

Towards Re-architecting Today's Internet for Survivability

NSF Workshop Report

Fabián E. Bustamante
Northwestern U.
fabianb@cs.northwestern.edu

Walter Willinger
NIKSUN Inc.
wwillinger@niksun.com

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

John Doyle
Caltech
doyle@caltech.edu

Marwan Fayed
Cloudflare Research
marwan@cloudflare.com

Steven Low
Caltech
slow@caltech.edu

Stefan Savage
UCSD
ssavage@ucsd.edu

Henning Schulzrinne
Columbia U.
hgs@cs.columbia.edu

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 It has NOT been peer reviewed.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

On November 28-29, 2023, Northwestern University hosted a workshop titled "Towards Re-architecting Today's Internet for Survivability" in Evanston, Illinois, US. The goal of the workshop was to bring together a group of national and international experts to sketch and start implementing a transformative research agenda for solving one of our community's most challenging yet important tasks: the re-architecting of tomorrow's Internet for "survivability", ensuring that the network is able to fulfill its mission even in the presence of large-scale catastrophic events. This report provides a necessarily brief overview of two full days of active discussions.

CCS CONCEPTS

• **Networks** → **Public Internet**; **Network properties**; **Network architectures**;

KEYWORDS

Internet, Survivability, Resilience

1 INTRODUCTION

Over the past decades, the Internet has undergone a major change from being primarily a research-oriented network for academics to becoming a cyber-physical infrastructure critical for modern society in general and the global economy in particular. This transformation has occurred mainly by happenstance rather than by design and under the assumption that the current architecture that has ensured its robustness in the past would be sufficient to provide the robustness now expected from it.

We believe that this organically-grown architecture of today's Internet cannot live up to this new role humanity has assigned it or withstand the types of threats that it now faces.

Re-architecting today's Internet as critical infrastructure requires a new understanding of the architectural principles on which it should be based. It demands a reassessment of the possible scenarios that can challenge the network's basic functioning and the

threats that can arise due to the network's constant evolution. At the same time, it must explore paths for incremental deployment that embed the necessary incentives for adoption. Given the expected tight coupling of tomorrow's Internet with the emerging smart grid, the analysis of potential threats and any re-design to enhance survivability must consider both systems in parallel and inform each other's progress.

The success of such an ambitious effort depends on close collaborations among a broad and interdisciplinary team of scientists, including networking researchers, power/smart grid experts, economists, resilience engineers, and control systems researchers.

With the generous support of NSF, a group of us organized a workshop on November 28-29, 2023. The workshop, entitled "Towards Re-architecting Today's Internet for Survivability," aimed to bring together an initial group of national and international experts in a range of these areas to sketch and start implementing a transformative research agenda for solving one of our community's most challenging yet important tasks: the re-architecting tomorrow's Internet for "survivability," ensuring that the network is able to fulfill its mission even in the presence of large-scale catastrophic events [5].

The workshop run for two days. Given the variety of topics, the first day focused on creating a shared understanding of the space with overview talks by leaders in the different areas we have identified: *Power Grid and the Internet*, *Control Systems*, *Threats to Internet Survivability*, *Resilience Engineering*, and *Perspective from the Public and Private Sectors*. Building on this, the second day was dedicated to short talks in each area, following more or less the same structure, and brainstorming sessions to derive a common research agenda.

This report follows the structure of the workshop as described in Table 1 that lists the presentations, speakers, and discussants that took part in the workshop agenda. Section 2 introduces a set of overview talks meant to build a common ground for in-depth discussions. Section 3 covers a number of discussion sessions lead by some of the participants. We close in Section 4 with some general

<i>Overview Talks</i>	
Reconsidering Internet Architecture	Presenter: John Doyle (Caltech) Discussants: Ramesh Govindan (U. of Southern California), Fernando Paganini (U. ORT Uruguay), Lixia Zhang (UCLA)
Threats to Internet Survivability	Stefan Savage (UCSD)
Powergrid and Internet	Dominic Gross (U. of Wisconsin-Madison), Steven Low (Caltech), Lang Tong (Cornell U.)
Control and Learning	James Anderson (Columbia U.)
Resilience Engineering	David Alderson (Naval Postgraduate School), John Allspaw (Adaptive Capacity Labs), and David Woods (Ohio State U.)
Public & Private Sector Perspective	Henning Schulzrinne (Columbia U.) and Marwan Fayed (Cloudflare)
<i>In-depth Discussions</i>	
Threats to Internet Survivability	Zakir Durumeric (Stanford U.), Stefan Savage (UCSD), Aaron Schulman (UCSD)
Control and Learning	Nik Matni (U. of Pennsylvania)
Resilience Engineering	David Alderson (Naval Postgraduate School), John Allspaw (Adaptive Capacity Labs), Lorin Hochstein (Coupang), Zoran Perkov (Super Stealth Startup Inc.), David Woods (Ohio State U.)
Powergrid and Internet	Dominic Gross (U. of Wisconsin-Madison), Steven Low (Cal Tech), Fernando Paganini (U. ORT Uruguay), Joshua Taylor (U. of Toronto), Lang Tong (Cornell U.), Le Xie (Texas A&M U.)
Public & Private Sector Perspective	Henning Schulzrinne (Columbia U.), Marwan Fayed (Cloudflare), Doug Montgomery (NIST), Yih-Chun Hu (UIUC)

Table 1: Overview of workshop topics and participants.

observations and future directions. We aim to provide a faithful summary of the workshop presentations and discussions and reflect the participants’ diverse views.

2 OVERVIEW TALKS

The set of overview talks began with a presentation of recent advances on a theory of architecture by John C. Doyle, followed by presentations from Ramesh Govindan, Fernando Paganini, and Lixia Zhang, linking some of the concepts discussed in the context of today’s Internet architecture.

The remaining talks included a general introduction to threats to Internet survivability, by Stefan Savage, the architecture and challenges of the current power grid, led by Dominic Gross, Steven Low, and Lang Tong, and introductions to control and learning, by James Anderson, and resilience engineering, by David Alderson, John Allspaw, and David Woods. Henning Schulzrinne and Marwan Fayed closed these overviews with discussions on the role of the public and private sectors.

The following paragraphs present brief summaries of these talks with references to the relevant material.

