

Contrasting Views of Complexity and Their Implications For Network-Centric Infrastructures

David L. Alderson, *Member, IEEE*, and John C. Doyle

Abstract—There exists a widely recognized need to better understand and manage complex “systems of systems,” ranging from biology, ecology, and medicine to network-centric technologies. This is motivating the search for universal laws of highly evolved systems and driving demand for new mathematics and methods that are consistent, integrative, and predictive. However, the theoretical frameworks available today are not merely fragmented but sometimes contradictory and incompatible. We argue that complexity arises in highly evolved biological and technological systems primarily to provide mechanisms to create robustness. However, this complexity itself can be a source of new fragility, leading to “robust yet fragile” tradeoffs in system design. We focus on the role of robustness and architecture in networked infrastructures, and we highlight recent advances in the theory of distributed control driven by network technologies. This view of complexity in highly organized technological and biological systems is fundamentally different from the dominant perspective in the mainstream sciences, which downplays function, constraints, and tradeoffs, and tends to minimize the role of organization and design.

Index Terms—Architecture, complexity theory, networks, optimal control, optimization methods, protocols.

I. INTRODUCTION

MANY cyber-technical visions convincingly suggest that net-centric technology will provide unprecedented levels of capability and efficiency to support the operation and management of modern society’s most vital functions—ranging from delivery of economic goods and services, business processes, global financial markets, education, health care, defense, and other government services [1]. A fundamental challenge is to understand and manage the growing complexity of these systems [2].

Historically, we have done a poor job in managing the fragilities created by our complex networks, from global warming to ecosystem destruction, global financial crises, etc. In many cases, past failures are due to fragilities that were direct side effects of mechanisms that promised to provide great benefits, including robustness. Thus, “robust yet fragile” (RYF) is an increasingly central issue in network-centric infrastructures and

other complex engineering systems, but much of advanced technology has, if anything, made things worse. Computer-based simulation and rapid prototyping tools are now broadly available and powerful enough that it is relatively easy to demonstrate almost anything, provided that conditions are made sufficiently idealized. We are much better at designing, mass-producing, and deploying network-enabled devices than we are at being able to predict or control their collective behavior once deployed in the real world. The result is that, when things fail, they often do so cryptically and catastrophically.

The growing need to understand and manage complex systems of systems, ranging from biology to technology, is creating demand for new mathematics and methods that are consistent and integrative. Yet, there exist fundamental incompatibilities in available theories for addressing this challenge. Various “new sciences” of “complexity” and “networks” dominate the mainstream sciences [3] but are at best disconnected from medicine, mathematics, and engineering. Computing, communication, and control theories and technologies flourish but remain largely hidden and isolated, even from each other. At the risk of oversimplifying a complex subject, our aim here is a minimal but universal taxonomy that cuts across all of these areas, connecting back to now-classical visions of “organized complexity” and cybernetics, and, today, to the latest insights from biology and technology. A theme in this paper is the central but poorly formalized role that network architecture plays as “constraints that deconstrain” [4] to organize functional complexity and robustness in highly evolvable systems. The compatibility regarding these topics within and between engineering, mathematics, biology, and medicine is striking and suggestive, and there has been substantial progress bringing greater rigor and relevance to this broad subject, although much remains to be done.

We will focus on comparing the architecture of advanced technologies to biological systems for several reasons. Biological metaphors are ubiquitous (e.g., the notion of a “cyber-physical ecosystem”), and so-called *bioinspired systems* are increasingly popular topics of research and development. Yet, biological systems themselves contain bewildering complexity, often making accessible only the most superficial connections to technology, thus compounding errors and confusion. As we will argue, the connection between advanced technology and biology is neither superficial nor accidental, and much can be learned from comparing and contrasting the organizational principles underlying complex biological and technological systems. With molecular biology’s description of components and the growing attention to systems biology, the organizational principles of biological networks are becoming increasingly

Manuscript received July 1, 2008; revised March 10, 2009. Date of publication June 3, 2010; date of current version June 16, 2010. This paper was presented in part at the 2007 IEEE International Conference on Systems, Man and Cybernetics, Montreal, QC, Canada, October 7–10, 2007. This work was supported by the Office of Naval Research under MURI Awards 0001408WR20242 and N000140810747. This paper was recommended by Associate Editor M. Ulieru.

D. L. Alderson is with the Operations Research Department, Naval Postgraduate School, Monterey, CA 93943 USA (e-mail: dlalders@nps.edu).

J. C. Doyle is with the California Institute of Technology, Pasadena, CA 91125 USA (e-mail: doyle@cds.caltech.edu).

Digital Object Identifier 10.1109/TSMCA.2010.2048027

apparent. Biologists are articulating richly detailed explanations of biological complexity, robustness, and evolvability that point to universal principles [4].

On the other side, the complexity of advanced technology is indeed now approaching that of biology. While the components differ, there is striking convergence at the network level of architecture and the role of layering, protocols, and feedback control in structuring complex multiscale modularity. New theories of the Internet and related networking technologies have led to the testing and deployment of new protocols for high-performance networking [5]. These are part of a new mathematical framework for the study of complex networks [6] that suggests that this apparent network-level evolutionary convergence both within and between biology and technology is not accidental but follows necessarily from the universal system requirements to be efficient, adaptive, evolvable, and robust to perturbations in their environment and component parts.

Various domains of systems engineering and applied mathematics, including computational complexity, numerical analysis, control theory, operations research, etc., have their own different but surprisingly compatible notions of architecture, robustness, and complexity despite enormous differences in components and terminology. Computer science (CS) theorists have extensively formalized and explored computational complexity, and concepts such as NP hardness and undecidability have become a universal language, and one we will take for granted. Less pervasive is the similar formalization that control theorists have for robustness to component and environment uncertainty. Particularly surprising is the conceptual consistency across domains [7]–[9], including biology [4], about issues that are “architectural,” despite a general lack of formalization, even within engineering. All of these provide an essential foundation, and our aim is not for novelty but for consistency with them and to broaden the concepts and provide context and connections rather than deepen them. While we do not assume the reader is versed in all such domain details, we do hope that it resonates with experts.

II. ROBUSTNESS, DESIGN, AND COMPLEXITY

A key driver of network complexity is the desire for *robustness*, which is defined as follows.

Definition: A [property] of a [system] is **robust** if it is [invariant] with respect to a [set of perturbations].

The use of square brackets emphasizes that formal treatment of robustness requires the specification of the system, the property, the set of perturbations, and some measure of invariance (e.g., relative to some norm). Fragility is the opposite of robustness (i.e., the lack of invariance) and must similarly be specified. The [property] usually involves some notion of system *function* and/or *purpose* and thus depends on domain-specific aspects of the intended behavior, not simply the structure, of the system of interest. Thus, without specifying properties and perturbations, to say that a system as whole is robust or fragile can only mean that the tradeoffs are handled well or poorly, respectively.

Many important properties of systems can be viewed as a specific type of robustness. *Reliability* involves robustness to component failures. *Efficiency* is robustness to resource scarcity. *Scalability* is robustness to changes to the size and complexity of the system as a whole. *Modularity* is robustness to structured component rearrangements. *Evolvability* is robustness of lineages to changes on long time scales.

A system can have a property that is robust to one set of perturbations and yet *fragile* for [a different property] and/or [perturbation]. Highly evolved systems, ranging from advanced technologies to biology, exhibit RYF features, and understanding RYF tradeoffs lies at the heart of the design challenges for network-centric infrastructures. An example of a possible RYF tradeoff is that a system with high efficiency (i.e., using minimal system resources) might be unreliable (i.e., fragile to component failure) or hard to evolve.

This discussion of robustness assumes that “system,” “property,” and “perturbation,” are well defined, and most engineering disciplines actually do share compatible, if not identical, notions of these terms, so there is hopefully little risk in relying on this shared understanding. Much more problematic is the meaning of “complexity” that is relevant for describing biological and technological networks. There does not exist a universally accepted definition of complexity, and considerable confusion within the broader scientific community can be traced to various usages that are not merely different but opposite. We discuss these opposing notions in Section IV and, for now, simply highlight three dimensions of complexity that are relevant to challenges in system design: large and/or diverse number of components, the complexity of their interconnections and interactions, and the complexity of the behaviors that result. In our view, however, complexity is most succinctly discussed in terms of functionality and its robustness. Specifically, we argue that complexity in highly organized systems arises primarily from design strategies intended to create robustness. We proceed, in the remainder of this section, with arguments to support this claim.

A. Organization and Constraints

Complex biological and technological systems share a feature we call *highly organized*, in that their components are arranged in a very specialized structure that enables their functionality and/or robustness features. Indeed, a feature of highly organized systems is that even minimal random rearrangement of that structure tends to destroy its most salient features. Then, a key observation in the understanding of highly organized systems is that their structure is a consequence of specific *constraints* that are placed on their functionality and/or behavior, and this is largely independent of the process by which this organization arises, whether by design or evolution. We will describe four kinds of constraints: component, system/environment, protocol, and emergent.

Component-Level Constraints. The components that comprise any system are typically constrained in terms of what they can do, even separately. For example, much of mechanics, electrical circuits, chemical processes, etc., can be described in terms of relationships such as $F = MA$ and $V = IR$. These

constraints are often expressed as differential or algebraic equations, but they can be more general. The uncertainty of components often imposes constraints on a complex system that are as important as the nominal idealized component behavior.

