Lessons from "A first-principles approach to understanding the Internet's router-level topology"

David L. Alderson
Naval Postgraduate School, USA
dlalders@nps.edu

John C. Doyle
Caltech, USA
doyle@caltech.edu

Walter Willinger
NIKSUN, Inc., USA
wwillinger@niksun.com

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Our main purpose for this editorial is to reiterate the main message that we tried to convey in our SIGCOMM'04 paper but that got largely lost in all the hype surrounding the use of scale-free network models throughout the sciences in the last two decades. That message was that because of (1) the Internet's highly-engineered architecture, (2) a thorough understanding of its component technologies, and (3) the availability of extensive (but typically noisy) measurements, this complex man-made system affords unique opportunities to unambiguously resolve most claims about its properties, structure, and functionality. In the process, we point out the fallacy of popular approaches that consider complex systems such as the Internet from the perspective of *disorganized complexity* and argue for renewed efforts and increased focus on advancing an "architecture first" view with its emphasis on studying the *organized complexity* of systems such as the Internet.

CCS CONCEPTS

• **Networks → Network Design Principles**; **Topology analysis and generation**;

KEYWORDS

Network topology, Network design principles

1 INTRODUCTION

Our SIGCOMM'04 paper on "A first-principle approach to understanding the Internet's router-level topology" [1] appeared at a time when the general excitement within the scientific community about the recent discovery of scale-free networks and their claimed universality had reached a fever pitch. In the particular case of the Internet, that discovery started in a line of research that portrayed this engineered system, including its router-level topology as "scale-free" (SF), with a central "hub-like" structure that makes the network simultaneously robust to random losses of nodes yet fragile to targeted attacks on the highly connected nodes or "hubs" in its core [2–4]. This combination of error tolerance and attack vulnerability, a tell-tale sign of SF structure, was subsequently referred to and popularized as the "Achilles' heel" of the Internet [3, 5], a property that reportedly went unnoticed by the engineers responsible for designing that very network.

Given that (1) SF methods are quite general and do not depend on any details of Internet technology, economics, or engineering, and (2) there were already decades of research on the structure and properties of the Internet, the broad appeal of that reported surprising discovery was understandable. At the same time, the SF approach as a whole and its resulting Achilles' heel claim also caused significant confusion among the networking community in general and the Internet community in particular. This confusion motivated closer scrutiny of this widely popular line of research that fascinated researchers across the different fields of science.

At a high level, this commentary serves as a reminder of an important message that we tried to convey in [1] but that got largely lost in all the ensuing SF-related "noise." That message was that because of the Internet's engineered architecture, a thorough understanding of its component technologies, and the availability of extensive (but typically noisy) measurements, this complex man-made system affords unique opportunities to unambiguously resolve most claims about its properties, structure, and functionality. In particular, the ability to rigorously validate most Internet-specific theories and models and their resulting claims creates the potential for any existing confusion or apparent controversies about the Internet to be settled once and for all, and our SIGCOMM'04 paper describes such an evaluation effort where the focus is on the theory of SF networks and its ensuing models and claims for the Internet.

More specifically, our intentions with this editorial are (i) to review the basic arguments that motivated the claims for the SF-view of the router-level Internet (along with a simple explanation why these are false) and (ii) to revisit our alternate perspective and arguments for a minimal model of the router-level Internet, one that explicitly includes fundamental design details and physical constraints that combine to ensure the intended functionality of the designed-for structure. We also comment on the impacts of this work for the Internet networking community (largely a success) and for the broader network science community (largely a failure). We conclude with a discussion of the lessons learned from both the success and failure of our work and the need for renewed efforts and increased focus on studying the organized complexity of the Internet.

2 SUMMARY OF ARGUMENTS AND RESULTS

This is a story in two acts. The first act involves arguments in favor of a SF perspective of the router-level Internet, and why these are false. The second act involves attempts to develop a minimal, yet explanatory perspective on the drivers of router-level network structure. We consider each in turn.

2.1 A Scale-Free Internet?

The starting point for the explosive growth in SF network models was the observation that a variety of naturally occurring and engineered systems exhibit power laws in the distribution of their connections [2–5]. Because this extremely skewed distribution is so different from the typical Poissonian distribution found in classic Erdös-Renyi random graphs, the "discovery" of power laws motivated an investigation of mechanisms that can give rise to it. One such mechanism is *preferential attachment*: a probabilistic growth process in which nodes are added one at a time and where each new node is more likely to connect to an existing node with more connections. This leads to a rich-get-richer process that gives rise to a *scale-free* distribution of node connections. It also yields a graph structure where the high-connectivity nodes (commonly called "hubs" in the SF literature) are crucial to maintaining the overall connectivity of the network; targeting these hubs can literally fragment the network.

Investigation into the connectivity of the Internet at different layers also showed signs of power laws [6]; this led to claims that the Internet also followed SF structure and also that its most salient features were the result of this connectivity structure. For the router-level Internet, this implied that error tolerance (i.e., minimal impact from loss of an arbitrary router) could be explained because most nodes in a SF network have little effect on the overall connectivity. However, it also suggested that targeting the high-connectivity routers could disconnect the router-level Internet, a vulnerability that got significant attention in the scientific and mainstream media because these studies showed "the removal of just a few key hubs from the Internet splintered the system into tiny groups of hopelessly isolated routers" [5].

