problem is "closed" or "open." In a *closed system*, there is a distinct (impermeable) boundary between the problem and the environment, so that the set of variables that constitute the problem are a finite closed set. This might typically reflect the legal moves of a game. In a closed system, it is, in principle, possible to create a finite map of all the possible states, although, when the dimensionality is high (as in chess), it is practically impossible to do.

In an *open system*, there is a permeable boundary with the environment, such that changes in the environment interact with behavior of the system. For example, with respect to the larger question of cardiac disease, the cholesterol space is an open system, since additional variables (e.g., blood pressure) need to be considered. Note that the state space for cardiac disease could be expanded to include systolic and diastolic blood pressure, as well as other variables including life style variables (diet, smoking, exercise, etc.). As more variables are added, the space becomes increasingly complex.

In an open system, the dimensionality of the problem space can effectively approach infinity—since there are often an uncountable number of environmental variables that can potentially impact the problem state. In the case of cardiac disease, there may be many factors, which cannot be easily specified, that may impact overall health. For example, in the aviation domain, Rochlin (1997) notes that "pilots and other operators live in a world of irreducible uncertainty, of 'unknown unknowns,' where they must deal with what was not known not to be known" (p. 215).

To the extent that one of the major lessons of cognitive science over the last 20 years is that *context matters,* it might be argued that all cognitive systems are open systems. Thus, the dimensionality of most cognitive systems is large, if not uncountable.

As we think about the design of functional cognitive systems, one important aspect of the dimensionality of these systems (or one subset of important variables) will be the ends or goals for the work domains (e.g., the ideal levels of cholesterol). Typically, this will not be a simple function of maximizing or minimizing a single variable. Rather, most domains involve multiple goals, and as we turn to the second dimension of complexity, interdependence, we will see that these goals may involve complex interactions. In fact, rather than talking about goals as if there is a fixed, multidimensional target associated with success in a work domain, perhaps it is more accurate to talk about systems of values (e.g., safety, efficiency, personal satisfaction, etc.) that provide general indices of the health of a work domain.

In most cases, success for these domains cannot be unambiguously specified a priori. The goal posts are often moving, and in many cases, success or failure can only be discriminated in hindsight. As Weick (1995) notes, "sensemaking" is typically retrospective. In other words, we discover where we were going from where we end up. For example, in medicine, the goals of treatment may change from curing a disease to managing pain, as a disease like cancer progresses. When considering the dimensionality of a problem or work domain, it is important to think about the variables over time. In many domains, the nature of the problem can change over time as a function of the interactions with the larger context, as discussed in the following subsection.

2.2 Interdependence

Interdependence refers to the nature of the mappings from one state to another or to how each component (e.g., state, task, or agent) within the work domain contributes to progress (over time) toward achieving the functional ends for the system. If the interdependence is low, then the mappings from state to state are additive or linear. For example, a clockwork mechanism or a string of dominoes may have many possible states, yet the transition from one state to another is a relatively simple function of the component changes. Systems that are low in interdependence can often be fairly well modeled using linear analytic techniques. In systems with high degrees of interdependence, the mappings from state to state are more uncertain or are more variable. For example, in even a problem like managing Type 2 diabetes (see Klein and Lippa, this issue), there can be many paths to the same goal and the same actions may not always achieve the same ends, depending on interactions with other variables, and as suggested above, the ends themselves may change as a function of variables such as age. Thus, progress toward achieving functional goals (e.g., healthy glucose levels) depends on interactions among the component variables (e.g., age, stress, exercise, food intake, medication). Highly interdependent systems will generally be very difficult to model using most linear analytic techniques. Although nonlinear techniques are available, these generally are only tractable for problems with low dimensionality.

In his classical work on organizational theory, Thompson (1967) discusses three common forms of interdependence within organizations: pooled, sequential, and reciprocal. When the interdependence is *pooled,* the output is a simple (e.g., additive or linear) combination of effort from within the organization. That is, there is little interaction between actors (or component actions). If each actor does his job properly, then the system will reach the goal. For example, in a crew shoveling coal, the total output would be the simple sum of each worker's contribution—and failure of any individual would have only a minor consequence on the total production.

When the interdependence is *sequential*, there are sequential or precedence constraints such that some jobs cannot be done properly unless other jobs are completed properly first. For example, assembly lines production systems often involve tight sequential constraints such that tasks down the line depend on proper completion of earlier tasks. Much of the Scientific Management approach to the design of assembly lines focused on insuring the proper partitioning of work to satisfy the precedence relations—to find the single best way for each component to fit into the sequence.