System-Level Constraints. What distinguishes biology and technology from other types of systems is that there are complex constraints on the system as a whole that are not consequences of those on the components, including functional requirements (i.e., what the system needs to do), as well as the environmental and operating requirements (e.g., the conditions under which the system must be able to achieve this function), and particularly robustness to uncertainty and perturbations from the environment.

Protocols. Protocols, typically in the form of rules for the configuration and/or interaction of system components, may impose additional constraints on the overall system. Although these additional constraints may reduce the number of possible system solutions, a “good” set of protocols minimally constrains these solutions so as to facilitate a focus on the feasible and robust solutions.

Emergent Constraints. The interaction of the aforementioned constraints can imply an additional set. We will (with some trepidation) refer to as “emergent” those constraints that are nontrivial consequences of the interaction between the system and component-level constraints, and possibly protocols. Perhaps the most important emergent property of any set of constraints is whether their intersection is (non)empty, so theory and methods to determine this are central to engineering specification and design.

Emergence is also associated with *unintended consequences* for either good (an emergent benefit) or bad (an emergent fragility). Emergent benefits in organized systems are sufficiently rare as to be peripheral to this paper, much wishful thinking notwithstanding, and claims regarding them usually result from a poor understanding of the other constraints. In contrast, emergent fragilities, whether unintended or the result of hard tradeoffs, are dominant problems in complex systems. Emergent fragilities are also particularly celebrated within the scientific literature on “disorganized complexity,” where the (often surprising) collective behavior arising from largely random interconnection of typically homogeneous components is called “emergent” (our focus in Section IV).

B. Dimensions of Design and Organization

Our view is that “organization” is essentially the specialized structure that allows a system to satisfy the aforementioned constraints. This notion applies to a wide variety of systems, and we will use “design” very broadly as both the resulting system organization and the process leading to it. While greater formalization is needed, most existing designs result from a mix of the *ad hoc*, trial and error, accident, and history, so design, in this sense, in no way implies a designer but merely some process, such as Darwinian evolution. Here, reverse engineering helps to evaluate emergent properties or even to discern what other constraints have driven an existing design.

Our use of system design is also quite narrow in some respects. There is a vast space of alternatives in systems

engineering, and we are focusing on an extreme “corner” of the design space, taking for granted many tradeoffs that are, in most settings, quite fundamental [7]. The most serious limitation in our focus here, and one we will return to briefly at the end, is that it does not explicitly address the broader human and/or policy issues in complex systems that can dominate the technology. However, within technology, we have the following assumptions about the designs of interest here: 1) They are highly modular (versus integrated); 2) they use diverse components that are imperfect (versus perfect); 3) they have late (versus early) binding of functionality. This allows for 4) a diversity of evolvable (versus frozen) capabilities and behavior; and 5) these systems have fast (versus slow) pace of change, and adaptive (versus preplanned) behavior via distributed (versus centralized) control, with extensive use of feedback. Most systems engineering starts with these (perhaps more fundamental) design decisions, each of which has its own associated tradeoffs. We focus on these specific features because they reflect the choices *already made* in existing net-centric technologies and biology, and our aim is to elaborate on the subsequent issues that arise. We also take for granted pervasive use of modeling and simulation.

C. Design Strategies for Robustness

Given these restrictions, we briefly review some aspects of strategies that are most commonly used to build robustness into system design. The most obvious is to reduce component uncertainty through “ultraquality” design and manufacturing. Engineers use digital electronics in part because they have near-perfect repeatability, with an acceptable overhead in terms of speed, energy, and precision, allowing systems of enormous complexity. However, not all components can be ultraquality. Fortunately, in many cases, the careful organization of unreliable components (e.g., in parallel) can increase the aggregate robustness of the system. For example, building functional redundancy (i.e., backups) of critical components or subsystems is a popular and well-studied approach to robustness, but redundancy *per se* is a limited and blunt tool. More broadly, there has been for some time a large academic literature on system reliability (e.g., [10] and [11]). These strategies can be effective at providing robustness in the face of component uncertainty, but they do not help to achieve robustness to the external environment, and the latter drives most of system complexity. To address this, and also provide effective redundancy management of diverse components, system designers also use a variety of protocol-based strategies.

System sensors can monitor system performance; detect individual component wear, damage, or failure; and/or identify external threats and perturbations to the system. While the ability to maintain an appropriate level of situational awareness of both components and environment is critical for complex systems of all types, such visibility is of limited value if the system cannot act upon that information, and taking appropriate action typically drives complexity far more than does sensing. Using feedback interconnection of sensors and actuators, it is possible to combine components with very different properties to create systems with robustness that far exceeds that of components

separately. For example, it is typically much cheaper to build ultraquality sensor and computation components that use energy efficiently and are extremely robust to noise and environmental shocks, than it is to build comparably inexpensive actuators, which often must have large power requirements and hard limits on achievable accuracy. Feedback control can blend powerful but sloppy actuators with ultraquality sensors to create systems that approach the power of the actuators and the ultraquality of the sensors.

This ability to mix extremely heterogeneous components gives system designers enormous flexibility in exploiting tradeoffs at the component levels to provide effective tradeoffs at the system level. Since this strategy is now ubiquitous, most of the complexity in highly engineered or evolved systems is in control processes that regulate the internal state and respond to external changes. For example, in biological regulatory systems of increasing functional complexity, recent work has shown that communication and control systems occupy an increasing fraction of the information and costs of the systems [12]. These automated control systems are the primary source of RYF in complex systems, since the same systems that provide robustness under normal operating conditions can yield extreme fragilities if they fail or are hijacked.

Software bugs have become a dominant source of cost in many advanced systems. “The Ariane 5 space launch vehicle was destroyed on its initial flight in a classic common mode failure. The software on the primary flight control computer caused the computer to crash shortly after launch. The dual redundant system then switched to the backup flight control computer, which had failed as well moments before for exactly the same reason that the primary computer failed. Ironically, the software failure was due to code leftover from the Ariane 4 and not actually necessary for the phase of flight in which it was operating” [7, p.16]. An important point here for net-centric infrastructures is that the advent of ultraquality digital hardware merely shifts fragilities from hardware component failures to software verification, which is driving the continued development of formal methods in CS [13].

In the end, creating nominal functionality under ideal conditions usually requires only a small fraction of the complexity required to make that functionality robust not only to scale, size, interconnection, and/or interaction but also to an enormous variety of perturbations and variations in the environment and components. Indeed, the emergence of complexity can often be seen as a spiral of new challenges and opportunities that organisms and/or technologies exploit, but which also lead to new fragilities, often from novel perturbations. When successful, fragilities are met with increasing complexity and robustness, which, in turn, creates not only new opportunities but also new fragilities, and so on. Managing or, ideally, preventing this “RYF complexity spiral” [14] remains a central challenge in engineering, medicine, and human society.

D. RYF in Practice

Biological and many technological systems are robust and evolvable in the face of even large changes in environment and system components, yet they can simultaneously be extremely

fragile to novel perturbations. The ubiquity of this RYF behavior is arguably their most crucial emergent property, since their persistence or evolution by natural selection is actually a system constraint. Whereas previous work [14], [15] has detailed key features of RYF in biological and technological systems, our aim here is to expand on the implications for RYF as an emergent constraint for design, as well as the challenges associated with managing RYF tradeoffs. While RYF extremes seem uniquely modern due to the scale of human technology, these RYF features of complex systems actually appear in some form on all time and space scales, from the tiniest microbes and cellular subsystems up to global ecosystems, and from the oldest known human social systems to our latest technologies and their impact on our ecosystems.

New evidence suggests that prokaryote life began deep underground, largely impervious to the late heavy bombardment (LHB) from 4.5–4 Gyr ago that routinely vaporized surface land and water (and also wiped out any records of life or of geological activity, making direct confirmation of these claims difficult) [16]. When the LHB ended, organisms evolved increased complexity to occupy new ecological niches, but with fragilities to new perturbations, which drove additional spirals of complexity, robustness, and new fragility. Most dramatically, the mass extinctions that punctuate the last 1 Gyr dramatize RYF tradeoffs between organisms specializing to exploit the rarity of a particularly perturbation (e.g., large impacts), but becoming catastrophically fragile and thus extinct when they finally do occur. The biosphere as a whole remains highly robust because by most measures—such as biomass, cell count, or genomic diversity—bacteria still dominate, although this may be little consolation to humans.

Modern humans have new (emergent) RYF medical problems, with metabolic syndrome (diabetes, obesity, and inflammation), autoimmune diseases, and cancer being side effects of biological control and compensatory mechanisms so robust as to normally go unnoticed. Symptoms such as insulin resistance, fat accumulation (including elevated intramyocellular triglyceride, or IMTG), inflammation, and proliferative growth are all also crucial elements of the robust control systems, essential for successful hunter gatherers, which maintain glucose for the brain while fueling muscles, even after periods of fasting, and orchestrate repair and regeneration in response to moderate and survivable damage and disease [17]. Cancer too involves the hijacking of development and/or regeneration controls to create cell-level robustness with emergent organism-level fragility [18].