Figure 1 summarizes the logical argument in favor of a scale-free view of the router-level Internet. Preferential attachment in network growth leads to power laws in network connectivity and also yields high connectivity hubs that make overall network connectivity vulnerable to attack. The argument is that the presence of power laws in the connectivity of the router-level Internet therefore implies the existence of a previously unknown vulnerability to attacks on the most highly connected routers.

However, Figure 1 also illustrates the logical flaw in this argument. Specifically, the presence of power laws in network connectivity does not necessarily suggest preferential attachment as the growth mechanism at work. In fact, there are many different mechanisms that can give rise to power laws in network connectivity, and preferential attachment is only one of them [7–9], albeit a well-known one that has been "rediscovered" multiple times in the past [10, 11]. Perhaps the most compelling image in our SIGCOMM'04 paper and subsequent work was [1, Figure 6] that showed five networks, each with the same number of nodes, same number of links, and same degree sequence, but each having obviously different structural properties. In particular, whereas SF networks resulting from preferential attachment had high-connectivity hubs in the center of the network, where they are essential to maintain connectivity (and become the characteristic Achilles' heel vulnerability), the figure illustrated other 'equivalent' networks where the high-connectivity nodes are at the network periphery such that their loss would be localized (and they would not pose such a threat). This, on its own, seemed to be enough to convey the logical fallacy in the SF story, as outlined in Figure 1.

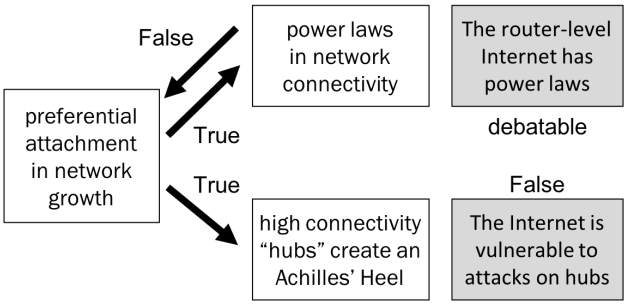

Figure 1: Basic Argument (and Logical Flaw) in the Scale-Free Story for the Internet. The basic claim was that a scale-free distribution in the connectivity of the Internet indicates the presence of high connectivity "hubs" that pose an overlooked vulnerability in the network. It is true that preferential attachment in network growth leads to power laws in network connectivity and also yields high connectivity hubs that create a vulnerability in the network if targeted. However, the converse is not true: the presence of a power law in network connectivity does NOT imply preferential growth in network formation (and therefore does not necessarily imply vulnerable hubs). The presence of power laws in network connectivity for the Internet is perhaps debatable, but the claim that the router-level Internet is vulnerable to attacks on its high degree hubs is false.

It is worth noting that the issue of power laws in Internet topology is a subject of its own debate, partly because of the idiosyncrasies and imprecision involved in large-scale network measurement studies [12, 13] and partly because of the ad-hoc nature of commonly-used statistical techniques for inferring power law-like distributions from noisy data [14]. This controversy regarding the presence of power laws in the topology of the Internet created additional confusion, but it ultimately does not change the analysis in Figure 1. In short, whether or not there are power laws in the topology of the router-level Internet, claims of scale-free vulnerability of high-connectivity hubs are not substantiated and inherently flawed.

2.2 An Engineering Perspective

From an engineering design perspective, preferential attachment was never a plausible explanation of topology formation for the router-level Internet, however, part of the confusion regarding the SF story was the absence of an alternative perspective. The second part of the story in our SIGCOMM'04 paper was to propose function, technology, and economics (as opposed to simple connectivity patterns) as the key drivers of structure for the router-level Internet.

In a nutshell, our "first-principles approach" describes an effort at identifying some minimal functional requirements and physical constraints needed to develop simple models of the Internet's router-level topology that are illustrative, representative, insightful, and consistent with engineering reality. Despite the large number of factors that undoubtedly play a role in the minutiae of real-world

router-level topology design, we were able to identify (a) hard technological constraints in the form of physical limits on router capacity, connectivity, and link bandwidth and (b) economics-driven considerations in terms of anticipated end user demands and expected network performance. These are key design contributors to the type of router-level structures that are possible and make sense from an engineering perspective.

This engineering-based approach has several advantages. It respects the highly designed nature of the network, it reflects the engineering intuition that exists about a great many of its parts, and it is fully consistent with readily available but noisy measurements of router-level topology. In essence, a limit on the throughput capacity of a router means that a router can have either a small number of high-bandwidth connections or a large number of small-bandwidth connections. Because of this fundamental tradeoff, high-connectivity hubs will be performance bottlenecks if they are used as key interconnection points in the center of a network. When we calculated the throughput capacity of different topology designs, we observed that SF model networks had throughput numbers that were two orders of magnitude smaller than others we had designed to be 'heuristically optimal' (also called a *heuristically optimal topology* or HOT model). Our designed networks used high-connectivity nodes at the edge of the network (for traffic aggregation purposes), but had a sparsely connected core network structure with high-speed connections. This design is consistent with (and inspired by) the design of real router-level networks that the various Internet service providers operated in the early 2000s.

Our approach ultimately resulted in a new means for rigorously discerning the key differences between different topology structures. By visualizing them in a 2-dimensional space where the coordinates refer to network performance (i.e., throughput capacity) and network likelihood (from a random graph construction perspective) [1, Figure 8], the differences between topologies that were generated by preferential attachment or other random graph generation techniques (designed to achieve a power law distribution in connectivity pattern) and our engineering-based approach became apparent and were telling.