When the interdependence is *reciprocal,* there are contingency constraints, such that the requirements for proper functioning of some tasks will change depending on how other tasks are performed. In such systems, there will often not be any *single best way*, but rather the proper functioning of each component will necessarily change, depending on the behavior of other components (including environmental or contextual components).

The descriptions of management of the shuttle flight (Patterson and Hoffman, this issue), of the emergency department (ED) (Perry and Wear, this issue), and even of Type 2 diabetes (Klein and Lippa, this issue) illustrate reciprocal dependences. For example, Patterson and Hoffman (this issue) describe the reciprocal coupling between the US and the Russian teams, where late decisions by the Russians required last minute changes in the docking procedures by the US (i.e., closing the vent doors prior to docking). Perry and Wear (this issue) describe how both the charts and the whiteboards are used to facilitate cooperation due to reciprocal dependencies within the emergency department (e.g., between treatment and billing, with consultant, and across shifts). They suggest how "shadow charts" have evolved to meet the needs for coordination (i.e., reciprocal adjustments) that are not satisfied by more structured electronic charts. Finally, Klein and Lippa (this issue) describe the reciprocal dependencies between stress, exercise, food intake, and medication in determining glucose levels such that "there is no single formula for glycemic control" (p. 2 of manuscript).

It is likely that all three forms of interdependence (pooled, sequential, and reciprocal) are present in most modern work domains. For example, it is likely that the whiteboards described by Perry and Wear (this issue) reflect all three forms: division of labor (pooled dependencies); precedence relations (sequential dependencies); and dynamic adjustments to real-time contingencies (reciprocal adjustments) in the ED. Thus, the interdependence dimension in Fig. 1 is not a discrete function of one form of interdependence or another, but rather it will reflect the relative dependency on these three forms of interdependence. Interdependence will be low to the extent that success depends only on the simple additive or pooled activity of components (e.g., agents or tasks) in the work domain; the more that success depends on interactions [e.g., either precedence (sequential) or contingency (reciprocal) constraints], the higher will be the degree of interdependency. That is, the degree to which the behavior of one component must be adjusted to reflect the behavior of other components, the higher will be the degree of interdependence.

2.3 Complexity: dimensionality $\times$ interdependence

In discussing complexity, Hollnagel (this issue) distinguishes between tractable and intractable systems:

> A system, or a process, is said to be tractable if the principles of functioning are known, if descriptions are simple and with few details, and most importantly if the system does not change while it is being described. Taken together these conditions mean that a system can be described both in principle and in practice. Conversely, a system or process is intractable if the principles of functioning are only partly known or even unknown, if descriptions are elaborate with many details, and if the system may change before the description is complete (Hollnagel, this issue, p. 9 of manuscript).

In Fig. 1, tractability is associated with the lower left. Tractable problems are either low dimensionality or low in interdependence. Low interdependence can be a result of either a pooled organization or a mechanistic (causal) organization. Pooled organizations are tractable as a function of the properties of large aggregates (see Weinberg 1975). In these organizations, the details of individual elements can be ignored, so that the behavior of the system can be described as a function of summing or averaging over the elements.

Alternatively, tractability can result from mechanistic or causal interdependencies among the components. The prototypical case for this is a clockwork mechanism. Here, the contingencies are fixed a priori so that there is no need for reciprocal adjustments during operations. Although the components are tightly coupled, the interdependencies are locked in—so these systems would be considered low with respect to the interdependence dimension of Fig. 1.

A key property of causal systems that makes these problems tractable is reversibility. In effect, it is possible to turn the crank forward or backward, and in either case, the behavior path will be identical, but reversed. This has major implications for theories of error (see Cacciabue, this issue; and Dekker and Nyce, this issue). In a causal system, it is in principle possible to trace back along the time history of behavior to find the *root cause* or to discover where to place the blame and who to punish. Although even in a causal system, there is the question of the appropriate stopping rule: Is the cause the proximal act that set the accident in motion or is it necessary to trace further back to *latent causes* at organizational or even societal levels?

As Fig. 1 illustrates, however, most modern work domains fall outside the region of tractability. In these work domains, the details matter, so summing or averaging over components results in the loss of critical information. Additionally, behavior in these systems is not reversible. It is not possible to turn the crank backward to discover a root cause. Tracing back along behavioral trajectories creates an illusion of order and regularity that is not justified—it leads to a hindsight bias (e.g., Fischoff 1975; Woods et al. 1994).

