Reliability Engineering and System Safety 141 (2015) 5–9

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Four concepts for resilience and the implications for the future of resilience engineering

David D. Woods

Initiative on Complexity in Natural, Social & Engineered Systems, The Ohio State University, United States

ARTICLE INFO

Available online 3 April 2015

Keywords:
Resilience engineering
Resilience
Sustainability
Resilient control
Robust control
Complexity
Complex adaptive systems
Socio-technical systems

ABSTRACT

The concept of system resilience is important and popular—in fact, hyper-popular over the last few years. Clarifying the technical meanings and foundations of the concept of resilience would appear to be necessary. Proposals for defining resilience are flourishing as well. This paper organizes the different technical approaches to the question of what is resilience and how to engineer it in complex adaptive systems. This paper groups the different uses of the label 'resilience' around four basic concepts: (1) resilience as rebound from trauma and return to equilibrium; (2) resilience as a synonym for robustness; (3) resilience as the opposite of brittleness, i.e., as graceful extensibility when surprise challenges boundaries; (4) resilience as network architectures that can sustain the ability to adapt to future surprises as conditions evolve.

1. Introduction

Today's systems exist in an extensive network of interdependencies as a result of opportunities afforded by new technology and by increasing pressures to become faster, better and cheaper for various stakeholders. But the effects of operating in interdependent networks has also created unanticipated side effects and sudden dramatic failures [42,1]. These unintended consequences have led many different people from different areas of inquiry to note that some systems appear to be more resilient than others. This idea that systems have a property called 'resilience' has emerged and grown extremely popular in the last decade (for example, articles in scientific journals on the topic of resilience increased by an order of magnitude between 2000 and 2013 based on search of Web of Science, e.g., Longstaff et al. [26]). The idea arose from multiple sources and has been examined from multiple disciplinary perspectives including: systems safety (see Hollnagel et al. (2006)), complexity (see [1]), human organizations (see [42,40,22,32,31]), ecology (see [41]), and others. However, with popularity has come confusion as the label continues to be used in multiple and diverse ways.

As multiple observers from different disciplines began to study the characteristics that affect the ability to create, manage, and sustain resilience, four core concepts appear and recur. This paper organizes the diverse uses of the label 'resilience' into groups based on these four conceptual perspectives. The paper refers to these four concepts as resilience [1] through [4]. First, people use the label resilience to refer to how a system *rebounds* from disrupting or traumatic events and returns to previous or normal activities (rebound = resilience [1]).

Second, people use the label resilience as the equivalent to the concept of system *robustness*. These two concepts have recurred repeatedly in work on resilience, especially in the early stages of exploring how systems manage complexity as they appear to provide a path to generate explanations of how some systems are able to manage increasing complexity, stressors, and challenges (robustness = resilience [2]).

As researchers have continued to study the problem of complexity and how systems adapt to manage complexity, two additional concepts have emerged. Upon further inquiry, the empirical results begin to reveal how some systems overcome the risk of *brittleness*, i.e., the risk of a sudden failure when events push the system up to and beyond its boundaries for handling changing disturbances and variations [7,43,44]. From the perspective of overcoming the risk of brittleness, a third use of the label resilience becomes the idea of *graceful extensibility* [47,45] — how a system extends performance, or brings extra adaptive capacity to bear, when surprise events challenge its boundaries (graceful extensibility = resilience [3]).

Another line of inquiry has pursued formal models of systems that have proved to be evolvable in biology and technology (e.g., the internet). A fourth use of the label resilience emerged from this work that focuses on the question: what are the architectural properties of layered networks that produce *sustained adaptability* —the ability to adapt to future surprises as conditions continue to evolve? [14,32,31]. This line of work centers on how networks can manage fundamental trade-offs that constrain all systems

---

E-mail address: woods.2@osu.edu

http://dx.doi.org/10.1016/j.ress.2015.03.018
0951-8320/© 2015 Elsevier Ltd. All rights reserved.

[9,13,5,18]. It seeks to identify governance policies that operate across layered networks in biological systems, social systems, and technological systems—what governance policies sustain the ability of the network to continue to function well and avoid falling into traps in the trade spaces as conditions change over long time scales (sustained adaptability=resilience [4]).

This paper briefly considers each of the four, in turn, to explore how each has stimulated lines of inquiry and led to new and sometimes unexpected results. The intent of the paper is to set a new baseline for future work. Whatever the historical contributions of each of these four concepts, the question is how to advance productive lines of inquiry. Organizing the numerous and continuing attempts to define resilience around these four concepts blocks out a great deal of noise (see the overview in [27]). The review of the four concepts sets the stage to debate which concepts have the potential to continue to advance our understanding of complex adaptive systems.