One of the key insights for our work was that if we want models that are truly explanatory and not merely descriptive, we need to broaden our perspective beyond that of simple graphs to include the drivers of this network structure [1, 15]. While the overall construction of the router-level Internet is decentralized and somewhat ad hoc, it is much more than the outcome of a random process. By choosing the language of constrained optimization (i.e., the objectives and constraints that guide this construction) instead of the language of random graphs, we transform network modeling into a problem driven by domain-specific graph annotation not mere graph connectivity, and we move beyond a simple exercise in data fitting. As described in [16]: "by focusing on optimization as a modeling process, not a specific modeling outcome (i.e., the solution to any one optimization problem), one can systematically study how particular objectives and constraints shape the large-scale structure and behavior of complex networks. With this perspective, optimization-based reverse-engineering approaches such as HOT serve best as a conceptual framework (or a modeling methodology), not a specific model for complex networks."

A second key lesson for modeling Internet structure was the simple reminder that there is no such thing as a unique "Internet topology" due to the many different layers of virtualization and abstraction, each of which is shaped by different objectives, constraints, and circumstances. Moreover, due to the inherent difficulties with Internet topology measurement, one cannot take measured topology data at face value, particularly when that data has high variability [14].

3 IMPACTS

This work represented the start of a line of research that had impacts on the Internet research community and the broader study of complex networks.

3.1 Direct Impacts

Honored with the Best Student Paper Prize Award at SIGCOMM'04, the paper's impact on the networking community was immediate. On the one hand, network engineers and operators were mostly surprised by a paper that, in their minds, stated the "obvious" and refuted the Internet Achilles' heel myth, but they generally appreciated the documented first-principles approach for its ability to identify and explain the cause-effect relationships present in the very systems that they design and operate. On the other hand, Internet researchers basically abandoned the SF approach for modeling the physical Internet, and those still convinced of the utility of SF networks shifted their focus to the higher layers of the Internet's architecture, where looking for SF structure in virtual topologies such as the Facebook friendship graph or the Twitter social network has become a minor industry in the last decade.

The few times the SF approach tried to make a comeback and reclaim its relevance for physical Internet structures such as datacenter topologies [17], the attempts quickly faded because technology constraints and economic considerations still restrict intra-datacenter connectivity even when that connectivity could otherwise be governed by randomness [18, 19]. In receiving the 2016 ACM SIGCOMM Test of Time Award, the paper was singled out for (i) *questioning the prevailing work on scale-free graph structure for network topologies that incorrectly speculated an "Achilles' heel" for the Internet and providing instead a methodologically sound basis to explain the observed structure of Internet topologies*, and (ii) *bringing a greater degree of rigor in network topology research and evaluation and informing the community of potential pitfalls in using black-box network models without a clear understanding of underlying structural effects in network design.*

3.2 Broader Impacts

Given the attention that SF networks received from the mathematical community in general and the graph theorists in particular, following our SIGCOMM'04 paper, we also wrote a number of papers that specifically targeted those audiences. Our motivation for doing so was two-fold. For one, by aiming at a more mathematically-inclined readership that can be assumed to be familiar with the basics of Internet technologies, we used the Internet example to demonstrate why the existing theory of SF networks has inherent inconsistencies which, in turn, lead to commonly-cited claims about the Internet that are verifiably incorrect [20]. At the same time, we

also wanted to demonstrate the great potential that a large-scale complex system such as the Internet has for developing a new mathematical theory that is scientifically more challenging, more relevant in practice, and ultimately more rewarding because of the new insights it provides [13]. While the papers were in general well-received, with [13] selected for inclusion in the *2010 Princeton Anthology of Best Writing on Mathematics* [21], our attempts at convincing the broader scientific community, particularly physicists and "network scientists," of the enormous benefits of enriching their models (that are largely governed by randomness) with "just the right amount" of domain knowledge have been mostly futile.

3.3 Catching the Network Science "Bug"

At the time of our SIGCOMM'04 paper, the debate surrounding the role of SF networks in the Internet seemed to be peaking, however in retrospect it was relatively little when compared to the explosive growth in popularity that the field of network science has experienced in the subsequent 15 years. In 2005, the National Research Council (NRC) released a report titled "Network Science" [22] that served as a rallying cry for increased investment for research and education in the tools and techniques of applied graph theory and statistical mechanics as applied to a wide range of complex networks. The growth in funding, publications, and citations since then have been breathtaking (see [16] for a review of this early history). It is worth noting that the original claims in the popular science literature for an Achilles' heel vulnerability of the router-level Internet [3–5] have received over 10,000 citations collectively (according to Google Scholar as of July 2019), and have never been retracted despite what most in the SIGCOMM community consider "a closed case." Our attempts to tell this story to the broader scientific community (e.g., [23]) have received limited success in reaching this audience, and anecdotally our experience is that the myth of the Achilles' heel vulnerability for the router-level Internet remains prevalent among the network science community.

Despite the sometimes sensational claims made about the Internet, network science as a discipline has managed to operate largely independently from the network research and operator community. This disconnection was established early in its history, as noted by the NRC [22]: "network science is distinct from both network technology and network research: It is characterized by the discovery mode of science rather than the invention mode of technology and engineering." Again, rather than looking to domain-specific details for an understanding of structure and function, the starting point is most often one of random graphs, with models that emphasize connectivity patterns over system performance. For example, a prominent 2008 retrospective "Scale-Free Networks: A Decade and Beyond" [24] opens with the statement: "For decades, we tacitly assumed that the components of such complex systems as the cell, the society, or the Internet are randomly wired together"—a claim that many in the SIGCOMM community would recognize as nonsensical.