2.1 Reconsidering Internet Architecture

The starting point of John Doyle’s presentation was that in the last decade, there have been significant advances in our understanding of how complex systems such as the human brain or the Internet work, both in terms of theory and its applications. In particular, this understanding has shed new light on complex system architectures in general, is particularly relevant for re-architecting the future Internet and power grid, and promises to be even more important in the context of envisioned cyber-physical systems (CPS) that use the Internet as “brain” for control of their physical networks (e.g., transportation network, public water systems).

To illustrate what new theory there is now and that might be relevant for the Internet or power grid, John Doyle focused on the (human) brain and especially on how it does sensorimotor control of the (human) body and used it as a canonical case study. In discussing his recent efforts on this topic (as described in [22]), among the key points he highlighted were (1) the need to understand how speed-accuracy tradeoffs at the level of individual components (i.e., nerves comprised of bundles of axons) connect to and characterize the speed-accuracy tradeoffs of the system that is comprised of these components (i.e., subsystems involved in sensorimotor control), (2) the ubiquity of diversity in living and engineered systems and the underlying mechanisms through which diversity in the delays and rates of sensing and signaling between layers improves the performance of (layered) control systems, and (3) the universal principle behind “diversity-enabled sweet spots (DESSs)” and the importance of examining what role this principle plays (or doesn’t play, and why not) in the exploration of layered architectures encountered in such diverse systems such as the bacterial cells, cell phones, and the Internet.

He argued that this is the richest existing case study demonstrating how sophisticated cyber systems (e.g., brains) control complex physical networks (e.g., human bodies) and how systems such as (human) brains have a richly layered architecture that has a far more sophisticated cyber control (e.g., Internet) of physical systems (e.g., CPSs) than anything that we have engineered/built yet. John Doyle concluded his talk with the ominous observation that while these richly layered architectures exhibit enormous robustness and evolvability, they are also prone to severe fragilities. In particular, he mentioned that ongoing efforts towards massive virtualization of much of modern technology make catastrophic failure events almost inevitable, just as our biological architectures make cancer, auto-immune disorders, and other life-threatening diseases largely unavoidable. On a more positive note, he expressed his hope that

once there will be enough “useful” case studies (such as the one that shows how the human brain does sensorimotor control) that demonstrate “the good, bad and the ugly” that our current architectures promote (and also show why), it will be possible to start leveraging the new theory and make attempts at re-architecting that tip the scale from new architectures that promote “the bad and the ugly” towards those that ensure “the good”

Ramesh Govindan, Fernando Paganini, and Lixia Zhang followed John Doyle with presentations that explored some of the design principles discussed in John’s talk in the context of two highly-engineered systems – the Internet and the power grid.

Ramesh Govindan’s presentation touched on the challenges of ensuring the availability of the global-scale infrastructures of hyper-scalers and the services they support. He presented findings of a “root-cause” analysis of large-scale failures in Google’s world-wide backbone network. The presented analysis could be considered as a motivation for the need for a broader and more in-depth understanding of identified root causes that goes beyond purely technical or engineering issues and explores how human decision-making (at layer 8 - the “social layer”), management decisions (at layer 9 - the “economic layer”), or regulatory policies (at layer 10 - the “political layer”) may turn out to be the ultimate culprits (individually or in combination).

Fernando Paganini’s presentation focused on decentralization of control architectures, drawing comparisons between the Internet and the power grid. For the Internet, a decentralized, layered architecture has operated well when there is abundance of bandwidth, together with buffering to manage transient traffic imbalances. In the power grid, various factors require the control of a centralized “system operator” entity. These factors include the peculiarities of AC power flow constraints, scarce transmission capacity, and the global dynamic effects of imbalance. However, both the Internet and the power grid are undergoing changes. While more centralized forms of control appear in the Internet (e.g. in cloud computing infrastructures), in the power grid, the massive deployment of distributed energy resources calls for increasingly more decentralized operations. For each of the two domains, the correct mix remains an open question, but will have to be recognizant of their mutual interdependence which, in turn, will be impacting their survivability under large-scale failure scenarios.

In the last presentation, Lixia Zhang challenged us to carefully consider what we mean by the current Internet architecture and how we envision any attempts at re-architecting it, especially when considering the ongoing changes to its layered organization, the evolving hourglass [27], and the seemingly endless layers of virtualization (e.g., RFC 9484 describing the tunneling IP through an HTTP server acting as an IP-specific proxy over HTTP [23]). Her talk was a reminder that there are really two alternatives to “re-architecting” the Internet – should we take a more *evolutionary* approach that is exemplified by the IP-over-HTTP example or is there a need to contemplate a more *revolutionary* approach such as the one articulated in [21])?

2.2 Threats to Internet Survivability

Stefan Savage gave an overview talk in which he discussed some of the different facets of threats to Internet security/survivability.

Starting with a historical perspective, he argued that while some of the core distributed Internet protocols and services were designed in a cooperative environment and were implemented in a similarly trusted world, subsequent effort to secure them against malicious intents by third parties (e.g., misusing DNS, hijacking BGP) have been largely unsuccessful.

As for the main reasons, Stefan pointed towards important trade-offs between distributed or decentralized and centralized designs. On the one hand, centralized designs are in general simple, cheap and practical but typically hamstringing innovation, limit expansion and scalability, and magnify the impact of problems or failures. On the other hand, while decentralized designs of protocols and services support innovation and expansion, they tend to cause complications (e.g. complex and unknown dependencies), create transitive trust relations that are both easier to attack and more difficult to scale, and result in limited visibility (i.e., difficult to audit). Importantly, as a community, we lack a good theory about where and when to use centralized vs decentralized designs.

In discussing the different aspects of this trade-off, Stefan first pointed out that economic forces favor centralized designs and described recent trends towards centralization in almost every aspect of the Internet ecosystem, from physical network infrastructure and access provisioning to service infrastructure and applications and services. For example, according to the 2019 Global Internet Report [1], at the service level, six companies deliver the majority of web resources, and the top three DNS, CA and CDNs cover between 50-70% of the top 100k sites. At the same time, a handful of operators run all gTLD registries, a few public resolvers are centralizing DNS resolution, and Microsoft and Google handle email for 30-40% of all domains [10, 16–18]. From a security/survivability perspective, this type of centralization clearly amplifies the impact of problems such as failures and attacks, as several recent events have shown us (e.g., Nashville bombing of 2020[25], Facebook incident of 2021 [13], Rogers’ 2022 outage [30]).