RYF complexity is not confined to biology. During the last century, engineering complexity, in all its aspects, has exploded wildly, with profound impact that is both positive and negative. Most obvious is the ubiquity, diversity, and complexity of digital hardware and associated software. Together with modern control theory, this enabled an aerospace industry with routine global and space travel, as well as ballistic missiles. The Internet has grown from a simple research network to a critical communication and computing system, enabling an enormous diversity of applications. Lean, global manufacturing, supply chain management, and sophisticated financial risk management strategies have made the whole world an efficient

integrated industrial platform. Yet, almost nothing appears sustainable in the long run, and catastrophic cascading fragilities seem increasingly commonplace.

Thus, a hallmark of both biology and advanced technologies is that they exhibit extremes of robustness and fragility, and highly evolved or well-designed systems at best effectively manage the resulting tradeoffs. The complexity of technology is exploding around us but in ways that remain largely hidden. Modern institutions and technologies facilitate robustness and accelerate evolution but enable catastrophes on a scale unimaginable without them (from cascading failures in networks to market crashes, war, epidemics, and global warming).

E. RYF in Theory

Despite its importance, the RYF nature of complex systems, and net-centric infrastructures in particular, remains poorly understood outside narrow technical domains. Yet, within these domains are theories of emergent constraints that are arguably among the most profound insights of the last century and have become standard material in undergraduate engineering courses. Examples (often associated with the names of researchers who were influential in their development) come from the fields of communications (Shannon's channel capacity theorem), control (Bode's integral formula), and computation (Turing, Gödel, and undecidability, as well as NP hardness), plus earlier constraints from thermodynamics (Carnot) and constraints on energy and entropy. In numerical analysis and operations research, high computational complexity is directly connected to problem fragility through the notion of ill-conditioning. Bode's integral formula is the most explicitly RYF as a constraint on the robustness of feedback control, but Shannon's channel capacity theorem limits the robustness of communication.

Turing and Gödel showed that even short theorems may have no short proofs and also that easily described tasks like checking whether a computer program will halt (the classical halting problem) can be undecidable. While two gravitationally interacting bodies yield simple orbits, Poincare showed that the motion of even three interacting bodies can be chaotic and hard to predict in the long term (i.e., fragile to initial conditions). Indeed, these theories all have explicit RYF hard limits on what is achievable, either absolute or in expectation, together with methods to design systems that approach these hard limits.

These emergent constraints are at the heart of their respective subjects and so important that we sometimes refer to them as *laws*. Unfortunately, while the aforementioned theories have been dramatic success stories individually, they provide a poor foundation for an integrated theory of robustness. Each depends on specific and largely orthogonal system, component, and protocol constraints. However, new unifications are encouraging. As one example from a growing literature, [19] integrates aspects of information and control theories. Ultimately, a single theory instead of four separate ones might provide a foundation for a theory of protocol design, so it is an important goal of current research.

While there has been some attempt to create a more integrated RYF complexity theory [14], [15], [20]–[26] that con-

tains each of the existing theories as a special case, much remains to be done. Recent developments in technology, biology, medicine, and mathematics are fragmented but offer the promise of a coherent subject. Despite its diversity, this direction shares the aim to create robust (e.g., evolvable, scalable, efficient, etc.) systems, rigorous theories, and efficient algorithms despite limited resources, uncertain environments and components, and hard limits on what is achievable. It is increasingly, if belatedly, focusing on predicting and avoiding the unintended consequences (a euphemism for ignored fragilities) of technological complexity. The details are unfortunately beyond the scope of this paper; however, there are common organizing principles that can aid in a primitive understanding and can be illustrated with familiar and ubiquitous systems and their architectures.

III. SYSTEM ARCHITECTURE

Understanding RYF and organization requires understanding *architecture*, a term that is widely used, but rarely formalized, when discussing networks of all kinds. Architectures can be designed, evolved, or both but fundamentally are about handling or creating the various constraints described earlier to facilitate “good” solutions among competing tradeoffs. While there does not yet exist a coherent theory of architecture, there are now enough examples from biology and advanced technology that a “comparative physiology” of architecture can illuminate the fundamental research issues.

We next present some examples from common experience that are intended to make the notion of RYF complexity, organization, and architecture more concrete. Most modern engineering usage of “architecture” focuses on the elements of structure and organization that are most universal, high level, and/or persistent. System architecture must facilitate system-level functionality as well as robustness and evolvability to uncertainty and change in components, function, and environment. As with design, architecture includes both systems (or sets of systems) as artifacts and studies architecture as the modularity, interfaces, functional decomposition, etc., that form their actual structure or organization, as well as architecture as a set of principles used to guide design and construction [7].

A. Case Studies

1) *Buildings*: If the *function* of a building is to provide structure for organizing human activity in a way that is naturally hierarchical and somewhat scalable—i.e., from furniture to rooms to buildings to cities—then much of what we require from buildings (e.g., protection from the environment) is really robustness. Robust buildings are complex because they are constructed from a highly structured interconnection of a mix of heterogeneous components having very different individual robustness features, and these components are integrated into a collective whole that draws on specific features of each. For example, the load-bearing structures might, by themselves, have almost no other functionality, providing only a scaffolding on which to hang additional functional elements. Roofs might bear loads poorly but are waterproof. Insulation materials used

to prevent heat transfer, or opaque materials for privacy, or transparent materials to permit light might be neither structural nor waterproof, depending on their location in the building. To get real robustness in a building, one must consider tradeoffs in the use of these components, and complexity comes primarily from their integration and any control systems for power, water, waste, and heating and cooling that are used to provide robustness for these dynamic processes.

In this domain, our thinking about function and robustness is so intertwined that it is hard to separate them and the aforementioned discussion may seem strange at first glance. However, if we imagine what a minimally functional building without robustness would be like, then providing this basic function is greatly simplified. For example, in virtual worlds (e.g., computer games such as *Second Life*), the essential features of building function (the spatial layout that organizes human activity) and aesthetics (the visual impact of the spatial layout) can be effectively simulated. Without the robustness needs of the real world, their structures can be implemented with extremely simple homogenous virtual materials having no hidden structure or complexity. More generally, buildings in regions where robustness to weather is not needed may have greatly simplified walls and/or roofs. Moreover, the main difference between a temporary (e.g., modular) building and a permanent structure is typically not minimal function but robustness (and aesthetics). The point is that the hidden complexity of a real robust physical building is vastly greater than what is required for a minimally functional, perhaps virtual, one.

The classical view of architecture has focused on not only buildings, their function, and robustness to perturbations from the environment (including possibly attack) but also aesthetics. A central debate in architecture has always been the relationship between function, robustness, and aesthetics, with some arguing that they are inextricably intertwined in ways that are easily perceived but are hard to articulate [27], [28]. Our aim is not to resolve these deep debates but to simply use this very familiar domain to highlight the interplay between robustness and complexity more generally. Note that engineers also often reject a design because it is “ugly” or “a kludge,” and many design philosophies argue for simplicity whenever possible [7], so the aesthetics of system architecture has broad interest. A central challenge for a theory of complexity and architecture is to formalize these intuitive notions that “complexity implies fragility.”

2) *Aerospace Systems*: A robust functional aesthetic building can have an overall structure that remains essentially unchanged for decades or even centuries and yet provides a flexible platform for organizing human activities that itself can change wildly and in ways that are utterly unpredictable in detail to the original architect. This view of architecture as “constraints that deconstrain,” as creating platforms that create robustness and evolvability by accepting certain fixed constraints, applies to advanced technologies and biology as well as to ancient buildings. For example, the technology involved in the Boeing 777 aircraft has proven to be remarkably robust in the sense that its design has enjoyed successful operation amid varying weather and environmental conditions without a single fatality since its launch over a decade ago. It is, like

buildings, a useful example because it is very familiar and there are a multitude of books, Web sites, and television shows that provide accessible expositions of the 777 and its design and manufacturing process.

In short, aircraft flight depends fundamentally on lift, drag, propulsion, and control. Lift and drag come primarily from the shape of the vehicle, and propulsion can be modeled (as in a wind tunnel) or simulated (in a computer) sufficiently that this relatively simple representation captures the basics of aircraft function in ideal conditions. The real complexity lies in the control of the aircraft, involving sophisticated feedback systems that make the system far more robust than its (also fairly reliable) components despite uncertain conditions, rich feedback diagnostics for preventative maintenance to replace components before they fail, and judicious use of ultraquality components.

The 777 is the canonical example of a vehicle platform that will persist for decades, with electronics and other subsystem updates but without significant change to the shape of the vehicle. Moreover, the need to make many copies allows one to amortize the cost of design, which has famously relied heavily on extensive modeling and simulation. Another important architectural feature of the 777 was the need to design simultaneously both the vehicle and its manufacturing process, which is itself a study in organized complexity (and the newer 787 illustrates its RYF features). The 777 is, in turn, part of a larger air transport system, and system integration is a greater problem in aerospace—where design is shifting from platforms to systems and “system of systems”—than for buildings. The overarching fragilities of our transportation infrastructure can be glimpsed at this vehicle level and involve more broadly the unsustainability of burning fossil fuels, of which commercial aircraft are only a small part.

The recent U.S. missions to Mars provide an equally familiar set of aerospace systems that had both spectacular successes and failures. Rovers worked beyond all expectations, yet whole missions were lost to software failures as apparently trivial as in the Ariane 5 failure. The components of these systems can be tested, but the full system operation cannot in a way far more extreme than the 777, and the complexity challenge is dominated by creating robustness to uncertainties of the full operational space environment.