4 BARRIERS TO PROGRESS

There are a variety of complicating factors that serve as barriers to progress in developing an Internet-inspired mathematical theory of organized complexity that advances the view of "architecture first" [46] to support design and development of future systems.

4.1 Power Laws

Power laws remain a source of fascination and confusion within the broader scientific community. Borrowing from our previous work [25] (but with reference numbers updated to align here),

> In the [new sciences of complex networks, NSCN] and modern physics literature, power laws are viewed as "signatures" of specific mechanisms, namely, critical phase transitions [26] and preferential growth [27], 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 frequencies plotted on log-log scale, make it easy to "discover" power laws where none exists or to mischaracterize their most salient features (see [20, Sec. 2.1]). However, a broader view of high variability and power laws reveals a long and rich history outside physics [28–30].

Ironically, many of the high-profile results celebrating power-laws and the need for special models to explain them actually arise as artifacts of statistical errors. These are not limited to SF networks for technological systems, but also include self-organized criticality (SOC) for wildfires (see original claims in [31] and a rebuttal in [32]) and edge-of-chaos for heart rate variability (see [33] for a discussion of claims and [34] for a rebuttal). In our work, we have emphasized mathematical, statistical, and data-analytic arguments advanced decades ago by B. Mandelbrot [36] that suggest that power-law (or more accurately, *scaling*) distributions should be no more suprising than Gaussians when dealing with high variance data [35]. In this sense, power laws are "full of sound and fury, signifying nothing" [42]. Nonetheless, the "rediscovery" of power laws in SF networks remains a topic of active research [37, 38] and still attracts attention within the high-impact scientific literature [39].

4.2 Architecture: More Than Connectivity

A second source of confusion within the broader scientific community comes from wildly different notions of *architecture*. Within the network science community, the term 'architecture' is used synonymously with 'connectivity'. As clearly articulated in a highly cited review [40]: "By definition, complex networks are networks with more complex architectures than classical random graphs with their 'simple' Poissonian distributions of connections." This is in large part due to the fact that when abstracting a complex system down to a simple graph, there is little left to study other than its connections, perhaps with simple annotations.

In contrast, the engineering perspective of architecture tends to focus on the *protocols* (persistent rules of interaction) over specific *modules* (implementations that obey protocols and can change) [25]. Here, the Internet shares remarkable similarities with biology, where there are plausibly heavy tails and even power laws all over. Similarly, where done correctly, biology provides a parallel story of highly structured systems which "invite" simple abstractions into a great many completely different "graphs" although graphs are not good models in any of these cases. Similar to the SF story for the

router-level Internet, many of the arguments in favor of SF biology, neurology, medicine, or even ecology turn out to be specious when examined from an architectural perspective that emphasizes protocols and modules over simple connectivity [43–45].

4.3 On the Role of Layering

A key architectural characteristic that is often under-appreciated and/or misunderstood by the broader scientific community is the role of *layering* in highly engineered or highly evolved systems. By design, a main feature of layering is that each layer exists for a purpose (i.e., provides its own functionality), and builds on the functional layers below it. As a consequence, each layer is typically shaped by different objectives and constraints, and each layer can therefore have its own connectivity. While the connectivity at a given layer is largely independent of the other layers, its physical manifestation at the lowest layer(s) turns strictly logical/virtual at the higher layers, with important implications for interpreting the resulting constructs and assessing their relevance, individually or as part of the system as a whole.

Despite the combination of potentially wildly different layer-specific connectivity structures in the Internet, the system as a whole is much more than random wiring with minimal tuning (e.g., to match one or more power law distributions at the different layers). However, to recognize this, it is important to have answers to questions such as *Why does one use layering as a design principle?* and *How does one layer?* To this end, recent work that is similar to our SIGCOMM'04 paper in the sense that it advances an architecture-centric approach and emphasizes the importance of functionality allocation (i.e., deciding on which layer does what) has established "layering as optimization decomposition" as a common "language" or top-down method for designing layered protocol stacks from first principles (see, for example, [46] and references therein). In short, by modeling an overall communication network in terms of a generalized network utility maximization problem, each layer gets mapped to a decomposed subproblem, and the interfaces among the different layers are quantified as functions of the optimization variables coordinating the subproblems. It is this combination of horizontal decomposition into distributed computation and vertical decomposition into functional modules (e.g., congestion control, routing) that formalizes the notions of "networks as optimizers" and "layering as decomposition" and provides answers to the *why*, *how* and *what* questions that arise in the context of existing or newly-proposed layered protocol stacks.

4.4 On the Role of Virtualization

Another important ingredient for understanding modern-day computing and communication systems is *virtualization*. Just as the Internet's layered TCP/IP stack enables the plugging-in of new applications above and new link technologies below it's narrow waist (i.e., "IP over everything, and everything over IP"), with virtualization, this plug-and-play capabilities become at the same time more ubiquitous and less visible. With the separation of compute resources (e.g., CPU, memory) or requests for services from the underlying physical delivery of those services as their main goals, virtualization techniques have been applied across the board, from memory to laptop or server hardware, to operating systems (OS) and applications, and to storage and entire networks.