He then elaborated on the fact that the systems that comprise today’s Internet have become increasingly inter-dependent, creating complex and often unknown dependencies, with no straightforward ways to produce dependency graphs (e.g., do two ISPs share physical infrastructure and where?). He commented on a lack of a real composition architecture for cloud services and emphasized the fact that the lack of resilience in such increasingly inter-dependent systems is largely invisible - until some failure event occurs. He concluded his presentation by pointing towards three main culprits for the current state of affairs in today’s Internet: (i) The current architectures of the Internet as a whole and of the various systems that comprise the Internet are not designed for audibility (so integrity failures can be invisible); (ii) the key protocol deployments are not well-tested against threats that compromise their correct use and operation (e.g., DDoS will always be with us), and (iii) the design for resilience and the detection/mitigation of problems are severely hamstrung by limited visibility and a lack of good theory.

2.3 Powergrid and Internet

Steven Low organized an overview presentation where he, Dominic Gross, and Lang Tang discussed basic aspects of the power grid and key differences between the power grid and the Internet and

described issues that arise in the context of the ongoing transformation of today's power grid into tomorrow's smart grid.

In his presentation, Steven Low articulated the key differences between today's Internet and today's power grid by asking (and answering) three key questions: (Q1) What is the function of the Internet (power grid); that is, what does the Internet (the power grid) provide for applications? (Q2) What are the challenges that the Internet (power grid) faces and must overcome to support its function? and (Q3) What type of control system does the Internet (power grid) use to overcome these challenges? In short, for the Internet the answers are (I1) its function is to transfer byte-streams reliably end-to-end from senders to receivers, (I2) the challenges include lost or out-of-order packets and bit errors during transmission, and (I3) it utilizes a control system that has a layered architecture and is fully decentralized. In contrast, for today's power grid, the answers are (P1) its function is to transfer power at nominal voltage and frequency from generators to loads according to Kirchhoff's laws, (P2) its challenges concern generation-demand imbalances that can result in safety and power quality issues, including violations of frequency limits, voltage limits, or line capacity limits; and (P3) it uses a control system for balancing generation and demand everywhere that exhibits a time-scale based hierarchy and is largely centralized. Table 1 provides a further differentiation between today's Internet and power grid and is reproduced here from Steven Low's presentation.

Steven Low then discussed aspects that are of critical importance for ongoing efforts to design, deploy and operate tomorrow's smart grid infrastructure. On the generation side, these aspects include the use of uncertain, not dispatchable and typically highly intermittent sources of energy (e.g., solar and wind power) and the rapid expansion of distributed energy resources (DERs) and inverter-based resources (IBRs) that have low or zero inertia and give rise to new dynamic patterns that are absent in today's power grid with its generator-based control with large inertia. Another critical aspect for the future grid is the potential for significant energy storage. Table 2 (also reproduced here from Steven Low's presentation) succinctly summarizes these key aspects that differentiate today's power grid from tomorrow's smart grid and will require a major overhaul of the current grid control paradigm.

In his short talk as part of this session, Dominic Gross focused on the interoperability of the Internet and the power grid and addressed three future grid-specific topics. In particular, he discussed (i) the resilience of emerging power systems where converter-interfaced generation, storage and transmission are expected to dominate and produce fast time-scale dynamics that remain poorly understood; (ii) the need for grid-supporting Internet infrastructure, including grid-forming data-center concepts that can provide grid support on fast time scales, scalable and secure communication networks, and communication functions and infrastructure tailored to power system control and coordination; and (iii) the need for Internet-supporting power systems functions and infrastructure such as energy storage and power flow control to prioritize the power supply for critical information and communication infrastructure and technology-specific equipment/traffic and power flow control and medium voltage direct current (MVDC) to inter-link data-centers, power generation, and storage efficiently and reliably.

Lang Tong gave the last short talk in this overview session and discussed the requirements for next-generation monitoring and control for grid resiliency. In particular, he addressed implications of the increasing use of uncertain, not dispatchable and typically highly intermittent sources of energy and the rapid proliferation of DERs on future grid monitoring and control architectures and commented on the impact that these developments have on the requirements for the Internet as far as its use for effective monitoring and grid control is concerned.

2.4 Control and Learning

In his presentation, James Anderson introduced the System Level Synthesis (SLS) framework, a novel perspective on constrained robust and optimal controller synthesis for linear systems [3]. This framework featured implicitly in John Doyle's presentation where he used the canonical example of how the (human) brain does sensorimotor control of the (human) body. James highlighted how by working directly with system responses, SLS provides transparency in how system constraints, structure, and uncertainty affect controller synthesis, implementation, and performance. He showed that it is this transparency that can be exploited to improve upon the state-of-the-art so as to be able to apply controller synthesis at Internet scales.

For illustrative purposes, James focused on two particular applications of SLS, namely large-scale distributed optimal control and robust control. In the case of distributed control, he showed how SLS allows for localized controllers to be computed, extending robust and optimal control methods to large-scale systems under practical and realistic assumptions. In the case of robust control, he described how SLS allows for novel design methodologies that, for the first time, quantify the degradation in performance of a robust controller due to model uncertainty and emphasized that transparency is key in allowing robust control methods to interact, in a principled way, with modern techniques from machine learning. In explaining these applications, he focused on practical and efficient computational solutions and demonstrated the methods on easy to understand case studies.

James concluded his introduction to SLS with a brief discussion of promising ongoing research efforts in this area, including integrating SLS into model predictive control algorithms, combining optimal control and machine learning (ML), further understanding the algebraic structure underlying localized controllers and their state-space realizations, and applying the resulting new tools to application areas spanning power-systems, the Internet, and other cyber-physical systems of societal or economic importance.