The tractable versus intractable distinction reflects a huge gulf between our analytical and experimental methods and the realities of complex, safety critical systems. In fact, the boundary between tractable and intractable in Fig. 1 reflects, to a large extent, the limits of Western scientific methods. As Dekker and Nyce (this issue) observe:

> Newtonian-Cartesian precepts, especially the ideas of causality and the decomposition of complex systems into constituent. elemental parts, have been long dominant. This has led us to think about human action in terms of linear sequence(s) of causes and effects and to seek “root causes” in the malfunctioning or breakage of component parts.

As Kirlik (this issue) observes, this view places serious constraints on empirical work as well. Empirical work tends to sample only a very small region in the lower left corner of the complexity space (Fig. 1). Yet, we have tended to extrapolate broadly from these experiments. In fact, the general consensus seems to be that the simpler more controlled the experiment, the more general the results. However, as Kirlik (this issue) and others (e.g., Bennett and Flach 2011) observe, for complex systems, there is no free lunch—experimental control generally comes at the cost of generality. The lesson is not to abandon controlled research, but to sample more broadly within the complexity space (e.g., utilizing synthetic environments, Flach et al. 2010; and naturalistic observation, Klein et al. 1993) and to be more cautious in our generalizations from any specific experiment.

2.4 The dynamics of complexity

As we consider the behavior of complex systems (e.g., reversibility), time is a critical dimension. However, this dimension is not well represented in Fig. 1. Systems theorists make an important distinction between stationary and nonstationary processes or systems. In a stationary system, the parameters of the process that determine the nature of the interactions among variables are invariant over time. Most of our analytical techniques for modeling systems assume that the process is stationary.

In nonstationary systems, parameters change over time. These changes may be a function of time (e.g., aging); they may be a function of prior activity (e.g., learning), or they may be a function of context (e.g., situation dynamics). Such changes over time are one of the reasons that cognitive systems are not generally reversible. You cannot undo experience—so order in time almost always matters!

Another important feature of the dynamics of complexity is the circular coupling among the variables that one typically finds. Whenever, there is a circular coupling within a network of variables, any decomposition associated with causality becomes suspect (Jagacinski and Flach 2003). When there is a circular coupling, all components (e.g., perception and action) are simultaneously cause and effect. The behavior of variables within the circular coupling is simultaneously shaping and being shaped by the behavior of all the other variables in the circle.

Perrow (1984) considers dynamic aspects of work domains in relation to the coupling axis. For example, in tightly coupled systems, delays in processing are not possible and the sequence of activities is invariant. These systems reflect clockwork types of organization. In loosely coupled systems, delays are possible and sequencing can change. This coupling in time tends to be correlated with the interdependence dimension in Fig. 1—reciprocal forms of interaction tend to be associated with loose couplings. However, in principle, the coupling over time might be best considered as an important third dimension of complexity, not well represented in Fig. 1.

3 Control

Ashby's Law of Requisite Variety is ultimately a prescription for control (i.e., for "destroying variety"). In simpler terms, the Law of Requisite Variety states that a controller must be at least as complex as the problem that is being controlled, otherwise there will be some regions of the problem space that will be unreachable—or there will be some sources of variety (you might read error here) that cannot be destroyed. This raises a very fundamental question: If the variety in complex, open systems is essentially uncountable, are these systems controllable? Is control possible in the face of irreducible uncertainty? The answer to this question will hinge on what we mean by "control."

Figure 2 provides a spatial representation for thinking about various forms of control systems that I hope will help to address the question of controllability with respect to complex systems.

The conventional view of control typically implies some local agency (e.g., homunculus or central committee) that sets the goal, creates a plan for reaching that goal, monitors behavior relative to the plan, and takes corrective action to minimize deviations from the plan (e.g., Miller et al. 1960). If this is the image of control, then the answer to the question about controllability is undoubtedly that most complex systems are NOT controllable.

There are several reasons for this, but the principal reason is associated with the capacity for any central agency to process information. It takes time to gather and integrate information, so that if the amount of information required is uncountable, then the time needed to process the information will surely exceed the limits of stability for any control system. The local agent will always be a day late and a dollar short, or in other words, the local agent will always be fighting the last war. Note that time delays are a primary threat to stability in any closed-loop network (e.g., Flach et al., in press). In the face of irreducible uncertainty, centralized, hierarchical control systems will fail as a result of "paralysis of analysis."