3) *Internet*: The strengths and weaknesses of modern organized complexity are exemplified by the Internet architecture, which will be a central case study in our review of complexity. The Internet was conceived and created, and it has grown from a research network to critical infrastructure with much less theoretical support than comparable aerospace systems. Recently, a coherent theoretical framework built on optimization and control theory has been developed, which retrospectively confirms the intuition of the original engineering (e.g., [6], [29], and [30], and the references therein). However, the path to this theory was not straightforward, and the engineers clearly led the theorists (this too is typical). On the surface, the Internet is a communication network linking computers to function as a distributed operating system. Its hidden core protocols [Transmission Control Protocol (TCP) and Internet Protocol (IP)] involve feedback control systems that dynamically

manage the network resources, and whose design made little use of theories from communication and computational complexity, although both inform current research. In fact, clever engineers ignored theoreticians' doubts, brilliantly proposed some basic design principles, and used their intuition and much tinkering (as always) to build an architecture well beyond any existing theory. Yet, the Internet's main fragilities, including the lack of security and flexibility, remain daunting challenges, and even more damaging fragilities will arise as the Internet's role expands.

As case studies, the 777 and the Internet have important differences. The 777 and Mars rovers are bounded systems in the sense that there is a hierarchical "parts list," and central control over when and where new parts can be added. In contrast, the Internet is unbounded in that any device conforming to the standard interfaces and protocols can be connected and function, which is a crucial design requirement. The 777 also has well-defined interfaces, but bounded so that all the parts and connections are known before assembly, unlike the Internet, which is more fundamentally about the interfaces and protocols than the physical artifact. As a result, the use of a well-understood mix of modeling and simulation, prototyping, testing, and experimentation famously enabling the 777 is relevant but very different for the Internet.

Just the parts list in modern buildings, the 777, or the Internet cannot be called simple by any measure, yet most of this complexity remains hidden in normal operation, with the systems typically far more robust than their components. This enormous internal heterogeneity and complexity, together with external functionality and robustness, is the modern hallmark of organized complexity. However, complexity in technological and biological networks is driven not by the number or diversity *per se* of the parts themselves but by the control systems that manage the interaction among components. In technology, this is evidenced by an explosion in complexity in computer networks, automobiles, airplanes, supply chains, package delivery, etc., where emphasis is increasingly placed on *where*, *when*, and *how* more than *who* and/or *what*. Similarly, in biology, a count of protein-coding genes is only weakly correlated with organized complexity of organisms [31]. Thus, organized complexity is not merely about robustness but about the management of functional robustness and the isolation of fragility.

B. Protocol-Based Architectures

Protocols are at the heart of architecture, and the parallels between the Internet and biology are particularly striking, although the level of detail that we can explore here is severely limited, and the ideas are new and admittedly controversial. Biologists Gerhart and Kirchner [4] nicely capture the role of protocols in the phrases "constraints that deconstrain" and "facilitated variation." They describe how constraints in the form of universal shared protocols provide a platform for diverse functionality and robustness (modularity, efficiency, and scalability) by *facilitating* ("deconstraining") large but functional *variation* on which selection can act. This controlled process on the variation side greatly accelerates evolution over

what is possible by the slow accumulation of small changes by selection acting on random point mutations. Antibiotic resistance by horizontal gene transfer (HGT) is perhaps the most familiar example, but biologists have now cataloged endless examples. Although the phrasing differs, this kind of modularity is exactly what engineers would describe as the essence of architectures [7], and expanding on this observation is a central theme of this paper. In *protocol-based architecture (PBA)*, the *protocols* (rules of interaction that persist) are more fundamental than the *modules* (which obey protocols and can change and diversify). PBAs facilitate coherent and global adaptation to variations in both components and the environments on a vast range of time scales despite implementation mechanisms that are largely decentralized and asynchronous.

1) *Protocols and the Internet*: The TCP/IP protocol suite enables adaptation and control on time scales from the sub-microsecond changes in physical media to the millisecond-to-second changes in traffic flow, to the daily fluctuations in user interactions, and to evolving hardware and application modules over years and decades. This dramatizes how a PBA's constraints facilitate both robustness and evolution, and thus, the nature of the protocol constraints is the essence of architecture. "Layering" is the most important architectural feature (although packets are the more obvious) and captures the idea that each layer in the protocol stack provides services to the layer above by utilizing and hiding the resources in the layers below [6], [32]. Reading this paper on a laptop and the Internet would require application software (e.g., web browser) and physical layer hardware (superficially a display and keyboard, but more essentially processing, memory, and interfaces) both obeying protocols on files and signals. These obviously allow access to an incredible diversity of software and hardware resources, but it is the hidden TCP/IP layers that are actually more fundamental.

The protocol stack is (perhaps unfortunately) called an "hourglass" because a thin hidden "waist" of universally shared feedback control (TCP/IP) layers sits between a vast diversity (e.g., wide range) of visible upper (application software) and lower (hardware) layers. Roughly, the IP layer controls the routes for packets on the physical layer below. TCP then uses this virtualized network resource and controls the rates of application flows from above and uses acknowledgement-based retransmission to guarantee delivery if any packets are lost. This allows "plug and play" between modules that obey the shared protocols; any set of applications that "talks" TCP can run transparently and robustly on any set of hardware that "talks" IP. A crucial point that we will return to later is that such layered PBAs deliberately do not constrain the topologies within layers, so there is emphatically no sense in which network connectivity is an aspect of architecture.

2) *Protocols in Biological Networks*: The differences between the Internet and biology are of course more obvious and profound than the similarities. Briefly, the upper Internet layers are all software, whereas in biology, all layers involve chemistry. The Internet imports electric power and hardware from manufacturing processes external to it, although the design and manufacturing of computer hardware, not discussed here, is also finely layered. In contrast, the bacterial cell must not only

have the instructions or recipes to make its components, such as enzymes, but also make and control its energy and redox status, cell envelope, etc. This massively autocatalytic nature of the cell has no parallels in the existing Internet, but this will change in future cyber-physical networks that must have sustainability as an additional robustness requirement.

A striking similarity between life and advanced technologies is their reliance on PBAs. Evolvability, from microbes to ecosystems to IP-based networks, illustrates how dramatic, novel, dynamic, and distributed changes on all scales of time and space can also be coherent, responsive, functional, and adaptive. New genes and pathways, laptops and applications, and even whole networks can all “plug and play” as long as they obey the appropriate protocols. Similarly, a PBA in biology and its control mechanisms facilitate both robustness and evolvability despite massive impinging pressures and variation in the environment. With the most obvious example involving the table of codons at the DNA layer, biology’s universally shared set of protocols is more fundamental and invariant than the modules whose control and evolution are facilitated by them.

Metabolic and protein–protein interaction networks form the bacterial cell application layer. These networks are constrained at the system level by unpredictable intracellular and extracellular environments and at the component level from physiochemical laws governing the individual parts (e.g., proteins and enzymes) and their reactions (i.e., conservation laws). Remarkably, despite these huge differences in details, the main architectural feature of a layered hourglass protocol stack has striking parallels to the Internet. Briefly, in the microbial biosphere, DNA is the “lowest” layer with the genomic information about how to manufacture proteins as a fundamental resource that must be managed. (This is very unlike Internet hardware as the resource although, interestingly, very much like the design of that hardware, a topic that is beyond the scope of this paper.) Thus, genes that “talk” the central transcription and translation protocols (the analog of TCP/IP) can be expressed on demand by the call, as well as move by HGT, greatly accelerating evolution in a kind of “bacterial Internet” [33]. However, also like the technological Internet, the functionality of new HGT genes and the proteins is enhanced by having additional shared protocols at the cell’s “application” layer, such as group transfers and carriers in metabolism, and conserved residue pairs in signal transduction. Thus, selection acting at the protocol level could evolve and preserve shared architecture, essentially “evolving evolvability” [4].

At biology’s application layer, enzymatically catalyzed reactions are controlled by this richly layered architecture. Allosteric control and a huge suite of post-translational modifications of proteins, and the rapid changes in location of macromolecular modules enable adaptive responses to environmental signals or alterations on rapid time scales. In the DNA and RNA layers, transcriptional and translational controls (and, in eukaryotes, regulation of alternative splicing and editing) act on somewhat longer time scales. On still longer time scales within and across generations, the sequences of the DNA itself can change not only through random mutation but also through highly structured and evolved mechanisms, such as HGT, that facilitate the generation of adaptive diversity.

Furthermore, as biologists dig deeper past the superficiality of sequence data into the complexity of regulation, they unearth additional layers of control that are fundamentally similar to those in advanced technologies, with particular interest recently in RNA-based regulation. There is seemingly no limit to the ingenuity that biology uses to create additional layers of sophisticated control [33].

3) *RYF in PBAs*: Bacterial and Internet PBAs underlie both their robustness and fragility. PBAs allow typical network behavior to be fine-tuned through elaborate but hidden control systems and thus appear boringly robust despite large internal and external perturbations. PBAs also facilitate evolution, from microbes to humans and from an academic research network to a global information infrastructure. As a result of layering and control, complexity and fragility remain largely hidden, often revealed only by catastrophic failures and often only after the system has absorbed multiple insults without apparent consequences. Large structured rearrangements are tolerated by control systems that reallocate network resources, easily conferring robustness to outright failures of components and brute force attacks, with disastrous fragilities to small random or targeted changes that subtly violate core protocols. The greatest fragility in PBAs is from parasites and predators, who hijack and consume universal and standardized interfaces and building blocks. Hosts and prey, in turn, leverage PBAs to create diverse immune and protective strategies, and a PBA-enabled arms race spiral ensues. Antibiotic success and resistance illustrate both sides of this RYF spiral, as does the enormously complex challenge of Internet security (and also perhaps the dominant role that both government and financial institutions play in modern economies).