This "virtualization everywhere" has created *virtual infrastructures* that provide a layer of abstraction between computing, storage and networking hardware on the one hand, and the applications running on them on the other hand. As a result, OSes and applications can be managed as a single unit by encapsulating them into virtual machines (VM) which in turn can then be provisioned to any system. While this hardware-independence provides great efficiency and flexibility, it further obscures from the user the underlying mechanisms and details and provides a user experience that is largely unchanged by the use of such highly architected virtualization. Put differently, it is typically only under failure scenarios and network impairments that users see signs of a system's reliance on ubiquitous and purposefully-applied virtualization.

Note that just as the concepts of layering and layered protocol stacks are foreign to the network science view of networks, so is the idea of virtualization. In this sense, the network science view of networks that has turned the study of abstract graphs (extracted from measurements collected from today's Internet with its ubiquitous use of highly engineered virtualization) into a minor industry necessarily boils down to finding hidden simplicity that "emerges" from enormous complexity. However, the applications of this approach to the Internet demonstrate at once both the irony and the fallacy of this view.

4.5 The Effect of Layering and Virtualization

One critical feature of architectures that emphasize layering and virtualization is that they allow the system to hide the details of one layer from another. On the one hand, this feature is essential from an engineering perspective because it greatly facilitates the interoperability of diverse technologies. For example, the "narrow waist" of the Internet's TCP/IP protocol stack allows for a diversity of applications to run on a diversity of physical link technologies. Moreover, it allows new applications or link technologies to be readily adopted as long as they "speak" TCP/IP. At the same time, there is also an irony to this hiding of details behind layers in the sense that it not only invites the use of abstract constructs such as random graphs but also makes them (and hence the network science view) seem plausible. More cynically, the network science view likes abstractions in the form of random graphs, in part, *because* they obscure layered architectures to the point where they are not even part of the existing vocabulary and can therefore be conveniently and completely ignored.

An important consequence of this outward appearance of simplicity that layering and virtualization convey about the underlying system is that it both invites "discoveries" of system vulnerabilities that, upon further scrutiny, turn out to be specious, and obscures the detection of real-world vulnerabilities that can have a devastating impact on the system as a whole. To illustrate, consider the widely-studied connectivity structure known as the Web graph where nodes represent websites and edges represent hyperlinks. This logical graph construct manifests itself at the application layer of the TCP/IP protocol stack, and measured instances of this Web graph have persistently shown power law properties in their connectivity [47, 48]. From the SF perspective, the Web's Achilles' heel

is its extreme vulnerability to attacks that target its most highly-connected websites. However, eliminating high-degree nodes from the Web graph would require bringing down websites that, because of their popularity, have been replicated on multiple servers located across the Internet, with the different servers being furthermore virtualized for greater efficiency and cost savings. Recognizing these and similar difficulties and complexities reveals at once the specious nature of this claimed vulnerability but requires a basic understanding of how the lower layers of the TCP/IP stack enable the required geo-replication of content and virtualization of essential network infrastructure. Both of these functionalities provided by the lower layers contribute to a well-performing Web but the details of their implementation are in general invisible to the users of the Web. Moreover, even if theoretically feasible, bringing down a popular website may inconvenience a set of users but has little if any impact on the overall functionality of the Internet.

In fact, from an engineering perspective that is fully aware of the Internet's layered TCP/IP protocol stack and its highly architected virtualization, one of the greatest vulnerabilities for the Internet is not attacks on its Web graph's high-degree logical nodes but hijacking any of the protocols that are critical for a functioning Web (or any other popular applications). Consider, for example, the Internet's control plane where the BGP protocol is responsible for routing the requests for accessing the various websites and viewing or downloading their content across the Internet. Irrespective of whether BGP is hijacked for malicious purposes or by accident (i.e., router misconfiguration), its effect is that popular websites such as www.google.com can be brought down [49]. Another well-known type of hijacking attacks are DNS-based amplification attacks. They can render entire websites inaccessible by hijacking the DNS protocol and using it as weapon-of-choice to distributed denial-of-service (DDoS) the website of interest [50, 51].

Note that both types of vulnerabilities rely on the full functioning of all the lower layers to achieve maximal impact or "success" on the application layer (e.g., delivering malware to as many end hosts as possible, infiltrating control systems of cybercritical infrastructure). Leveraging the key functionalities of the very system they try to attack and ensuring their continued availability during the attack to achieve maximum damage is not only a hallmark of the Internet's most serious attacks but has also striking similarities with some of the most devastating attack scenarios encountered in biological systems (e.g., cancer). Yet, network science with its focus on random graphs, possibly in conjunction with random dynamic processes over such graphs, not only lacks the taxonomy to describe such attacks but is incapable of even perceiving such scenarios due to its inherently limited view of what constitutes a complex system.

5 FROM "FIRST-PRINCIPLES" TO "ARCHITECTURE FIRST"

Despite the success of the first-principles approach to understanding the Internet's router-level structure advanced in [1], our SIGCOMM'04 paper has largely failed to impress upon the science community in general and the networking research community in particular the full potential of an "architecture first" view of the Internet beyond its physical infrastructure. We argue that science will be better served by pursuing an "architecture first" approach to understanding highly architected virtualization that mimics the efforts described in [1] and will be ultimately more interesting, rewarding, insightful, and relevant than further musings about power laws, error tolerance, or attack vulnerabilities.

5.1 On Laws, Layers, and Levels

For networking researchers, the most familiar layered architecture is arguably the layering of application (app)/operating system (OS)/hardware (HW) in our phones and computers, and their extensions to networks. Within the HW layer, the memory hierarchy can also be thought of as layers of tape/disk/RAM/cache/register from high, big, cheap, and slow to low, small, expensive, and fast. Memory also illustrates "levels" with each layer being implemented in lower level components, such as transistors. Of course, the apps and OS layers have sublayers and levels, and so on. Layered architectures are everywhere in technology and biology but virtualization means that details are often hidden from nonexperts. We adopt an admittedly greatly oversimplified view to highlight broad similarities and differences, and encourage experts to add additional details.