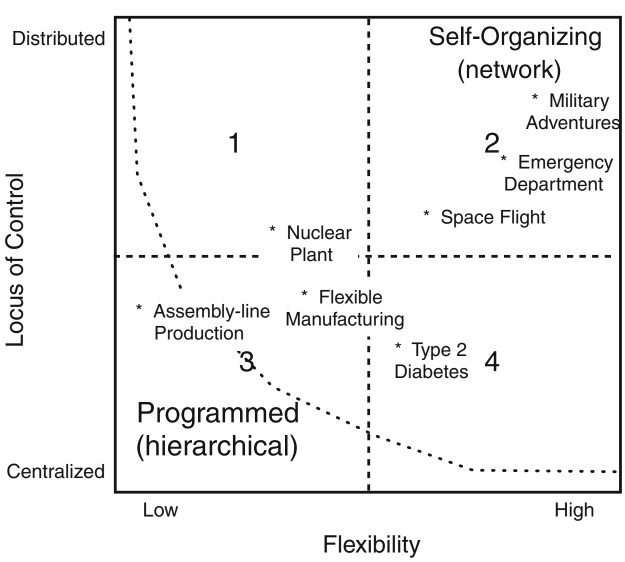

Fig. 2 This space illustrates varieties of control systems as a function of locus of control (centralized vs. distributed) and the degree of Flexibility (or discretion)

One of the first to recognize the limitations of centralized control systems with respect to complex systems was the economist Hayek (1945). He argued that centralized control of economic systems would lead to instability due to the inability to access the information needed:

> The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess (p. 519).

In other words, Hayek argues that no centralized control agency could match the variety of an economic system. However, this does not mean that economic systems could not be stabilized. Hayek believed that stability could be achieved through a more distributed form of control as reflected in market driven systems. In these systems, control is distributed across individuals, each responding based on specific, local information. While the variety of each individual agent is limited with respect to the problem complexity, there are conditions where the individuals might be coordinated so that the system will be "smarter" than any of the individuals within it. Hayek (1945) writes:

> We need decentralization because only thus can we insure that the knowledge of the particular circumstances of time and place will be promptly used. But the "man on the spot" cannot decide solely on the basis of his limited but intimate knowledge of the facts of his immediate surroundings. There still remains the problem of communicating to him such further information as he needs to fit his decisions into the whole pattern of changes of the larger economic system (pp. 524–525).

This need for "further information as he needs to fit his decisions into the whole pattern" articulates the coordination problem discussed by Patterson and Hoffman (this issue) with respect to managing space flight. The major point, relative to managing complexity, is that it is possible to match the variety needed to control complex systems, but the types of control systems that can match this variety are very different than the conventional hierarchical or servomechanism images of control.

An important distinction with respect to types of control systems that is reflected in Hayek's contrast between centralized and market economies is the distinction between *hierarchical* (or top down) versus *network* (heterarchical or bottom-up) forms of control systems. The conventional notion of control typically implies a hierarchical structure—specific goals or instructions are passed down the chain of command from a commander (central agency or brain) to the individual troops as implied by the formal structure of military systems. The troops follow the plan presented by the commander, adjusting behavior as a result of feedback relative to error from the prescriptions of the plan. In such a system, it is reasonable to say that the behavior of the troops is "caused" by the instructions from the commander. The troops have little flexibility, but to follow the plan (i.e., to do or die). An adaptation in this context simply means to correct errors relative to the plan.

A market economy, however, does not have a commander passing down instructions. It is a collection of individual agents all acting based on their local self-interests, and yet in some cases, the system as a whole can behave, as if it is moving toward some common good. A market economy illustrates a network form of organization. Each individual in this network is simultaneously constraining and being constrained by the other components of the network. Consistent with Hayek's intuition that market economies are more stable or more resilient than centralized economies, there is a growing consensus that given sufficient information, the network forms of control are generally more stable than hierarchical forms of control. For example, Arquilla and Ronfeldt (1997) write that "the information revolution favors and strengthens network forms of organizations, while making life difficult for hierarchical forms" (p. 5).

Between the poles of strict formal hierarchies and flat networks, there is a great range of flexibility in organizations. For example, a heterarchical form of organization is one that functions like a hierarchy at any specific moment, but the locus of control shifts over time, so there is no single fixed authority. This form of organization has been associated with high-reliability systems (e.g., Rochlin 1997). Rochlin (1997) observed this type of organization in aircraft carriers, where the formal military hierarchy gave way to a "flat" or heterarchical form of organization, where "even the lowest rating on the deck has not only the authority but the obligation to suspend flight operations immediately, under the proper circumstances, without first clearing it with superiors" (p. 83–84).