2.5 Resilience Engineering

David Alderson, John Allspaw, and David Woods introduced "Resilience Engineering (RE)" and provided an RE perspective on the goal of the workshop, namely re-architecting today's internet for survivability. Alderson began with the simple point that Internet function is much more than routing, to include all the value-added layers above routing that now work together to provide an ecosystem of Critical Digital Services. That is, one can identify a number of failure scenarios where Internet routing works perfectly fine, but the broader ecosystem of services is severely disrupted. As a

Internet	Power grid
Layerd architecture	Time-based hierarchical control
Decentralized control	Centralized control
Storage everywhere	No significant storage
Dynamics & control: fast and narrow timescale (congestion control ~100ms, routing ~mins)	Dynamics & control: slower and wider timescale (power electronics ~ms, AGC ~sec-mins, market ~hours-days)
Packets follow routing algorithms	Power flows according to Kirchhoff's laws
Control & economics are decoupled	Markets are integral part of control

Table 2: Comparison: Today's Internet vs today's power grid.

Today's grid	Future grid
Generator-based control with large inertia	IBRs and DERs with zero to low inertia
Few large control points	Many small control points
Slow dynamics and control (~sec-mins)	IBR enables fast control
Frequency deviation is global control signal	Greater reliance on the Internet for denser communication
No significant storage (at timescales above ~30sec)	Potential for significant storage (e.g., EV, H2, flexible loads)
Market conditions: dispatchable generation and high marginal costs	Market conditions: uncertain/intermittent generation and ~zero marginal costs

Table 3: Comparison: Today's grid vs future grid.

result, the stated goal of "Internet survivability" in the presence of incidents needs to be much more than continued routing.

Responding to the argument in the workshop prospectus that "[the] transformation of the Internet into a critical infrastructure has occurred largely by happenstance, rather than by design," Alderson argued that this transformation has not actually been happenstance, but representative of broader patterns in adaptive behavior found in biology, cognitive systems, economics, engineering, social systems, etc. This "slide-to-criticality" for technologies—from nice-to-have, to front-line, to mission-critical, to essential—is ubiquitous in human and human-technology systems. For example, the recent discovery of the liblzma backdoor [11]—from evolution, vulnerability, hijacking to recognition—demonstrates a general pattern of adaptation for advantage that also produces mal-adaptive patterns that cross many layers well beyond the usual representations of the Internet, software, or technology stacks. Moreover, the real world provides continuing streams of incidents that invite study into these patterns, specifically as: (1) an empirical opportunity for learning about dealing with complexity, (2) context for developing theory to understand how resilient systems survive, and (3) a platform for engineering new architectures with adaptive capacity.

The concern in the workshop prospectus that "[the] evolved architecture of today's Internet cannot live up to this new role humanity has assigned it or withstand the types of threats that it now faces" is consistent with the challenges faced by other systems whose growth has led to increased complexification. That is, such systems must face growing system complexity (stimulated by new technologies and opportunities), new conflicts and threats

(as others 'hijack' capabilities for their own purposes), a changing environment with external events at scale (e.g., climate-driven extremes), and changing tempos of activity and larger shifts in tempo (as the world pushes to do things 'faster, better, and cheaper'). A major challenge in today's ecosystem of Critical Digital Services is whether we can learn how to offset changing risks before failures occur as growth continues. Or more specifically: Can we build capabilities to be poised to adapt to keep pace with and stay ahead of the trajectory of growing complexity and the penalties that arise as a result [38]?

Woods provided a brief introduction to Resilience Engineering, which has evolved over the last twenty years as a field [14, 15] and a community [24] devoted to understanding how adaptive systems, at all scales, possess the capacity to stretch or extend performance and avoid brittle collapse when events challenge their normal competence for handling situations. In particular, the Theory of Graceful Extensibility (TGE) [37] derives three subsets of principles (Subset A: risk of saturation, Subset B: networks of adaptive units, Subset C: constraints on maneuver) faced by all entities in the adaptive universe. These principles follow from three fundamental and inescapable constraints: (1) resources are finite (and therefore, conflict is ubiquitous); (2) change is continuous (therefore, models become stale and surprise recurs); (3) other units at other layers are adapting for advantage from their perspective. Collectively, TGE lays out a foundation for architecting systems that can adapt to challenges ahead, even when the exact challenge to be handled cannot be completely specified in advance. The pursuit of such a system architecture remains an important research challenge, and it is

particularly critical for the ongoing design and management of infrastructure systems [38]. To date, the adoption of RE design principles in critical infrastructure is nascent and remains a significant line of effort [2].

Despite the attention and progress, *resilience* as a concept remains noisy in the literature, largely due to its recent popularity across disparate communities as an organizing principle for managing stress and/or change. As described by Woods [36], there are four distinct notions commonly associated with resilience—rebound, robustness, extensibility, and sustained adaptability—with implications for how to engineer these features into complex systems [8]. For a review of these concepts as they have been studied in the context of network optimization, see also [26].

Resilience Engineering for the Internet has focused primarily on software, with past successes the result of a consortia of academia and industry studying how Critical Digital Services cope with complexity over cycles of growth, adaptation, challenge and surprise. For example, the STELLA Report [34] was the first result of a multi-year project called “Coping With Complexity” in which Ohio State’s Cognitive Systems Engineering Lab partnered with IBM, IEX, Etsy, and other organizations critically dependent on software infrastructure up and down the stack.

Allspaw provided a brief history of the DevOps movement and how it has led to a key acknowledgment: *how software behaves in the real world cannot be predicted or anticipated comprehensively*. That is, practitioners now believe that software cannot be built “correctly,” rather it must be operated. In turn, this means that there is no crisp boundary between ‘application developer’ and ‘systems engineer’ roles. Moreover, the rise of continuous deployment in Critical Digital Services has necessitated the use of various hedging strategies for managing the risk of brittle failure, as well as novel techniques for understanding disruptive events.

A starting point for re-architecting today’s Internet is a true understanding of the factors contributing to the incidents that cause Internet service disruption. Here, classic results in cognitive systems engineering distinguish between *Work as Imagined (WAI)* versus *Work as Done (WAD)*; see [41], with quotes from [35].

<u>Work as Imagined (WAI)</u>	<u>Work as Done (WAD)</u>
<ul style="list-style-type: none"> • System is built and operated as designed • Components of the system (humans, algorithms, devices) behave as specified • Exceptions/Anomalies are relatively few and usually well anticipated. 	<ul style="list-style-type: none"> • “Adaptations tailored to contingencies and context are always going on” • “The adaptations that make the system function also hide the systems weaknesses.” • “Management often can’t see the gaps so it seems that the system is functioning as designed.” • Anomalies and surprises are continuous.