An ideal memory would be fast, large, and cheap, but no such individual component exists. This tradeoff can be loosely thought of as a *law*, but one that depends on available technology in addition to more fundamental laws of physics. Here we use "law" and tradeoff broadly (and somewhat loosely) to describe constraints, including legal or regulatory, on what components are available. What virtual memories do is exploit the diversity in the memory hierarchy to create a *diversity-enabled sweet spot (DeSS)* that is fast, large, and cheap despite having no such individual parts, with the fragility that certain rare access sequences could be very slow.

Such tradeoffs are pervasive throughout Internet technologies. For example, virtual memory is cheap and plenty, but accessing it takes time; CPU resources on modern programmable devices are scarce but allow for certain compute tasks to be performed at line-rate; and cloud resources are essentially unlimited but it takes time to transfer all the necessary training data and retrieving the required learning model. Despite the highly domain specific nature of these laws, successful architectures to cope with them are remarkably universal in their use of layers and levels to create a DeSS. Indeed, our view is that the most important function of a complex architectures is to create a DeSS where simpler uses of components cannot.

5.2 From Networking to Biology ...

Our cells, organs, bodies and brains also illustrate layered architectures in myriad ways. Our central nervous system has the brain/brainstem/spine as its minimal layers, with neurons and other cells as the main lower level. What the brain/brainstem/spine and apps/OS/HW have in common is a DeSS for flexibility and speed from highly constrained parts. Maximum speed requires special purpose hardware in our spines or computers, but the greatest flexibility is in our apps and brains layer, where apps, memes, and ideas are highly swappable via shared languages. Most hidden but essential are the OS and brainstem to virtualize our fast hardware. Experts estimate that conscious thought is limited to 100 bits/sec bandwidth, whereas the optic nerve alone runs at 10Mbits/sec [52]. Most of this enormous amount of information is not discarded but

is used unconsciously and automatically to drive sensorimotor control, and direct the attention of that tiny and precious conscious thought.

Our cells have a familiar DNA/RNA/protein layering that also emphasizes the molecules that implement the layers, with a universally shared genome by incredibly diverse cell types which are organized to create remarkable DeSS. A major difference is that cells and bodies have their own internal supply chains to make new proteins and cells, whereas our computer hardware is produced "out of band." A similarity is despite the importance of "code" in DNA and software, hardware and proteins are necessary for the code to function. And some terminally differentiated cells in our blood, eyes, and skin jettison their DNA, RNA, and associated molecular machinery. As in our technologies, specialized hardware solutions provide superior performance but with a loss of flexibility.

5.3 ... and back to Networking

One area where the proposed "architecture first" view of the Internet promises to have significant impact on future Internet research concerns the long-standing problem of network automation and the ultimate challenge of developing autonomous or "self-driving" networks [53]. In fact, given the recent advances in programmable, protocol-independent data planes (with languages for programming them) and the emergence of scalable platforms for processing distributed streaming data (capable of leveraging the latest Artificial Intelligence/Machine Learning tools), networking researchers have finally technologies at their disposal that have long been viewed as pre-requisites for automating networking tasks at scale; that is, executing hundreds or thousands of highly diverse network automation tasks concurrently and as fast and accurately as possible despite the uncertainties in the environment (e.g., traffic load, application mix, failure scenarios).

When considering, for example, automating a network management task such as detecting and mitigating DDoS attacks in real time, programmable data planes afford a number of different choices when it comes to allocating the compute resources (e.g., CPU, memory) and communication resources (e.g., connectivity, bandwidth) that are required to perform the task-specific telemetry (e.g., sensing), in-network computation (e.g., automatic inference), and forwarding (e.g., actuating). The important observation is that each of these choices comes with its own hard tradeoff. In this case, the tradeoff is concerned with balancing hardware component-induced speed vs. hardware component-specific accuracy. When combined, the ensuing component-specific speed-accuracy tradeoffs result in system-level speed-accuracy tradeoff that determines how well a given network automation task can be performed with no human operator in the loop. While the idea of exploiting the diversity in available compute resources and programmability capabilities among the different hardware components of a network to achieve the "best case" scenario (i.e., performing the task of interest both fast and with high accuracy) is already being explored (e.g., see [54] and references therein), how to get the network to recognize and then operate at such a DeSS remains an open problem, even for the case of a single (sufficiently complex) network automation task.

6 LOOKING AHEAD: THE GOOD, THE BAD, AND THE UGLY

From a strictly technological perspective, an "architecture first" approach to realizing the vision of autonomous or self-driving networks can be expected to focus on new architectural designs that will facilitate DeSS-seeking network automation at scale. Here, "looking over the fence" shows that biology has already worked out an architecture for human sensorimotor control that exhibits remarkable speed and accuracy despite being implemented by imperfect hardware (e.g., nerves and muscle components with their own speed-accuracy tradeoffs). A recently developed theoretical framework [55, 56] provides a first explanation how a highly effective layered control architecture in conjunction with diversity among the available hardware components enables fast and accurate system-level performance to be achieved using slow or inaccurate hardware components. By building on this new theoretical framework and adjusting it to network-specific conditions, it may be possible to provide a scientifically-sound foundation for developing new technologies in support of autonomous networks that can serve as a practical guide for engineers in their DeSS-driven quest to build fast and accurate systems from imperfect hardware.