From today’s perspective, the Internet architecture was brilliant in many respects, such as using packets and layering, but it had little theoretical support for ongoing design, which has primarily been tinkering and intuition with considerable experimentation via simulation or prototype. This is much like evolution but without the bacterial biosphere’s relatively unlimited time and scope (billions of years and 10^{30} cells) to explore design alternatives. It has aspects of human evolution where our historical legacy from fish to primates endures. In both cases, the core protocols have been such amazing platforms for innovation that changing them is either truly impossible or appears nearly so, but a fundamental reassessment might benefit the health of both humans and their technology [8]. For sure, as technological visions increasingly emphasize ubiquitous control, communications, and computing—with systems requiring a high degree of autonomy, adaptation, evolvability, scalability, and verifiability—a more predictive, rigorous, coherent, and reasonably complete mathematical theory underpinning Internet technology and its future architecture is needed [6].

4) *Designing Protocols*: Designing architectures is fundamentally about the choices associated with all four types of constraints (systems level, component level, emergent, and protocols) that give rise to organized system structure. This constraint-based view of architecture naturally invites *optimization* as an important tool for the analysis and design of network protocols. Recent analysis and design of IPs has benefited tremendously from a developing mathematical

framework that views various protocol layers as carrying out asynchronous distributed computation over the network to optimize a global objective function subject to resource constraints in the network [6]. Under this view, different layers iterate on different subsets of the decision variables using local information to achieve individual optimality. Taken together, these local algorithms attempt to achieve a global objective. Such a framework exposes the interconnections between protocol layers as different ways to modularize and distribute an otherwise centralized function, formalizes the common practice of breaking down the design of a complex system into simpler modules, and provides a top-down approach to both the systematic design of layers and the tradeoffs between competing design objectives. The application of this framework to reverse engineering in biology is promising but so far relatively less developed [21], [22], [34]–[36].

IV. ALTERNATIVE VIEWS OF COMPLEXITY

Organized complexity, as described in this paper, is so fundamental and consistent in the way that engineers, biologists, and physicians see the world that it is sometimes hard to imagine a different view. However, in fact, much of the mainstream effort within the physical sciences on the study of complexity and most of the popularizations over the last several decades have followed a view that is not merely contradictory but profoundly incompatible. This alternate view downplays function, organization, systems constraints, and design, focusing instead on emergent features of disorganized systems, particularly their fragilities. The result is a disconnect that is far more severe than the technical fragmentation within engineering, and it remains a persistent source of confusion. In this section, we put this disconnect in an historical context, develop a broader taxonomy for complexity that includes both threads, and highlight some recent research that underscores the difference using illustrative examples from the Internet and biology. We expect that this section will have the most divergent responses from readers, an issue that we will address throughout.

A. Historical Perspective

The use of the term “organized complexity” can be traced to Warren Weaver [37], [38], who, in 1948, contrasted three classes of problems facing science: problems of *simplicity* versus *disorganized complexity* versus *organized complexity*.

Simplicity: Weaver briefly mentions simple problems as those involving a small number of variables, and he attributes much progress in the physical sciences and technology advances to successful application of the scientific method to simple problems. With modern computation, the number of variables *per se* has receded in importance, but Weaver’s other distinctions remain relevant. We can update Weaver by saying that a system is simple if it has “simple questions” (i.e., models, theorems, experiments, and computations) to which there are “robust answers.” By *simple questions*, we mean only that the questions of interest can be posed using models that are readily manageable and easy to describe, that theorem statements are short, and that experiments are elegant, are easily described, and

require minimal interpretation. By *robust answers*, we mean that theorems have simple counterexamples or short proofs, algorithms scale, and simulations and experiments are reproducible with predictable results.

There are many classical examples of such simplicity: the pendulum as a simple harmonic oscillator; simple RLC circuits; the interaction of two bodies via gravity; and simple Boolean logic circuits as implemented in much digital hardware. Many questions regarding their behavior are simple, as are the answers. Most early educational exposure to mathematics involves simple theorems with short proofs, and this is always the goal of even the most advanced mathematics. What we view as simple necessarily evolves over time, and modern technology has greatly expanded what we might think of as simple. The triumph of reductionist science has been to reduce the *apparent complexity* of the world directly to an underlying simplicity. Physics has always epitomized this approach, and recently, molecular biology has successfully mimicked physics.

Disorganized Complexity: Weaver also describes an important development in scientific thinking that was taking hold in the early 1900s. Rather than pursuing simplicity with a few variables, “imaginative minds went to the other extreme,” focusing on problems with asymptotically infinite dimensions and developing “powerful techniques of probability theory and of statistical mechanics to deal with what may be called problems of *disorganized complexity*.” In Weaver’s canonical example of billiards, classical dynamics accurately predict trajectories of a small number of balls on a table, and expert players can robustly control these trajectories by keeping them relatively simple. As the number of interacting balls increases, robust predictions become intractable, either computationally or for players. However, as the size of the table and the number of balls become very large, specific problems involving ensemble average properties actually become easier and more robust, and statistical methods apply. In essence, what Weaver called disorganized complexity was ultimately a way to extend the “simple” to large ensembles and, in his thinking, was not really about complexity at all.

Organized Complexity: Weaver used the term “disorganized” here to emphasize that “the methods of statistical mechanics are valid only when the [billiard] balls are distributed, in their positions and motions, in a helter-skelter, i.e., a disorganized, way. For example, the statistical methods would not apply if someone were to arrange the balls” and their movements in some highly organized manner. While Weaver acknowledges the prevalence of disorganized complexity in many important systems, he notes the importance of an intermediate class that “does not depend primarily on the fact that the number of variables is moderate. . . . The really important characteristic of the problems of this middle region, which science has yet little explored or conquered, lies in the fact that these problems, as contrasted with disorganized situations with which statistics can cope, show the essential feature of *organization*. In fact, one can refer to this group of problems as those of *organized complexity*.”

Weaver’s motivating examples are primarily those of biological systems, and he points hopefully to early computers and operations research as potential tools to address organized

TABLE I
TWO DIMENSIONS OF COMPLEXITY

	simple models	complex models
robust behavior	simplicity	organized complexity
fragile behavior	disorganized complexity plus chaos, criticality, scale-free	irreducible complexity

complexity. Weaver's vision was shared by many of his contemporaries, and diverse authors used his paper as a departure point. Organized complexity was formalized by Norbert Wiener in his development of *cybernetics* as a common theoretical core to the integrated study of technological and biological systems [39]. Jane Jacobs in 1961 [40] argues that cities should be "identified, understood, and treated as problems of organized complexity" but that "... the theorists of conventional modern city planning have consistently mistaken cities as problems of simplicity and of disorganized complexity, and have tried to analyze and treat them thus. ... These misapplications stand in our way; they have to be hauled out in the light, recognized as inapplicable strategies of thought, and discarded." Thus, by 1960, it was clear to many engineers, mathematicians, biologists, and social scientists (at least) that organization was more than the statistics of random ensembles.

B. Unified Taxonomy

Table I extends Weaver's taxonomy minimally to include two recent developments: 1) the explosion in organized complexity and modern technology and 2) the rich subject, primarily within physics, that grew out of chaos, fractals, and criticality and has been dominating the study of disorganized complexity since 1960. One dimension in Table I is the complexity of the system descriptions and/or models, and the other is behaviors in response to perturbations in descriptions, components, or the environment (robust versus fragile). Simple models can have a few components or an asymptotically large number of randomly interacting homogeneous components, while complex models may have a large number of heterogeneous components.

Even simple models (and theorems) can exhibit extreme fragility and unpredictability (and long proofs). Organized complexity produced modern network technology by managing the resulting fragility–complexity spiral. Yet, it is possible for systems to have gratuitous fragility unrelated to any tradeoffs. This is very unlike RYF organization, where fragilities are the side effects of mechanisms for robustness, and the tradeoff between the two is the focus of design or evolution. Such gratuitous fragility is celebrated in modern treatments of disorganized complexity, which are more commonly known as the "new sciences of complex networks" (NSCN) or some variant. Although organized complexity pervades mainstream science in the form of modern technology, the NSCN view of disorganized complexity dominates modern scientific thinking about complexity in a way that would likely have surprised and concerned thinkers like Weaver, Wiener, and Jacobs.

NSCN offers an integrated, unified, and popular view of "simple yet fragile" complexity and has yielded a vast literature on power laws, fractals, self-similarity, and self-organization.

Among the most popular of these theories, *self-organized criticality (SOC)* [41], *edge-of-chaos (EOC)* [42], and *scale-free networks (SFN)* [43] are different in detail but share fundamental features, which are summarized in Table II. For example, their analysis focuses on random ensembles (of lattices or cellular automata, Boolean networks, and graphs, respectively) that are minimally tuned via a single parameter (an order or "friction" parameter, connectivity or correlation structure, and preferential attachment, respectively) to particular fragile configurations (a critical phase transition, a bifurcation point, or a power-law node degree distribution). The differences in details are indeed minor compared with their shared similarities and the extreme differences with organized complexity.