2. Four concepts for resilience

2.1. Resilience as rebound (or resilience [1])

The rebound concept begins with the question: why do some communities, groups, or individuals recover from traumatic disrupting events or repeated stressors better than others to resume previous normal functioning? A representative example of this approach is a recent compilation of papers assembled when an organization asked the Institute of Medicine to help it answer the above question [6]. We also find this question asked by business continuity centers as organizations confront extreme weather events that can produce surprising cascades of effects [11].

This use of the label resilience as [1] – *rebound* – is common, but pursuing what produces better rebound merely serves to restate the question. Where progress has been made, the focus is not on the period of rebound but on what capabilities and resources were present before the rebound period. Finkel's analysis of contrasting cases of recovery from or inability to recover from surprise provides compelling evidence [16]. First, it is not what happens after a surprise that affects ability to recover; it is what capacities are present *before* the surprise that can be deployed or mobilized to deal with the surprise. This issue was noted early on by Lagadec with respect to major external trigger events [20, p. 54]: “the ability to deal with a crisis situation is largely dependent on the structures that have been developed before chaos arrives. The event can in some ways be considered a brutal and abrupt audit: at a moment's notice, everything that was left unprepared becomes a complex problem, and every weakness comes rushing to the forefront”.

Second, rebound considers responses to specific disruptions, but much more importantly the disrupting events represent *surprises*, that is, the event is a surprise when it falls outside the scope of variations and disturbances that the system in question is capable of handling [43,46]. In other words, the key is not simply the attributes of the event in itself as a disruption or its frequency of occurrence, but how the event challenges a model instantiated in the base capabilities of that system. The surprise event challenges the model and triggers learning and model revision—a kind of *model surprise* [48]. There are patterns to surprise, or, as Nemeth puts it, there are regularities to what on the surface appears to be irregular variations in terms of how disturbances challenge normal functioning [30].

These two points highlight a paradox about resilience, that shifts the focus from resilience [1] to resilience [3] (graceful extensibility) as research begins to consider resilience as multiple forms of adaptive capacity. To overcome the risk of brittleness in the face of surprising disruptions requires a system with the potential for adaptive action *in the future* when information varies, conditions change, or when new kinds of events occur, any of which challenge the viability of previous adaptations, models, plans, or assumptions. However, the data to measure resilience as this potential comes from observing/analyzing how the system has adapted to disrupting events and changes *in the past* [44].

There are other limits to the line of inquiry based on resilience [1], for example, the concept of recovery to normal or previous function (return to equilibrium) has not held up to inquiry (see for example, [41]). The process of adapting to disruptions, challenges and surprises over time changes the system in question in multiple ways. In adapting to new challenges, systems draw on their past but become something new. Even when adapting to preserve, the process of adapting transforms both the system and its environment. Continuity occurs over a lineage of challenge and adaptive response, a series of adaptive cycles that compose an adaptive history.

It is historically interesting that questions about resilience are often formulated around finding a way to explain variations in how systems rebound from challenge. But research progress has left this framing behind to focus on the fundamental properties of networks, systems and organizations that are able to build, modify and sustain the right kinds of adaptive capacities [14]. Studies of biological systems [17] and evolutionary computational modeling of biological systems [23,24] have shown that properties that will sustain adaptive capacity in the future can be selected for [4]. These are examples of results that shift in focus the focus from resilience [1] to resilience [4] —architectures for sustained adaptability.

2.2. Resilience as robustness (or resilience [2])

Resilience [2] – *increased ability to absorb perturbations* – confounds the labels robustness and resilience. Some of the earliest explorations of resilience confounded these two labels, and this confound continues to add noise to work on resilience (as noted in [43,29]).

An increase in robustness expands the set of disturbances the system can respond to effectively. This simple definition is the basis for the success in robust control as a subset of control engineering [15]. “Robust control is risk-sensitive, optimizing worst case (rather than average or risk-neutral) performance to a variety of disturbances and perturbations” ([14, p. 15624]). Alderson and Doyle [1] point out that robustness is always of the form: system X has property Y that is robust in sense Z to perturbation W. In other words, robust control works, and only works, for cases where the disturbances are well-modeled.

If an increase in robustness expands the set of disturbances the system can respond to effectively, the question remains what happens if the system is challenged by an event outside of the current set? If the system cannot continue to respond to demands and meet some of its goals to some degree, then the system will experience a sudden failure or collapse – that is, the system is brittle at its boundaries—resilience [3]. In other words, resilience comes to the fore when the set disturbances is *not well modeled* and when this set is changing. And ironically, the set of poorly modeled variations and disturbances changes based on a record of past success which triggers adaptive responses by other nearby units in the layered network of interdependent systems. As a result of this fundamental result, and in a direct analogy to robust control, a new line of inquiry has emerged to develop resilient control systems for applications such as cybersecurity and cyber-physical systems (e.g., [36]).

Confounding resilience and robustness turns out to be erroneous in another way. If an increase in robustness expands the set