At the same time, biology also has a warning for networking as far as the vulnerability of newly developed technologies is concerned. For example, like genes, apps and memes are easily mutated and even swapped, enabling software and genetic engineering as well as education and culture. However, a strikingly common fragility is that parasites can also "swap in" their code. This infectious hijacking leaves the architecture largely intact but switches the function to benefit the parasite. Bacteria use gene swapping (called "horizontal gene transfer") to rapidly share antibiotic resistance genes that hijack other elements of the architecture (e.g., G-protein coupled signal transduction protocols). The eukaryotic malaria parasite has a complex lifestyle involving infectious hijacking of mosquitos and mammals. Experts claim that nearly half of the 100B humans who have ever lived were killed by malaria, with modern treatments reducing that rate [57]. However, other so-called "superbugs" are evolving with swappable code that defeats available medical treatments.

However, perhaps the most ominous and dire lesson that biology can teach networking concerns societal issues such as humanity's affinity for sharable "memes" containing ideas that are attractive and highly infectious, but false and dangerous. Assuming that social scientists are correct in arguing that the very falsehood of memes facilitates their transmission (because their mostly deleterious effects on their hosts make them honest signals of group membership), advanced technologies in support of ever more powerful communication and computing become enablers of future scenarios where "bad memes" may make malaria's history seem relatively tame. It is thus of critical importance for networking researchers to be held accountable for their design of new technologies by requiring them to examine in detail their technologies' exposure to infectious hijacking at all layers of interest, especially the non-standard "layers" 8 (i.e., individuals/users), 9 (i.e., organizations) and 10 (i.e., governments/legal entities). Unfortunately (and ironically), this call-to-action for networking researchers can be expected to be undermined by the fact that science itself is not immune to such

infectious hijacking, as the enormous and continued popularity of scale-free "memes" illustrates.

REFERENCES

[1] L. Li, D. Alderson, W. Willinger, and J.C. Doyle. A First-principles Approach to Understanding the Internet's Router-level Topology. *ACM SIGCOMM Comput. Commun. Rev. (Proc. ACM SIGCOMM'04)*, 34(4):3–14, 2004.

[2] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[3] R. Albert, H. Jeong, and A.-L. Barabasi. Attack and error tolerance of complex networks. *Nature*, Vol. 406, 378–382, 2000.

[4] S.-H. Yook, H. Jeong, and A.-L. Barabasi. Modeling the Internet's large-scale topology *Proc. National Academy of Sciences*, 99(21):13382–13386, 2002.

[5] A.-L. Barabasi and E. Bonabeau, "Scale-Free Networks," *Scientific American*, May 2003, 63-69.

[6] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology, *ACM Comp. Comm. Review*, 29(4), 1999.

[7] M.E.J. Newman. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 2005.

[8] A. Clauset, C.R. Shalizi, and M.E.J. Newman. "Power-law distributions in empirical data." *SIAM Review* 51.4, 2009, 661-703.

[9] M.P.H. Stumpf and M.A. Porter. "Critical truths about power laws." *Science* 335.6069, 2012, 665-666.

[10] G. U. Yule. A mathematical theory of evolution, based on the conclusions of Dr J. C. Willis, F.R.S. *Phil. Trans. B*, 213(21), 1924.

[11] H. A. Simon. On a class of skew distribution functions. *Biometrika* 42(3-4), 1955, 425-440.

[12] Willinger, W. and Roughan, M. Internet topology research redux. *ACM SIGCOMM eBook: Recent Advances in Networking*, Haddadi, H. and Bonaventure, O. (Eds.), 2013.

[13] W. Willinger, D. Alderson, and J. C. Doyle. Mathematics and the Internet: A Source of Enormous Confusion and Great Potential. *Not. of the AMS*, 56(5):586–599, 2009.

[14] Willinger, W., Alderson, D., Doyle, J., and Li, L., 2004. A pragmatic approach to dealing with high variability in network measurements. Proc. ACM SIGCOMM Internet Measurement Conference 2004, Taormina, Sicily, Italy, Oct. 25-27, 2004.

[15] L. Li, D. Alderson, W. Willinger, and J.C. Doyle. Understanding Internet Topology: Principles, Models, and Validation. *IEEE/ACM Trans. Networking*, 13(6):1205–1218, 2005.

[16] D.L. Alderson, "Catching the 'Network Science' Bug: Insight and Opportunity for the Operations Researcher," *Operations Research* 56(5): 1047-1065, 2008.

[17] L. Gyarmati and T.A. Trinh. Scafida: A Scale-free Network Inspired Data Center Architecture. *ACM SIGCOMM Comput. Commun. Rev.*, 40(5):4–12, 2010.

[18] A. Singla, C.-Y Hong, L. Popa, and P.B. Godfrey. Jellyfish: Networking data centers randomly. *Proc. USENIX NSDI'12*, San Jose, CA, 2012.

[19] A. Singla, P.B. Godfrey, and A. Kolla. High Throughput Data Center Topology Design. *Proc. USENIX NSDI'14*, Seattle, WA, 2014.

[20] L. Li, D. Alderson, J. Doyle, and W. Willinger. Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications. *Internet Mathematics*, 2(4):431–523, 2005.

[21] The Best Writing on Mathematics 2010, M. Pitici (Editor). Princeton University Press, Princeton, NJ, 109–133, 2011.