A major source of confusion is that NSCN and organized complexity typically use different and often opposite meanings for the same words, such as architecture, robustness, fragility, networks, and "not random" (even different notions of "proof" and evidence). Consider the start of a recent review article in NSCN [44]: "By definition, complex networks are networks with more complex architectures than classical random graphs with their "simple" Poissonian distributions of connections. The great majority of real-world networks. . . are complex ones. The complex organization of these nets typically implies a skewed distribution of connections. . . ." Note here that "architecture" and "organization" have nothing to do with system function, protocols, or any aspect of design but simply mean "graph topology."

Power laws (or *scaling distributions*) have been another source of confusion because of their opposite interpretations and incompatible statistical treatment. In the NSCN and modern physics literature, power laws are viewed as "signatures" of specific mechanisms, namely, critical phase transitions [41] and preferential growth [43], because these mechanisms can generate power laws in simple models of disorganized systems. Adding to the confusion is that common NSCN statistical techniques, such as the use of binned frequency plots, make it easy to "discover" power laws where none exists or to mischaracterize their most salient features (see [45, Sec. 2.1]). However, a broader view of high variability and power laws reveals a long and rich history outside physics [46]–[48].

Interest in preferential growth models (called the Matthew effect, or "rich get richer") actually peaked in the 1950s in a fierce debate regarding technical details (see [49] and [50]), and studying them became almost passé by 1960 due in part to earlier results on the strong statistical invariance properties of power laws. Specifically, the central limit theorem without a finite-variance assumption shows that power laws are the natural "null hypothesis" for processes involving high variance, just as Gaussians serve as a null hypothesis for low-variance processes. This view suggests that the mere presence of power laws alone implies nothing about an underlying mechanism, nor does it require special explanation beyond high variability (see modern reviews [51], [52] and related discussions [45], [48], [53]).

At issue is how to explain the ubiquity of power laws in real-world data from linguistics, economics, biology, the Internet, disaster statistics, etc. For researchers trained to look at the world through the eyes of disorganized complexity, criticality

TABLE II
CONTRASTING VIEWS OF NSCN AND ORGANIZED COMPLEXITY

	(a) New Sciences of Complex Networks	(b) Organized Complexity
Primitives	random ensembles	structured networks, architecture, design, evolution
Function	none	domain-specific performance and robustness
Architecture	graph topology, connectivity	protocols, layering, constraints that deconstrain
Components	largely homogeneous	heterogeneous, diverse
“Not random”	random but skewed, clustered	highly organized, structured, far from random
Tuning	<i>minimally tuned</i> via an order parameter	<i>highly tuned</i> via protocols
Power laws	“signatures” of criticality, SFN	“more normal than Normal”
Uncertainty	<i>minimal</i>	<i>large</i> (in environment and components)
Robustness	to random rewiring	to common perturbations, targeted attacks
Fragility	to common perturbations, initial conditions, attack	to random rewirings and rare perturbations
RYF	gratuitous	tradeoffs, unintended consequences
Examples	Citations, bacterial lineages, city sizes	Internet, biology, technology

and preferential attachment provide simple explanations. However, from the perspective of organized complexity, it is not power laws *per se* that are important but the presence of underlying high variability. High variability can arise naturally as the result of highly organized/optimized tradeoffs in the design of a variety of systems [54]–[57], and thus, for high-technology and biological systems, power laws serve as the natural statistical null hypotheses for the high variability in RYF systems [15].

C. Case Studies: Opposite Views of the Internet and Biology

NSCN offers a coherent and consistent world view about what is essential and universal in complex networks, and their model abstractions have an appealing simplicity, which is itself always an attractive goal. However, on almost every conceivable dimension, NSCN and organized views of complexity are not merely different but opposite. This can be seen in the books on EOC, SOC, and SFN, in the science literature that preceded them, in broader popularizations (e.g., [58]), and, most saliently, in the familiar case studies of the Internet and biology. Thus, it is important to be clear exactly what are the consequences of these differences and what is crucial versus incidental. Power laws feature prominently in much of the NSCN, along with claims that their presence implies certain underlying mechanisms. As noted previously, this claim is false and can lead to a host of methodological errors, but power laws are not the crux of the issue. The crucial distinction is really the NSCN emphasis on emergent fragilities in disorganized systems, as captured by the other dimensions in Table II.

The Internet provides a clear illustration of the differences between the NSCN view of architecture as graph topology versus the organized view of architecture as layering and protocols. SFN models of other networks, while different in detail, share all of these features, as do earlier SOC and EOC models (e.g., for the canonical SOC example of wildfires, see [57]). In particular, many SFN (and EOC and SOC) models analogous to the Internet have been made for a variety of biological networks. The flaws in this work are similarly the result of modeling biological structure as a disorganized system, but their clarification is technically more difficult (see [35] and [59]–[61]), so we focus on the Internet.

Four of the six most cited papers in *ISI Web of Science* searched with topics “Internet” and “network*” focus primarily

on SFNs [62]–[65]. What launched the SFN genre (grabbing the cover of *Nature* as well as popular news outlets as a previously overlooked “Achilles heel of the Internet”) is a scale-free model of the Internet router layer physical topology [65], in which the following hold: 1) The *primitives* are random router and web graphs without system-level *function* other than graph connectivity; 2) *architecture* simply means graph topology; and 3) SFN *components* are homogeneous, functionless links, and nodes (for routers but also web pages and autonomous systems). There is no mention of the layering, protocols (e.g., no TCP/IP), or extreme diversity and heterogeneity in link and router bandwidths and geography, web and application page size, content, and functionality (both in space and over time as applications and hardware change) that personify the organized Internet.

In SFNs, “not random” means probabilistic but power law distributed, achieved as an emergent property via minimal *tuning* of an order parameter, i.e., in this case, a linear term that determines preferential attachment. SFNs have minimal *uncertainty*, only that nodes and links can be removed or rewired. By definition, SFNs are *robust* to random degree-preserving rewirings, are minimally perturbed by random link or node removal or addition, yet are extremely *fragile* to attacks on high-degree central hubs. One important feature from organized complexity that the SFN router model introduced into NSCN is the idea of a worst case attack rather than merely random node removal. Otherwise, there is no sense in which “not random” in NSCN usage suggests highly organized, designed, evolved, or structured.

In organized complexity, SFNs provide simple models in many domains, including paper citation graphs (ironically, including the scale-free literature itself), bacterial clonal lineages in culture, and city sizes and many other social and economic systems (all well known for decades even before 1960). Since the Internet facilitates enormous diversity at the application layer, good SFN models should be found there as well, but they are very poor models of the Internet itself and most biology. An essential feature of the Internet’s true architecture is that it deliberately places no constraints on its physical or application connectivity. Thus, the connectivity could be anything in principle, but because of high variability in connectivity, link speeds, file sizes, etc., power laws do exist everywhere in the Internet (and in biology). However, this connectivity can be difficult to

measure and is full of complexities symptomatic of its highly organized structure, and both have been the source of errors when studied through the lens of disorganized complexity [66].

When viewed as organized complexity, the structure of the router-level Internet is easily understood as highly constrained at both the system and component levels [25], [45], [67]. Most simply, routers have hard limits on the rate of processing data packets, creating a design tradeoff—router connectivity can range between a small number of high-bandwidth connections and a large number of low-bandwidth connections. It is impossible for a router to handle a large number of high-bandwidth connections (although “large” and “high” are relative and always growing with technology). As a result, any high-degree “hubs” (if they exist at all and regardless of whether their degrees followed a power law) would directly constrain the achievable network performance. In SFN models, these hubs are centrally located, and in random graphs, this is where they must be with high probability [45]. Such central hubs would be such a severe performance bottleneck that their loss would be a secondary concern. While TCP/IP can, in principle, run on any router topology, an SFN would have performance so poor and costs so high as to be prohibitive in practice. In reality, any high-degree hubs in the router-level Internet would necessarily be located at the periphery of the network, where they would aggregate low-bandwidth users and connect them to the high-speed core. The loss of such edge hubs could mean local disruptions for disconnected users, but the core of the network would remain unaffected. Whether edge hubs create power laws in the real Internet is doubtful but also irrelevant [66].

Irreducible Complexity: The “irreducibly complex” combines the complicated descriptions of organized complexity with excessive fragility. It is associated with another contemporary of Weaver, namely, Rube Goldberg, whose famous cartoons depict “comically involved complicated invention[s], laboriously contrived to perform a simple operation” [68], which are themselves far too fragile to persist in the real world but have still become important metaphors. Engineers use Rube Goldberg comparisons as an insult similar to “kludge,” yet engineering students engage in fierce competitions to build their own Rube Goldberg-like contraptions. These are fun, amusing, and even educational because they share features of engineered organization, such as many heterogeneous parts, complex interconnection, high variability and dynamics, and sophisticated control, plus ingenuity, workmanship, and patience. Yet, the outcome is so unlike real engineering: trivial functionality, with hyperfragility to almost everything.