Also, discussions of "command intent" in relation to military leadership reflect the need for commanders to leave sufficient discretion to junior level officers so that the system can adapt to changing contingencies that could not be anticipated in a fixed formal plan. Shattuck (2000) describes the tension that military commanders experience between centralization and flexibility:

> The senior commander must make an inherent tradeoff which impacts the subordinate commander's ability to adapt to battlefield conditions. The battlefield is a highly complex, uncertain environment where a commander matches wits with his opponent while coping with such variables as terrain, weather, morale, fatigue and equipment. Providing subordinate commanders a large degree of flexibility is critical to success (p. 68).

Stability in network forms of organization is not a given. As Hayek noted, it depends on access to information. For example, Shattuck (2000) addresses this with the construct of *imparting presence*. This refers essentially to the ability to view the situation through the commander's eyes. In psychological terms, the point is that perception and action are intimately coupled. In control theoretical terms, this means that control depends critically on observation. Feedback is one example of this information coupling. But it does not reflect the more global significance of the role of observation in control systems. An important aspect of the observer problem is to identify the state variables to be fed back. In engineering systems, this observer problem is typically solved in the design of the control system. However, for biological systems, the choice of which variables to attend to is often made in parallel with the control problem. In essence, as March (1991) observed, control (exploitation) and observation (exploration) go hand in hand in managing complexity:

> A central concern of studies of adaptive processes is the relation between the exploration of new possibilities and the exploitation of old certainties … Exploration includes things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation. Exploitation includes such things as refinement, choice, production, efficiency, selection, implementation, execution. Adaptive systems that engage in exploration to the exclusion of exploitation are likely to find that they suffer the costs of experimentation without gaining many of its benefits. They exhibit too many undeveloped new ideas and too little distinctive competence. Conversely, systems that engage in exploitation to the exclusion of exploration are likely to find themselves trapped in suboptimal stable equilibria. As a result, maintaining an appropriate balance between exploration and exploitation is a primary factor in system survival and prosperity (p. 71).

Adequate resolution of the observation problem—to identify the variables that need to be fed back in order to assess progress toward an objective—is a necessary prerequisite for stable control. In a distributed control system, such as a free market, solving the observation problem will typically be distributed over the system. In this context, solving the observation problem requires communications to establish a kind of common ground for coordination among the distributed agents. Thompson (1967) described three types of coordination or control mechanisms in relation to his three types of interdependence: standardization (pooled); plans (sequential); and mutual adjustment (reciprocal).

Standardization reflects a minimal requirement for coordination—that is, there must at least be some common language for communication so that the outputs from one component can be interpreted by other components. Consistency has long been a hallmark of good design. For example, where would you find the print function for a new piece of software? Or why is the QWERTY keyboard a d*e facto* industry standard for word processing?

If there are precedence constraints, then standardization may not be sufficient. To manage the order, sequence, or timing of distributed activities, a *plan* or schedule can be useful for coordination—so that everyone is "on the same page."

Standardization and planning are coordination instruments that are typically used by central agencies to "control" an organization. However, both are forms of constraints that limit variety—and together they cannot meet the demands of many complex systems. For complex systems, such as military conflicts, it is common knowledge that "no plan survives contact with the enemy." Or in other words, the variety of conflict will always exceed the variety of any set plan. Thus, both Klein and Lippa (this issue) and Patterson and Hoffman (this issue) include "re-planning" as an essential feature of macrocognition.

Thompson's (1967) third means of coordination, *mutual adjustment*, is the capability that distinguishes free markets from centralized economic control systems. Free markets achieve stability as a result of individuals mutually adjusting to the evolving market situation. However, as noted above, this requires some degree of shared information or common ground, so that each individual can have some sense of how they "fit" into the global process. In some cases, standardization and plans (even if tentative) can be important contributions to creating that common ground.

As with the earlier discussion of complexity, Thompson's three categories of coordination are typically not associated with discrete forms of natural control systems. Rather, most natural control systems will reflect a combination of all three forms of coordination.