This important distinction has revealed itself empirically in the handling of real Internet outages, with a large and evolving community of effort organized under the heading of “Learning From Incidents (LFI),” see [19]. At the core of this approach is a focus on incident analysis with “blameless” postmortems, using near misses to understand success, and moving beyond “human error” as a scapegoat

that precludes learning about system fragilities [41]. Collectively, the insights from the LFI community about how to manage Critical Digital Services have grown out of a disconnect between the way that Internet services are imagined versus the way that they actually are provisioned and operated.

<u>How we imagine incidents</u>	<u>How incidents actually happen</u>
<ul style="list-style-type: none"> • Need to find the root cause • Can be categorized in a taxonomy, measured, and usefully described with statistics • Humans are seen as the problem because they make mistakes 	<ul style="list-style-type: none"> • Things are always messy • Root cause analysis is a fallacy that hides the real problems lurking in system complexity • Taxonomies often hide rather than reveal; statistics like availability and mean time to failure (MTTF) are not useful • Humans are seen as a resource necessary for system flexibility and resilience

Of note and like many other artifacts that have resulted from the ongoing development of the Internet, the best practices being discovered and practiced by the LFI Community fall outside of any formal architecture for the current Internet.

People who operate Critical Digital Services confront forms of complexity and uncertainty under pressure. Here at the sharp end, there is a regular flow of incidents that threaten loss of valued services to stakeholders, and usually operations handle these threats successfully. The critical information about risks, threats, change, adaptation, growth—and therefore about architecture now and in the future—arises in studying how this sharp end adapts to cope with complexity [34]. The last two speakers in this overview session were practitioners from industry, who (a) tangibly experience the pressures, (b) develop means to better cope with the complexities for their organizations and industry segments, and (c) are thought leaders among the practitioner communities. They used incident vignettes to illustrate the evolution of tactics and strategies to cope with the complexities.

Lorin Hochstein used a particularly difficult case drawn from his experiences and reflections on incidents while at Netflix. The anomaly in this incident highlighted many findings about the cognitive and collaborative demands these situations present and the sources for resilient performance [7, 39].

Zoran Perkovic who has developed and managed the infrastructures enabling modern financial exchanges including IEX and NASDAQ, highlighted the web of complex interdependencies that spread from the base of the technology stack up to regulatory policies with financial and criminal penalties in a fiercely competitive, high stakes, massively autonomous, distributed environment. Interestingly, every change, every regulation, every competitive move, every new technique ends up being expressed and deployed as software with some autonomous capability providing the potential for many ‘strange’ interdependencies to emerge and combine across layers of the technology stack and engaging other layers of human goals, roles, and organizations. He recounted an incident episode which demonstrates the critical role of human expertise when the network of automated systems misbehaved — unfortunately — by behaving exactly as they were designed.

The talks and discussion in this session provided a real world sample of what it means to cope with complexities and the bottom-up adaptive innovations in knowledge, collaborations, policies, tactics and tools. This reality check grounds explorations of fundamental top-down theories for architecting the internet for the future.

2.6 Public & Private Sector Perspective

In this overview session, Henning Schulzrinne and Marwan Fayed presented a public and private sector perspective of the challenge of Internet survivability.

Henning Schulzrinne, who served as chief technology officer (CTO) for the United States Federal Communications Commission from 2011 to 2014, discussed the importance of the reliability and survivability of the Internet as a core civil infrastructure. His presentation highlighted the interdependencies of communication networks with other critical infrastructures like energy, transportation, and emergency services. He addressed the value of regulatory tools and policies to enhance network reliability and explored economic concepts such as asymmetric information and moral hazard, which affect market dynamics and infrastructure resilience. As part of his talk, Henning also commented on the need for regulatory intervention to mitigate market failures and ensure robust network performance during disasters.

Marwan Fayed, who is (acting) head of research at Cloudflare, a large network and content services operator as well as a faculty researcher, began by referring to a recent ACM co-sponsored research panel [28] to set context. Following that panel he reiterated that (i) packets are required for routing, but value is drawn from connections; (ii) exposing IP addresses to applications was a mistake of the socket interfaces [9]; and (iii) key management for routing and connection security remains a hard problem. Fortunately these have been and continue to be active areas of research.

Looking ahead, Marwan suggested three imminent Internet-wide challenges. First among them is that the Internet is relatively opaque. Unlike power grids and other critical infrastructures focused on improving instrumentation and visibility, the Internet anecdotally seems harder to understand and comes with less visibility – crucial elements for trust and ecosystem health. He also discussed the regionalization or sovereignty challenges emerging around the globe. Existing solutions strive for logical isolation via DNS and unicast and regional anycast, or physical isolation achieved with in-region datacenters and cabling. The former is known to affect resilience and performance [43], while the latter requires billions in capital and changing the Earth. A suggested design principle for future should be to devise mechanisms that enable data to flow where it chooses, with safeguards that can be trusted or verified.

Lastly, Marwan proposed a revised “narrow waist” model of the Internet in which edge networks and services have an opportunity to establish a unified interoperability layer for Internet infrastructure services, e.g. caching, DDoS, hosted firewalls, zero-trust, and others [29]. Prior narrow waists consist of the Internet Protocol between end-to-end and point-to-point protocols, as well as HTTP between client-server pairs and the networks that connect them. In the proposed model the edge services layer protects and improves performance of private infrastructure, from and with managed and

unmanaged devices on the public Internet. This presents opportunities to establish common edge service interfaces so that providers can differentiate on value, and that facilitate new entrants into the ecosystem.

3 IN-DEPTH DISCUSSIONS

The second day was organized around a series of in-depth discussions led by the workshop’s co-organizers. The following paragraphs attempt to summarize the key observations and arguments that were made in the course of these discussions.