There are other engineering circumstances, such as in security and, particularly, cryptography, where the intractability and fragility of irreducible complexity (although the term is not used here) are actually valuable resources. This notion of “fragility by design” is a rich subject but unfortunately beyond the scope of this paper. As for biology, some properties of particular individuals that seem gratuitously fragile may be “frozen accidents” that have temporarily avoided selective elimination, and these may be intrinsically unpredictable. Biology may have parallels with cryptography, but as in engineering, most fragilities are the inevitable consequences of robustness tradeoffs. Indeed, the extraordinary robustness and evolvability of the

biosphere as a whole suggest that biology is rarely irreducibly complex. In general, the extreme and cryptic fragility of irreducibly complex systems, whether in biology or technology, suggests not intelligent design but the lack of it.

V. IMPLICATIONS, INCLUDING POLICY

This attempt to contrast and ultimately reconcile these very different views of complexity would be a mere academic exercise were it not for the prevalence of complexity in our most pressing challenges where technology and policy intersect. Scientific illiteracy of the public and of policy makers can make it easy for ideologues to zero in on areas of complexity, where explanations are difficult to understand, and muddy the water. However, this is also exactly what happens within science itself around “complexity.” That the broader scientific community can diverge so thoroughly on a topic as relatively straightforward as the Internet is discouraging for reaching the needed consensus on much thornier challenges such as sustainable infrastructure, medicine and health care, ecosystem degradation, global warming, and even evolution.

Thus, there is an urgent need both technically and educationally for a more integrated and accessible treatment of complex systems, and there is some hope. Existing organized theories of a system-level architecture for complex networks are nascent, but it is clear that design of architectures is a topic of increasing importance. As sketched in this paper, progress is being made in our understanding of complex system analysis and design, but considerable work remains. We need more accessible explanations of technical organization and tradeoffs, as well as a greater understanding of how the principles of organized complexity can be used to provide guarantees on the performance and reliability of the system as a whole. We cannot afford to wait to see what emerges from the ongoing integration of the Internet and other critical infrastructures, nor can we trust the validation of such systems to heuristics and simulation. Worse, the existing economic incentives in technology markets, which disproportionately reward time-to-market over quality and thus lead to a “hack and ship” mentality, are a disaster for net-centric infrastructures.

Fortunately, there is much recent work, compatible with the highly organized view of complexity taken here, that focuses on large institutions and their human actors, policies, and economics. It goes by various names combining resilience, security, and/or safety with engineering and/or economics (and which we will denote collectively as RSE) [9]. RSE recognizes that policy can often trump technology and, more importantly and broadly, that technological solutions are inadequate and incomplete if the human actors do not have their incentives aligned with the technology. A particularly consistent RSE theme cautions that “human error,” particularly when persistent, is often actually a symptom of system design problems. RSE urges looking past the proximal cause of any one observable failure to see whether it has merely triggered a fragility inherent to the system, one that should really be addressed with a system solution. RSE emphasizes that any new theory of network architecture must include a realistic treatment of policy and incentives.

In the meantime, NSCN remains a growing activity in academia and funding agencies, and there remains considerable hope (even among folks who concede that the field has not lived up to its hype) in the promise of fundamental and practical discoveries. Presumably, this is a symptom of a real need for better methods to deal with complexity and the frustration that everyone has with the obscure and RYF nature of technology. Unfortunately, the organized and disorganized views of complexity continue to be so incompatible that there is no real dialogue between them, and thus, it is difficult to build on the accessibility and successful popularizations of NSCN. Advocates of NSCN are quick to admit that their models are simple and thus have flaws (as do all models), but argue that this is only because (e.g., like thermodynamics) they avoid nonessential details. Yet, as we have argued, *features of organization are essential details*. Moreover, equally simple models of router and network constraints can more accurately describe realistic topologies and also be extended to describe the far more important elements of layering and protocols.

To Internet technologists and operators, what is truly beyond doubt is that the SFN models and the real Internet are so obviously and extremely incompatible [25], [45], [67] that further discussion would seem unnecessary, and we should move on. Yet, within NSCN, this consensus among experts is completely ignored, and the Internet remains the canonical success story where established tools from physics for studying power laws and disorganized systems uncovered a hidden but catastrophic flaw in the Internet (first overlooked and then denied by the experts). This is underscored in a recent special issue in the journal *Science*: "Today, the scale-free nature of networks of key scientific interest, from protein interactions to social networks and from the network of interlinked documents that make up the WWW to the interconnected hardware behind the Internet, has been established beyond doubt" [69]. Here, the RSE emphasis on systemic flaws in human systems raises provocative questions about the tendency for too many scientists (and not just the public and policy makers) to prefer simple myths over complex but inconvenient truths.

In the real world of network-centric infrastructures, the gratuitous (and presumably avoidable) fragilities highlighted in either NSCN or irreducible complexity would be disastrous far beyond the intrinsic RYF nature of highly organized systems. Yet, there is little consensus on even the most basic strategies to avoid this in real-world networks. In the worst case, we might build increasingly complex and incomprehensible systems that will eventually fail completely yet cryptically, and recent ecological disasters and financial crises are not encouraging. At best, we need to better manage the tradeoff between functional robustness and emergent fragility so that we can appropriately balance the benefits of increased efficiency and convenience with the potential for large-scale disruption. We may not be able to overcome intrinsically RYF aspects of the Internet and other net-centric infrastructures, but we should be able to minimize the potential risk of catastrophic failure.

ACKNOWLEDGMENT

The authors would like to thank the two anonymous referees for their considerable help in organizing this paper and for

their direct contributions to the articulation of ideas on systems engineering and design for robustness.

REFERENCES

- [1] National Science Foundation, NSF workshop on cyber-physical systems: Research motivation, techniques and roadmap, Austin, TX2006, Oct. 16–17. [Online]. Available: <http://varma.ece.cmu.edu/cps/>
- [2] K. W. Hipel, M. M. Jamshidi, J. M. Tien, and C. C. White, III, "The future of systems, man, and cybernetics: Application domains and research methods," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 37, no. 5, pp. 726–743, Sep. 2007.
- [3] D. L. Alderson, "Catching the 'network science' bug: Insight and opportunity for the operations researcher," *Oper. Res.*, vol. 56, no. 5, pp. 1047–1065, Sep./Oct. 2008.
- [4] J. Gerhart and M. Kirschner, "The theory of facilitated variation," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 104, pp. 8582–8589, May 2007, suppl. 1.
- [5] D. X. Wei, C. Jin, S. H. Low, and S. Hegde, "FAST TCP: Motivation, architecture, algorithms, performance," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 1246–1259, Dec. 2006.
- [6] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition," *Proc. IEEE*, vol. 95, no. 1, pp. 255–312, Jan. 2007.
- [7] M. W. Maier and E. Rechtin, *The Art of Systems Architecting*, 3rd ed. Boca Raton, FL: CRC Press, 2009.
- [8] J. Day, *Patterns in Network Architecture: A Return to Fundamentals*. Upper Saddle River, NJ: Prentice-Hall, 2008.
- [9] E. Hollnagel, D. D. Woods, and N. Leveson, Eds., *Resilience Engineering: Concepts and Precepts*. Hampshire, U.K.: Ashgate, 2006.
- [10] C. L. Hwang, F. A. Tillman, and M. H. Lee, "System-reliability evaluation techniques for complex/large systems—A review," *IEEE Trans. Rel.*, vol. R-30, no. 5, pp. 416–423, Dec. 1981.
- [11] R. Billinton and R. N. Allan, *Reliability Evaluation of Engineering Systems—Concepts and Techniques*, 2nd ed. New York: Plenum, 1992.
- [12] J. S. Mattick and M. J. Gagen, "Accelerating networks," *Science*, vol. 307, no. 5711, pp. 856–858, Feb. 2005.
- [13] C. Baier and J.-P. Katoen, *Principles of Model Checking (Representation and Mind Series)*. Cambridge, MA: MIT Press, 2008.
- [14] M. Csete and J. Doyle, "Reverse engineering of biological complexity," *Science*, vol. 295, no. 5560, pp. 1664–1669, Mar. 2002.
- [15] J. M. Carlson and J. C. Doyle, "Complexity and robustness," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 99, pp. 2538–2545, Feb. 2002.
- [16] L. J. Rothschild, "Earth science: Life battered but unbowed," *Nature*, vol. 459, no. 7245, pp. 335–336, May 21, 2009.
- [17] S. R. Stannard and N. A. Johnson, "Insulin resistance and elevated triglyceride in muscle: More important for survival than thrifty genes?" *J. Physiol.*, vol. 554, pt. 3, pp. 595–607, Feb. 2004.
- [18] H. Kitano, "Cancer as a robust system: Implications for anticancer therapy," *Nat. Rev. Cancer*, vol. 4, no. 3, pp. 227–235, Mar. 2004.
- [19] N. C. Martins, M. A. Dahleh, and J. C. Doyle, "Fundamental limitations of disturbance attenuation in the presence of side information," *IEEE Trans. Autom. Control*, vol. 52, no. 1, pp. 56–66, Jan. 2007.
- [20] M. Csete and J. Doyle, "Bowties, metabolism, and disease," *Trends Biotechnol.*, vol. 22, no. 9, pp. 446–450, Sep. 2004.
- [21] J. Stelling, U. Sauer, Z. Szallasi, F. Doyle, and J. Doyle, "Robustness of cellular functions," *Cell*, vol. 118, no. 6, pp. 675–685, Sep. 2004.
- [22] H. El-Samad, H. Kurata, J. C. Doyle, C. A. Gross, and M. Khammash, "Surviving heat shock: Control strategies for robustness and performance," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 102, no. 8, pp. 2736–2741, Feb. 2005.
- [23] J. C. Doyle and M. Csete, "Motifs, stability, and control," *PLOS Biol.*, vol. 3, no. 11, p. e392, Nov. 2005.
- [24] J. Doyle and M. Csete, "Rules of engagement," *Nature*, vol. 446, no. 7138, p. 860, Apr. 2007.
- [25] J. C. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, "The 'robust yet fragile' nature of the Internet," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 102, no. 41, pp. 14 497–14 502, Oct. 2005.
- [26] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ: Princeton Univ. Press, 2008.
- [27] N. A. Salinger, *Anti-Architecture and Deconstruction*, 3rd ed. Solingen, Germany: Umbau-Verlag, 2008.
- [28] C. Alexander, *The Nature of Order*. Berkeley, CA: CES, 2004.
- [29] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, no. 3, pp. 237–252, Mar. 1998.
- [30] S. Shakkottai and R. Srikant, "Network optimization and control," in *Foundations and Trends in Networking*. Hanover, MA: NoW, 2007.