The key point to make about control is that in dealing with complex systems, it becomes necessary to reconsider our image of what it means to control a system. This vision has to appreciate the intimate relation between control (exploitation, action) and observation (exploration, perception), and it has to appreciate that control will typically not be localized. Control does not necessarily mean to guide a system along a clearly defined path to a clearly defined goal. More typically, control in complex systems will mean finding a stable balance among multiple competing interests. In these cases, the stability point will rarely be a predefined set point, rather it will be an equilibrium point that emerges dynamically as a consequence of learning and experience.

4 Controlling complex systems

In general, as the complexity of the natural control problem increases, one can expect that the there will be a need for increased reliance on mutual adjustment to meet the demands of the Law of Requisite Variety. For example, for complex problems, replanning will be essential. For complex problems, there will be increased demands for distributed, flexible forms of control. Thus, there is a parallel construction in Figs. 1 and 2 such that the demands created by complexity in Fig. 1 are best met by the control capabilities in the corresponding regions of Fig. 2. For example, the ED discussed by Perry and Wear (this issue) would be the type of system that would be vulnerable to instability without sufficient capability for mutual adjustment. Thus, one would expect that high functioning (resilient) EDs would exhibit flexible, distributed styles of control. In this context, the white board and the charts are artifacts that help serve the observation or information requirements of stability for that style of control.

Today, systems that rely heavily on mutual adjustment to achieve stability are often called *self-organizing systems*. Perhaps, the most powerful image for self-organization is the theory of evolution. In this control solution, species achieve increasingly functional levels of organization without any explicit goal or plan to guide them. Order emerges or self-organizes as a result of the interactions between variation at the micro-level (e.g., genetic variations) and variation at the macro-level (e.g., selective survival) each simultaneously functioning as both cause and effect.

Note that evolution is sometimes discussed as a generator of variety since Darwin reported on the different varieties within a species that reflected local adaptations. However, we now understand that the variation of species is low, relative to variation at the genetic level. The species selected as a result of functional fitness with ecological constraints reflects only a subset of the possible variations at the genetic level. In effect, natural selection filters variation at the genetic level, selecting only those variations that "fit" the functional demands of the ecology, thus effectively destroying variety. Other examples of self-organizing systems can be found in the literature on Artificial Life—which includes many examples of systems were functional order emerges, without any specific goal, plan, instruction set, or program (e.g., Langton et al. 1992; Kugler and Turvey 1987).

In many respects, self-organizing systems will be a better metaphor for guiding intuitions about the control of complexity, than the conventional images (e.g., the hierarchy or the servomechanism). This has important implications for how we think about the control of complexity. First, it is important to realize that there is no single best way, nor is there a fixed target. Goals emerge and change over time. They are more typically a retrospective product of sensemaking, than an a priori guide to action (e.g., Weick 1995). For example, controlling glucose levels (e.g., Klein and Lippa, this issue) for people with Type 2 diabetes cannot be achieved by following a simple recipe. Both the means and the ends of stable glucose levels must be discovered through trial and error. In this case, sensemaking involves a retrospective analysis of patterns over time to learn about the complex interactions involved and to keep pace with changes associated with age and stress.

Despite the desire for and logical appeal of more normative models for regulating diabetes, the information demands and complexity of the interactions make it difficult or impossible to realize such a solution. Thus, the control dynamic involved in managing diabetes suggests a trail and error process akin to Lindblom's (1959) "muddling through" or "incremental adjustment." This involves a gradual tweaking of experience to balance between the stability of things that have worked in the past and the need to adapt to changing circumstances. Lindblom (1979) observes that "many critics of incrementalism believe that doing better usually means turning away from incrementalism. Incrementalists believe that for complex problem solving it usually means practicing incrementalism more skillfully and turning away from it only rarely" (p. 517).

I wonder whether the evolving constructs of situation awareness (Banbury and Tremblay 2004) and macrocognition (Klein and Lippa, this issue; Patterson and Hoffman, this issue, Patterson and Miller 2010; Schraagen et al. 2008) may, at least in part, reflect a growing realization among researchers that cognition is less about conformity to normative models of decision and choice (that could never be realized in a complex problem space), than it is about skilled incrementalism, that is, the skill of balancing between reliance on experience (i.e., generalizing from solutions that have worked in the past) and openness to learn from the errors and surprises that are inevitable in a changing environment. It is about finding the right balance between exploitation and exploration (March 1991). It involves an ecological rationality (e.g., Todd and Gigerenzer 2003), where heuristics are grounded in experience, rather than in context-free logic. It demands an abductive form of rationality (e.g., Bennett and Flach 2011; Flach 2009) where hypotheses guide actions, while they are simultaneously tested through those same actions.