3.1 Threats to Internet Survivability

The session started with a presentation by Zakir Durumeric who revisited the trust and visibility issues alluded to in Stefan Savage’s overview presentation and focused on the problem of trust and transparency. Using WebPKI to illustrate the current foundation for trust on the Web, he reminded the audience that authentication on the Web is based on validating X.509 certificates signed by Certificate Authorities (CAs) and described the Internet’s CA ecosystem, some 1,300 organizations that are trusted to validate the ownership or control of a domain. He then shared a number of critical observations: (i) Prior to 2012, the community had zero visibility into this ecosystem, (ii) only through relatively recent Internet scanning efforts did the community discover most CA certificates, (iii) these efforts revealed that pretty much everyone had the ability to sign certificates for any website and that CAs had been selling CA certificates to anyone who would pay for one. While the bad news is that without some sort of certificate transparency, CA certificates can’t be assumed to be trustworthy, the good news is that since 2017, Google Chrome requires all certificates to be logged in public Certificate Transparency (CT) logs, which in turn has dramatically improved the CA ecosystem. Zakir concluded by pointing out that transparency is a strong security primitive, requires that distributed trust can be appropriately monitored and verified, and may be a promising approach in other contexts as well (e.g., DNS, Internet routing).

In the second presentation in this session, Aaron Schulman returned to the problem of centralization in the physical network infrastructure and the risks that the resulting physical concentration poses for Internet access networks. Using the example of the outage that was caused by the 2020 Nashville bombing and damaged an AT&T network facility, he argued that while many of these edge facilities are repurposed houses or commercial buildings, they have been transformed over time into small data centers capable of supporting an increasing number of services. However, designed to withstand at best independent failures, these access facilities typically lack the means to survive targeted attacks intended to cause physical damage (e.g., fire or other intentional physical attacks). At the same time, because they are critical for supporting ever more services, they have also become more tempting targets for nefarious actors. One possible solution to make Internet access more robust to physical attacks on access facilities is for enterprises to utilize multiple independent access networks (including cellular providers) and for regulators or the market to incentivize multi-carrier access interconnections.

In her presentation, Morley Mao discussed the fragility of the Internet’s control plane (e.g., BGP, DNS) and argued that Internet survivability ought to mean more than just network connectivity but should also include the continued provision of basic and especially critical services. She then outlined some initial ideas about how the use of AI/ML might help make the Internet’s control plane more secure. To this end, she described a case study where connected and autonomous vehicles (CAVs) share sensor data to enhance perception capabilities (i.e., collaborative sensing), where the threat model considers a malicious participant that sends falsified data to the other participants, and where AI/ML-based approaches for detecting and mitigating such attacks on collaborative sensing have been considered. Morley suggested that similar approaches where CAVs are viewed as routing peers may be suitable for developing a more secure Internet control plane but may have to be considered in conjunction with the use of digital twins of the Internet cyber-physical system of interest to increase the robustness of the decision making(e.g., route selection).

In the last talk in this session, Alberto Dainotti came back to the visibility issue highlighted in Stefan Savage’s talk and described his groups’ recent work on observing Internet infrastructure failures (that sometimes coincide with power grid failures) and doing so at scale. In particular, he argued that understanding when, where and how Internet connectivity fails is challenging, mainly because network operators are reluctant to share failure data (may not even be aware of certain failures) and the core Internet protocols have not been designed with monitoring or auditing failure events in mind. Alberto described the design of IODA, a real-time system for monitoring Internet connectivity at the global scale, at the scale of individual countries and regions, and at the level of individual ASes. For each observed event, IDOA provides detailed information about the cause of the event, the operators/networks affected, how the communication stack was disrupted, and a timeline (including onset of event and restoration efforts). He finished his presentation with some illustrative examples, including a large CenturyLink outage in late 2018, the Venezuela blackout in 2019, the damage caused by the Russian war on Ukraine on network infrastructure and the power grid in Ukraine, and a timeline of measured Internet connectivity in Gaza since October 8, 2024.

3.2 Control and Learning

This session was intended to be less of a discussion-style session and provide instead a second overview talk on the topic of "Control and Learning". This second overview presentation was given by Nik Matni and was titled "System level synthesis and learning-based control". Building on the recent advances in the area of constrained robust and optimal controller synthesis for linear systems (collectively referred to as System Level Synthesis, or SLS) that were discussed by James Anderson in his overview talk on Day 1 of the workshop, Nik presented some of his recent work that combines robust control and machine learning (ML). On the one hand, robust control is needed because using feedback is one way to mitigate the effects of dynamic uncertainty (and provide worst-case and deterministic guarantees), especially when uncertainty is ubiquitous, not just in the environment but also in the utilized sensing methods/components and the considered models. On the other hand,

when faced with increasingly challenging environments, ever more difficult sensing tasks, and growing model complexity, using ML is a promising way to use past data to learn about and/or act upon the world, but deploying ML in the real world requires being able to provide stability, performance, robustness, and safety guarantees.

Nik showed how ML can be combined with robust control so as to reduce uncertainty by means of using more data to achieve better models/predictions) and at the same time mitigate uncertainty by improving performance thanks to better models/predictions. In particular, he argued that uncertainty is inherent in the output of any ML model, elaborated on what kind of uncertainty quantification is useful for control, and described how to explicitly account for this uncertainty when designing control policies. He then presented a case study that concerned the optimal control of an unknown system (I.e., instances with full information but unknown dynamics), mentioned a second case study involving perception-based control of a known system (i.e., instances with partial information obtained via complex sensing but known dynamics), and concluded his presentation with an illustration of a third case study that featured the problem of distributed optimal control of an unknown system (I.e., instances with asymmetric information and unknown dynamics).

In particular, he used this last case study to (i) highlight the difference between centralized dense control, sparse and distributed and localized control with delayed communications, and scalable learning-based distributed control; (ii) consider as a concrete instance the in-network congestion management problem, wherein a software-defined network is used to implement a distributed optimal controller designed to mitigate the effects of in-network congestion caused by rapid variations in traffic demand, and (iii) show that the design of such dynamic link-service rate policies can be cast as a learning-based distributed optimal control problem. Among the key lessons learned from these case studies were the observation that quantifying uncertainty in learned dynamics and sensing allows for leveraging tools from robust control and the insight that SLS makes transparent the effects of structure and uncertainty on controller implementation, complexity, sensing, performance, and safety.