- [31] R. J. Taft, M. Pheasant, and J. S. Mattick, "The relationship between non-protein-coding DNA and eukaryotic complexity," *Bioessays*, vol. 29, no. 3, pp. 288–299, Mar. 2007.
- [32] D. D. Clark, "The design philosophy of the DARPA Internet protocols," *ACM Comput. Commun. Rev.*, vol. 18, no. 4, pp. 106–114, 1988, Proc. ACM SIGCOMM 1988.
- [33] J. C. Doyle, M. E. Csete, and L. Caporale, "An engineering perspective: The implicit protocols," in *The Implicit Genome*, L. Caporale, Ed. London, U.K.: Oxford Univ. Press, 2007.
- [34] M. Morohashi, A. E. Winn, M. T. Borisuk, H. Bolouri, J. C. Doyle, and H. Kitano, "Robustness as a measure of plausibility in models of biochemical networks," *J. Theor. Biol.*, vol. 216, no. 1, pp. 19–30, May 2002.
- [35] R. Tanaka, M. Csete, and J. Doyle, "Highly optimised global organization of metabolic networks," *Proc. Inst. Elect. Eng.—Syst. Biol.*, vol. 152, no. 4, pp. 179–184, Dec. 2005.
- [36] H. El-Samad, S. Prajna, A. Papachristodoulou, J. Doyle, and M. Khammash, "Advanced methods and algorithms for biological network analysis," *Proc. IEEE*, vol. 94, no. 4, pp. 832–853, Apr. 2006.
- [37] W. Weaver, "Science and complexity," *Amer. Sci.*, vol. 36, pp. 536–544, 1948.
- [38] E. F. Keller, "Organisms, machines, and thunderstorms: A history of self-organization, part two. Complexity, emergence, and stable attractors," *Historical Stud. Natural Sci.*, vol. 39, no. 1, pp. 1–31, Winter 2009.
- [39] N. Wiener, *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press, 1961.
- [40] J. Jacobs, *The Death and Life of Great American Cities*. New York: Vintage, 1961.
- [41] P. Bak, *How Nature Works: The Science of Self-Organized Criticality*. New York: Copernicus, 1996.
- [42] S. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. London, U.K.: Oxford Univ. Press, 1993.
- [43] A.-L. Barabási, *Linked: The New Science of Networks*. Cambridge, MA: Perseus, 2002.
- [44] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, "Critical phenomena in complex networks," *Rev. Mod. Phys.*, vol. 80, no. 4, pp. 1275–1335, Oct.–Dec. 2008.
- [45] L. Li, D. Alderson, J. Doyle, and W. Willinger, "Towards a theory of scale-free graphs: Definition, properties, and implications," *Internet Math.*, vol. 2, no. 4, pp. 431–523, 2005.
- [46] V. Pareto, *Cours d'Economie Politique*. Geneva, Switzerland: Droz, 1896.
- [47] M. Mitzenmacher, "A brief history of generative models for power law and lognormal distributions," *Internet Math.*, vol. 1, no. 2, pp. 226–251, 2004.
- [48] E. Keller, "Revisiting 'scale-free' networks," *Bioessays*, vol. 27, no. 10, pp. 1060–1068, Oct. 2005.
- [49] H. A. Simon, "On a class of skew distribution functions," *Biometrika*, vol. 42, no. 3/4, pp. 425–440, Dec. 1955.
- [50] B. Mandelbrot, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. New York: Springer-Verlag, 1997.
- [51] A. Bookstein, "Informetric distributions, part I: Unified overview," *J. Amer. Soc. Inf. Sci.*, vol. 41, no. 5, pp. 368–375, 1990.
- [52] A. Bookstein, "Informetric distributions, part II: Resilience to ambiguity," *J. Amer. Soc. Inf. Sci.*, vol. 41, no. 5, pp. 376–386, 1990.
- [53] W. Willinger, D. Alderson, J. Doyle, and L. Li, "More 'normal' than normal: Scaling distributions and complex systems," in *Proc. Winter Simul. Conf.*, 2004, pp. 130–141.
- [54] J. M. Carlson and J. C. Doyle, "Highly optimized tolerance: A mechanism for power laws in designed systems," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 60, no. 2, pp. 1412–1428, Aug. 1999.
- [55] M. Manning, J. Carlson, and J. Doyle, "Highly optimized tolerance and power laws in dense and sparse resource regimes," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 72, no. 1, p. 016108, Jul. 2005.
- [56] T. Zhou, J. M. Carlson, and J. C. Doyle, "Evolutionary dynamics and highly optimized tolerance," *J. Theor. Biol.*, vol. 236, no. 4, pp. 438–447, Oct. 2005.
- [57] M. A. Moritz, M. E. Morais, L. A. Summerell, J. M. Carlson, and J. C. Doyle, "Wildfires, complexity, and highly optimized tolerance," *Proc. Nat. Acad. Sci. U.S.A.*, vol. 102, no. 50, pp. 17912–17917, Dec. 2005.
- [58] M. Buchanan, *Nexus: Small Worlds and the Groundbreaking Theory of Networks*. New York: Norton, 2003.
- [59] R. Tanaka, T.-M. Yi, and J. Doyle, "Some protein interaction data do not exhibit power law statistics," *FEBS Lett.*, vol. 579, no. 23, pp. 5140–5144, Sep. 2005.
- [60] R. Tanaka, "Scale-rich metabolic networks," *Phys. Rev. Lett.*, vol. 94, no. 16, p. 168 101, Apr. 2005.
- [61] G. Lima-Mendez and J. van Helden, "The powerful law of the power law and other myths in network biology," *Mol. Biosyst.*, vol. 5, no. 12, pp. 1482–1493, Dec. 2009.
- [62] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.
- [63] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Rev. Mod. Phys.*, vol. 74, no. 1, pp. 47–97, Jan.–Mar. 2002.
- [64] M. Newman, "The structure and function of complex networks," *SIAM Rev.*, vol. 45, no. 2, pp. 167–256, Jun. 2003.
- [65] R. Albert, H. Jeong, and A.-L. Barabási, "Attack and error tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378–382, Jul. 2000.
- [66] W. Willinger, D. L. Alderson, and J. C. Doyle, "Mathematics and the Internet: A source of enormous confusion and great potential," *Not. AMS*, vol. 56, no. 5, pp. 586–599, May 2009.
- [67] L. Li, D. Alderson, W. Willinger, and J. Doyle, "A first principles approach to understanding router-level topology," *ACM Comput. Commun. Rev.*, vol. 34, pp. 3–14, 2004, Proc. ACM SIGCOMM 2004.
- [68] Webster's New World Dictionary definition, Examples of Rube Goldberg's art and details of the annual contest to build a Rube Goldberg device are available from www.rubegoldberg.com.
- [69] A.-L. Barabási, "Scale-free networks: A decade and beyond," *Science*, vol. 325, no. 5939, pp. 412–413, Jul. 2009.



David L. Alderson (M'05) received the B.S.E. degree in civil engineering and operations research from Princeton University, Princeton, NJ, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA.

He has been a Postdoctoral Scholar at the California Institute of Technology, Pasadena, and has held Research Fellow positions at the Institute for Pure and Applied Mathematics, University of California, Los Angeles, and at the Santa Fe Institute, Santa Fe, NM. Since 2006, he has been an Assistant Professor with the Operations Research Department, Naval Postgraduate School, Monterey, CA. His research focuses on the function and operation of critical infrastructures, with particular emphasis on how to invest limited resources to ensure efficient and resilient performance in the face of accidents, failures, natural disasters, or deliberate attacks.



John C. Doyle received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1977 and the Ph.D. degree in mathematics from the University of California, Berkeley, in 1984.

He is currently the John G. Braun Professor of Control and Dynamical Systems, Electrical Engineering, and Bioengineering with the California Institute of Technology, Pasadena.

Dr. Doyle received the IEEE Control Systems Field Award, the IEEE Centennial Outstanding Young Engineer Award, the IEEE Hickernell Award, the American Automatic Control Council (AACC) Eckman Award, and the Bernard Friedman Award. He was also a co-recipient of the IEEE Baker Prize, the AACC Schuck Award, the IEEE Axelby Award (twice), and the Sigcomm Student Award. He has held national and world records and championships in various sports.