A second important implication of the self-organizing system metaphor is that the concept of error becomes problematic. Classically, error has been viewed as a deviation from either a specific target/plan or as deviation from a more general normative standard (e.g., prescriptions of logic or economic models of rationality). These standards are inconsistent with the self-organizing system metaphor. There is no static target or plan, and the context independent standards of classical logic no longer apply.

Cacciabue and Cassani (this issue) recognize this concern with the addition of a motivational component to their model of error. However, the model still assumes a causal dynamic that is not consistent with the self-organizing system metaphor. This does not mean that such models cannot be extremely useful for developing insight into the complex interactions of many important variables associated with error and performance. But it does mean that one should be very skeptical about any specific predictions based on this model. The model should be viewed as a tool for exploring the complexity, but not for explaining it.

Dekker (2011) and others (e.g., Woods et al. 1994) take a more radical approach, questioning the very construct of human error as an explanation for system failure. Dekker (2011) writes that in complex systems "the harmful outcome is not reducible to the acts or decisions by individuals in the system, but a routine by-product of the characteristics of the complex system itself" (p. 14). Yet, as Dekker and Nyce (this issue) observe, the desire for causal explanations and the desire to attribute blame to some individual are ingrained in Western culture.

5 Summary and conclusion

The fundamental question for this special issue is "coping with complexity" or "coping with uncertainty." Hollnagel (this issue) provides an historical context for the evolving appreciation of this problem. Where early approaches, such as Scientific Management, attempted to reduce the complexity to conform to a single best or optimal way, today there is a growing appreciation that in domains such as medicine, process control, flexible manufacturing, space exploration, and military operations, we are often facing irreducible complexity. In these domains, it is no longer possible to eliminate the complexity, rather the challenge is whether we can provide the tools so that the people in these systems are better able to cope with the complexity in real time.

Klein and Lippa (this issue), Patterson and Hoffman (this issue), and Perry and Wear (this issue) provide illustrative examples of coping with complexity. Klein and Lippa (this issue) and Patterson and Hoffman (this issue) suggest that understanding human performance in these systems requires a reconceptualization of cognition to consider more “macro” aspects of coping. This includes constructs such as replanning, problem finding, sensemaking, and coordination.

Perry and Wear (this issue) describe some of the information artifacts that provide the information coupling necessary to support coordination in a distributed control system. In particular, they illustrate how human experts creatively respond to fill information gaps that can be an unintended consequence of automated solutions.

Cacciabue and Cassani (this issue) suggest that in the face of complexity, it becomes necessary to consider the intentional or “motivational” dynamics that shape human choices. While Dekker and Nyce (this issue) explore why, in the face of this complexity, analysts continue to cling to causal explanations that tend to trivialize the dynamics of complex systems.

Finally, Kirlik (this issue) considers the implications for experimental programs of research. How do we pursue experimental programs of research that can generalize to the complex worlds of safety critical systems? Following Brunswik’s (1956) lead, he recommends that researchers attend more to the ecological constraints that shape performance, to insure that the experimental contexts are more representative of the domains to which they hope to generalize. If we design our experimental programs around narrow images of control (e.g., with simple, clearly specified goals), then any generalizations to complex work domains will be suspect.

In conclusion, I would like to shift focus from complexity to “coping.” How is it possible to “control” in the face of irreducible complexity? Despite increasing awareness of the power of networks and distributed processing, much of the research on human performance is still guided by a fairly conventional image of control. It is important to realize that hierarchical, centralized control systems, in principle, cannot cope with complexity. Thus, it becomes necessary to expand our vision of control to include more network forms of organizations. We need to shift focus from goal following behavior to stability around emerging targets of opportunity. We need to shift focus from error, to consider resilience. In the face of complexity, we need to think seriously in terms of a “science of muddling through” (Lindblom 1959). The goal of design needs to shift from protecting the system from human variability to amplifying human variability in ways that destroy situation variability. That is, we need to support the creative capacities of the humans so that they can invent solutions in real time to problems that could not have been anticipated in advance. In the end, the most reliable antidote to unresolvable uncertainty may be human creativity. In other words, only human variability can destroy the variability (i.e., uncertainty) that is endemic to complex work domains.