3.3 Resilience Engineering

David Woods began his presentation with a review of past work on survivability and complex systems, specifically how complex systems fail. A key finding across engineering disciplines is that failure is due to brittle systems, not limited components, subsystems, or human beings. One such example is the signature of "Robust Yet Fragile (RYF)"—i.e., surprising sudden collapse against backdrop of continuous improvement and/or new capabilities—because systems "are robust to perturbations they were designed to handle, yet fragile to unexpected perturbations and design flaws" [6, p. 2529]. Such brittle failure can often be explained by one of several patterns of adaptive breakdown [40]:

- *Getting stuck in outdated models*: the world changes but the system remains stuck in what were previously adaptive strategies.
- *Working at cross-purposes*: behavior that is locally adaptive, but globally maladaptive. This results from an inability to coordinate across roles, units, and echelons as goals conflict.

- *Decompensation*: exhausting capacity to adapt as disturbances and/or challenges cascade. Breakdown occurs when challenges grow and propagate faster than responses can be decided on and deployed to effect.

Collectively, TGE (Woods) and DESS (Doyle) provide the start of architectural principles to overcome risks from the brittleness that arises naturally from having to operate in a high-dimensional tradeoff space. Some architectural principles have demonstrations in real but bounded settings where responsible human roles supervise highly autonomous operations. The principles provide general policies for how to behave when approaching saturation and when neighbors are approaching saturation. In some cases these policies can take mathematical form [26]. In others they take the form of new software protocols that modify late and counter productive behavior when approaching saturation [8]. Practically, progress in Resilience Engineering has provided concepts and/or techniques to design or modify operational practices to be more continuously adaptive as they provide valued services. Both the more formal and more immediately pragmatic steps fall outside the usual frameworks for Internet dependent architectures.

The results from studying how people adapt to cope with complexity appear to be couched in the language of cognitive, social and organizational perspectives — new layers added to the technology stack. But this is not really the case. These studies reveal fundamental patterns and laws about adaptive behavior in general across the biological, technological and human spheres. These regularities apply everywhere across the technology stack regardless of which layer is chosen as a point of departure. The regularities are about more than people as they capture issues about architectures, growth, interdependencies, complexification, and trade-offs that influence adaptive capacities in the face of uncertainty and change. Technology advances stimulate these processes to transform human worlds of activity, purposes, risks, conflicts and cooperation, and the consequences that follow.

John Allspaw continued the discussion by talking about the Internet as a critical infrastructure as capabilities are developed, modified, deployed and operated over time to provide valued services to stakeholders. Ironically many of the services support other service providers and expand and hide interdependencies from stakeholders. Critical Digital Services have adapted over time to produce growth and handle new challenges demonstrating many principles of adaptive systems. One of the adaptations to handle complexities was to switch from separate silos for development, deployment and operations of critical software services. This structural partition had too little adaptive capacity to handle the pace of change and pressure to deploy advantageous services. Instead, he argued for (and pioneered) continuous development and deployment linking feedback, risk, gain, change into a fluent process both stimulating growth but also handling the complexities that accompany growth.

Operations and design need to be tightly connected in future architectures. The adaptive path of Critical Digital Services highlights several underlying principles that are surprising. Complete knowledge and testing of the system (components, software, users) is not possible without contact with the full complexities of production traffic. Inevitably, events will challenge its operation. The system is always adapting locally under pressures to be better, faster, and

cheaper. Ultimately, we can learn about the boundaries of a design’s competencies only by operating it. The key question is whether we can learn fast enough to keep pace with change and growth.

David Alderson summarized the current strategy for mitigating risks in infrastructure systems—through the use of modeling and simulation to find vulnerability gaps and then plug them—and led a discussion about why this will not work for the Internet. Because there is no staging environment that is representative of real production systems, digital twins will not suffice to uncover the edge cases that potentially lead to large scale failure. Despite recent emphasis on stress testing for financial systems by the US Federal Reserve and others, the RE perspective suggests there is perhaps little that can be learned from stress tests. Moreover, what is “critical” in the system is going to be dynamic, further complicating this challenge.

The RE lens re-conceptualizes “Internet survivability” as how to sustain long-term viability of Internet dependent critical services as growth produces new types and scales of challenges. One might see the descriptive language that results from studies of coping with complexity as characteristic of cognitive, human, and organizational layers. But this occurs because the patterns of adaptation—experienced by the networking research community as the “happence” evolution of the Internet—are derived from regularities of people in systems exemplifying these patterns. The patterns of adaptation are about much more than people as they capture issues about architecture, layering, interdependencies, trade-offs, saturation, tempo, synchronization, reframing, and more in a dynamic, limited resource world. Among the drivers of challenge and adaptation, deploying new technologies flows through and transforms human worlds of activity, purpose and consequences.

Moreover, the Resilience Engineering perspective, as practiced in the LFI Community for Critical Digital Services and elsewhere, serves also as the basis for empirical study of how engineers must confront the complexity that arises when Internet architecture comes into contact with real-world pressures for performance. Such an empirical grounding is an essential ingredient for any future re-architecting and is not currently being addressed elsewhere.

3.4 Powergrid and Internet

The session started with a presentation by Le Xie who used the cryptocurrency mining operations in Texas as an illuminating case study for illustrating the interaction between large flexible computing loads and the power grid. He presented data showing the impact of energy consumption of cryptocurrency mining data centers on the peak electric demand in Texas in the summer of 2022 and argued that the rapid growth of large flexible computing loads could bring both operational challenges and market design opportunities for power systems. Open research questions he posed included how to design the market signals so that flexible large computing loads could contribute maximally as demand response resources, especially during stressed grid operating conditions; and how to design proper incentive mechanisms in electricity markets to maximize the value and participation of cryptocurrency mining data center loads in provision of demand flexibility.

In the second presentation in this session, Josh Taylor talked about how the grid is used to transmit both energy and information. Sending information leverages the physics of the grid and can help with several different tasks such as fault protection and decentralized control. He argued that as more converter-interfaced resources are added to the system, this use of the grid is becoming more relevant because (i) it is easier for a converter to add small perturbations (e.g., by adding them to its controller setpoints), and (ii) at present, converters do not behave as predictably as synchronous machines. He illustrated this use with two examples: (i) fault detection, in which converters inject negative sequence current to make it easier for relays to distinguish between normal and faulty operation; and (ii) islanding detection, in which converters inject negative sequence current to detect when a portion of the grid has unintentionally disconnected and formed a self-powered island. Some of the open research questions he mentioned are how to optimize such perturbations so they are minimally disruptive, and for which tasks this use of the grid can and should be considered.