[22] National Research Council, Committee on Network Science for Future Army Applications. 2005. *Network Science.* The National Academies Press.

[23] J.C. Doyle, D.L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The "robust yet fragile" nature of the Internet. *Proc. National Academy of Sciences*, 102(41):14497–14502, 2005.

[24] A.-L. Barabasi. Scale-Free Networks: A Decade And Beyond. *Science* 325, 24 July 2009, 412.

[25] D.L. Alderson and J.C. Doyle. Contrasting Views of Complexity and Their Implications For Network-Centric Infrastructures. *IEEE Trans. Systems, Man, and Cybernetics - Part A*, 40(4):839–852, 2010.

[26] P. Bak. How Nature Works: The Science of Self-Organized Criticality. New York: Copernicus, 1996.

[27] A.-L. Barabási, Linked: The New Science of Networks. Cambridge, MA: Perseus, 2002.

[28] V. Pareto. Cours d'Economie Politique. Geneva, Switzerland: Droz, 1896.

[29] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Math.*, vol. 1, no. 2, pp. 226ÃŘ251, 2004.

[30] E. Keller. Revisiting "scale-free" networks. *Bioessays*, vol. 27, no. 10, pp. 1060-1068, Oct. 2005.

[31] Malamud, B.D., Morein, G. and Turcotte, D.L. Forest fires: an example of self-organized critical behavior. Science, 281(5384), pp.1840-1842, 1998.

[32] 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. 17 912âĂŞ17 917, Dec. 2005.

[33] Glass, L. Introduction to controversial topics in nonlinear science: Is the normal heart rate chaotic? CHAOS 19, 028501, 2009.

[34] Li, N., Cruz, J., Chien, C.S., Sojoudi, S., Recht, B., Stone, D., Csete, M., Bahmiller, D. and Doyle, J.C.. Robust efficiency and actuator saturation explain healthy heart rate control and variability. Proceedings of the National Academy of Sciences, 111(33), pp.E3476-E3485, 2014.

[35] W. Willinger, D. Alderson, L. Li, and J.C. Doyle. More "Normal" Than Normal: Scaling Distributions and Complex Systems. Proc. 2004 Winter Simulation Conference. Ingalls, Rossetti, Smith, and Peters, eds., 2004.

[36] B.B. Mandelbrot. Fractals and Scaling in Finance: Discontinuity, Concentration, Risk. New York: Springer-Verlag, 1997.

[37] Scale-free networks well done (2018) https://arxiv.org/abs/1811.02071

[38] Rare and everywhere: Perspectives on scale-free networks. (2019) https://www.nature.com/articles/s41467-019-09038-8

[39] Scale-free networks are rare (2019) https://www.nature.com/articles/s41467- 019-08746-5

[40] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes. Critical phenomena in complex networks. *Reviews of Modern Physics* 80, October-December 2008, p. 1275.

[41] W. Weaver. Science and complexity *Amer. Sci.,* vol. 36, pp. 536ÃŘ544, 1948.

[42] S. H. Strogatz. Romanesque networks. *Nature*, vol. 433, 2005.

[43] R. Tanaka. Scale-rich metabolic networks. *Phys Review Letters* 94 (16), 168101, 2005.

[44] R. Tanaka, T.M. Yi, J. Doyle. Some protein interaction data do not exhibit power law statistics. *FEBS Letters* 579 (23), 5140-5144, 2005.

[45] R. Tanaka, M. Csete, J. Doyle. Highly optimised global organisation of metabolic networks. *IEE Proceedings-Systems Biology* 152 (4), 179-184, 2005.

[46] M. Chiang, S.H. Low SH, A.R. Calderbank, J.C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE.* 95(1):255-312, 2007.

[47] A. Broder et al. Graph structure in the web. *Proc. WWW*, 2000.

[48] M. Newman. Networks (2nd Edition). Oxford University Press, 2018.

[49] D. Goodin. Google goes down after major BGP mishap routes traffic through China, 2018. https://arstechnica.com/information-technology/2018/11/major-bgp-mishap-takes-down-google-as-traffic-improperly-travels-to-china/

[50] KrebsOnSecurity hit with record DDoS, 2016. https://krebsonsecurity.com/tag/dns-amplification-attack/

[51] DDoS attack against Spamhaus was reportedly the largest in history, 2013. https://www.networkworld.com/article/2164810/ddos-attack-against-spamhaus-was-reportedly-the-largest-in-history.html

[52] M. E. Raichle. Two views of brain function. *Trends in Cognitive Sciences* 14(4):180–190, 2010.

[53] N. Feamster and J. Rexford. Why (and How) Networks Should Run Themselves. http://arxiv.org/abs/1710.11583, 2017.

[54] A. Gupta, R. Harrison, M. Canini, N. Feamster, J. Rexford, and W. Willinger. Sonata: Query-Driven Network Telemetry. ACM SIGCOMM, 2018.

[55] Y. Nakahira1, Q. Liu, T. J. Sejnowski, and J. C. Doyle. Fitts' law for speed-accuracy trade-off is a diversity sweet spot in sensorimotor control. *https://arxiv.org/pdf/1906.00905.pdf*, 2019.

[56] J. C. Doyle et al. Diversity sweet spots in layered architectures and speed-accuracy trade-offs in sensorimotor control. *Preprint*, 2019.

[57] S. Shah. The Fever: How Malaria Has Ruled Humankind for 500,000 Years. Sarah Crichton Books, 2010.