References

Arquilla J, Ronfeldt D (1997) In Athena’s camp: preparing for conflict in the information age. Rand, Santa Monica

Ashby R (1956) An introduction to cybernetics. Chapman & Hall, London

Banbury S, Tremblay S (2004) A cognitive approach to situation awareness: theory and application. Ashgate, Aldershot

Bennett K, Flach J (2011) Display and interface design: subtle science, exact art. CRC Press, Boca Raton

Brunswik E (1956) Perception and the representative design of psychological experiments. University of California Press, Berkeley

Dekker S (2011) Drift into failure: from hunting broken components to understanding complex systems. Ashgate, Aldershot

Fischoff B (1975) Hindsight-foresight: the effect of outcome knowledge on judgment under uncertainty. J Exp Psychol Hum Percept Perform 1:288–299

Flach J (2009) The dynamics of experience: a search for what matters. Proceedings of the European conference on cognitive ergonomics (ECCE 2009). Helsinki, Finland, pp 11–18

Flach JM, Bennett KB, Jagacinski RJ, Mulder M, van Paassen MM. The closed-loop dynamics of cognitive work. In: Lee JD, Kirlik A (eds) The Oxford handbook of cognitive engineering (In press)

Flach JM, Schwartz D, Bennett A, Behymer K, Shebilsi W (2010) Synthetic task environments: measuring macrocognition. In: Patterson E, Miller J (eds) Macrocognition: metrics and scenarios: design and evaluation for real world teams. Ashgate, Aldershot, pp 201–284

Hayek F (1945) The use of knowledge in society. Am Econ Rev 35(4):519–530

Jagacinski R, Flach J (2003) Control theory for humans: quantitative approaches to modeling performance. Erlbaum, Mahwah

Klein G, Orasanu J, Calderwood R, Zsambok C (1993) Decision making in action: models and methods. Ablex, Norwood

Kugler P, Turvey M (1987) Information, natural law, and the self-assembly of rhythmic movement. Erlbaum, Hillsdale

Langton C, Taylor C, Farmer J, Rasmussen S (1992) Artificial life II. Addison-Wesley, Redwood

Lindblom (1959) The science of “muddling through”. Public Admin Rev 19(2):79–88

Lindblom C (1979) Still muddling, not yet through. Public Admin Rev 39(6):517–526

March J (1991) Exploration and exploitation in organizational learning. Organ Sci 2(1):71–87

Miller G, Galanter E, Pribram K (1960) Plans and the structure of behavior. Henry Holt & Company, New York

Patterson E, Miller J (2010) Macrocognition metrics and scenarios: design and evaluation for real-world teams. Ashgate Publishing, Aldershot

Perrow C (1984) Normal accidents. Living with high-risk technologies. Basic Books, New York

Rasmussen J, Lind M (1981) Coping with complexity (Risø-M-2293). Risø National Laboratory, Roskilde
Rochlin G (1997) Trapped in the net. The unanticipated consequences of computerization. Princeton University Press, Princeton
Schraagen J, Militello L, Ormerod T, Lipshitz R (2008) Naturalistic decision making and macrocognition. Ashgate, Aldershot
Shannon C, Weaver W (1963) The mathematical theory of communication. University of Illinois Press, Urbana (original publication 1949)
Shattuck L (2000) Communicating intent and imparting presence. Military Rev, pp 66–72
Thompson J (1967) Organizations in action. McGraw-Hill, New York
Todd P, Gigerenzer G (2003) Bounding rationality to the world. J Econ Psychol 24:143–165
Weick K (1995) Sensemaking in organizations. Sage Publications, Thousand Oaks
Weinberg GM (1975) An introduction to general systems thinking. Wiley, New York
Wiener N (1948) Cybernetics: or control and communication in the animal and the machine. MIT Press, Cambridge
Woods D, Johannesen L, Cook R, Sarter N (1994) Behind human error: cognitive systems, computers, and hindsight. CSERIAC, Wright-Patterson AFB
Cacciabue P, Cassani M. Modeling motivations, tasks and human errors in a risk based perspective (this issue)
Dekker S, Nyce J. Cognitive engineering and the moral theology and witchcraft of cause (this issue)
Hollnagel E. Coping with complexity: past, present, and future (this issue)
Kirlik A. Relevance versus generalization in cognitive engineering (this issue)
Klein H, Lippa K. Assuming control after system failure: Type II diabetes self-management (this issue)
Patterson E, Hoffman R. Visualization framework of macrocognition functions (this issue)
Perry S, Wears R. Underground adaptations: case studies from healthcare (this issue)