In the last talk in this session, Dominic Gross discussed in more detail aspects that concern the integration of power systems and the Internet infrastructure. In particular, he addressed a core question that arises in this context, namely whether or not the fact that the two systems become more dependent on each other necessitates a closer integration at the operational level and, if so, what entity or entities should drive such closer operational integration efforts. He suggested two plausible pathways to closer integration of power systems and Internet infrastructure: (i) closer collaboration of hyperscalers (i.e., large cloud service providers) and power system operators, and (ii) dedicated power infrastructure for hyperscalers. In particular, he commented on the facts that hyperscalers have already made significant Internet infrastructure investments (e.g., data-centers and subsea cables) and that a lack of recognition of the need for reliable power supply by power system operators and utilities may prompt them to invest into dedicated power infrastructure such as microgrids, renewable power generation and energy storage, and even dedicated power distribution infrastructure.

3.5 Public & Private Sector Perspective

As part of this session, organized by Marwan Fayed and Henning Schulzrinne, Yih-Chun Hu provided an introduction to SCION [42], a clean-slate secure Internet architecture designed to provide high availability in the presence of adversaries, trust and path transparency, and inter-domain multipath routing. It offers security, path-aware networking, and multipath communication, and has already adopted by operators like Swisscom and financial institutions such as the Swiss National Bank. SCION organizes ASes into isolation domains (ISDs), managed by a core set of ASes that establish trust roots and issue certificates. This path-based architecture allows end-hosts to select from multiple end-to-end paths, enabling rapid failover, dynamic traffic optimization, and robust DDoS defenses.

Doug Montgomery's presentation explored the current Internet architecture, primarily defined by protocols like TCP/IP, DNS, and BGP. He highlighted gaps in standardization, particularly with middleboxes, security functions, and network virtualization, and emphasized the need for more cohesive standards. Doug also discussed

security, contrasting protocol-specific measures with comprehensive network security, and advocated for a Zero Trust Architecture, where the default stance is to deny access unless explicitly authorized. He closed his presentation by questioning the current process of how the Internet's architecture is defined and standardized, and asking us to consider re-evaluating it if we are to ensure the network's long-term survivability.

Part of the discussion focused on open roaming in wireless networks during disasters. Open roaming in the US, especially in the event of natural disasters, has evolved significantly, transitioning from a largely voluntary practice (e.g., the arrangement between AT&T and T-Mobile USA during Hurricane Sandy [33]) to a mandatory requirement to improve the resiliency and reliability of mobile wireless networks before, during, and after emergencies (e.g., FCC-22-50 [31]). On June 4, 2024, the US Homeland Security Bureau announced procedures for states requests to activate the FCC Mandatory Disaster Response Initiative [32].

4 CLOSING THOUGHTS AND FUTURE DIRECTIONS

As organizers, we approached the workshop planning with the understanding that tackling a problem of the scale and complexity of "Internet survivability" mandates a cross-disciplinary effort that includes, among others, networking researchers and control theory experts, power/smart grid researchers and economists, political and social scientists, and public policy experts associated with either various government agencies or relevant private organization. The workshop discussions reinforced this understanding: any community-driven research agenda aimed at meaningfully addressing the workshop's stated challenge must be cross-disciplinary at its core. At the same time, we realized that no single meeting dedicated to this workshop's topic could cast a wide enough net to craft a detailed research agenda. Nevertheless, we believe that this initial workshop succeeded in identifying some key directions that should be part of any such agenda. The following is a selected list of lessons learned from these two intense days of discussions and possible future directions.

For networking researchers, an important item on their future research agenda is distilling the essence of ongoing foundational approaches to re-architecting today's Internet. These approaches include the consideration of a new economic architecture of the Internet that entails the creation of a "public option" for the Internet's core backbone [12] and the proposal for enabling a permanent revolution in Internet architecture via Trotsky [21], a novel architectural framework that provides a backwards-compatible path (i.e., ensuring the continued functioning of legacy applications or hosts) to an extensible Internet where both new architectures can be deployed in a backwards-compatible manner and multiple architectures can exist side-by-side – something that cannot be achieved with our current notion of IP as the Internet's narrow waist.

Viewed through a cross-disciplinary lens, these and similar approaches give rise to new questions of fundamental importance. For example, since effectively and efficiently operating and managing tomorrow's massively distributed power/smart grid relies increasingly on a well-functioning Internet that can provide provably secure communication for controlling the power/smart grid,

can we expect the economic incentives to be aligned with the technologies capabilities so as to support an architectural framework where one of the different co-existing architectures is secure by design and therefore satisfies the requirements that are necessitated by the growing mutual reliance between the future power/smart grid and tomorrow's Internet? Similarly, does the emerging theory of graceful extensibility advocated by Resilience Engineering [37], which lays out a foundation for architecting systems that can adapt to partially unspecified challenges ahead, inform us in meaningful and effective ways about ongoing efforts to design and manage critical infrastructure systems [38]? The adoption of Resilience Engineering-based design principles in critical infrastructure remains a significant line of effort [2].

For control theory experts, complex engineered and control systems, such as those used in the power/smart grid or the Internet, are characterized by needing to operate robustly and reliably across many spatio-temporal scales, despite being implemented using highly constrained hardware components and software. Moreover, control methods are, in general, only used to design algorithms in these components, typically with minimal or no theory, and the larger system that is comprised of these components is often designed by others. Despite these challenges, recent advances in control theory have identified a universal design pattern that centers around the notion of layered control architectures (LCAs) and has the potential for natural but large extensions of robust performance from control to the full decision and control stack [20]. Building on the "model LCA" described in [20] to initiate a quantitative study of LCAs, another critical item on a proposed future research agenda will be to identify the occurrences of different LCAs in the Internet and the power/smart grid (as well as other cyber-physical infrastructures), understand the underlying universal mechanisms and design patterns, and leverage this knowledge to outline tentative paths towards a useful design theory. The insights from such a theory will enable us to understand the many tradeoffs of complex engineered systems such as the Internet and the power/smart grid (see also [4]).

Last but not least, from a public policy perspective, as the importance of the Internet as a cyber-physical infrastructure critical for modern society and the global economy at large is increasingly recognized by the various stakeholders in both the public and private sectors, it seems fitting for local and federal governments to take on a more visible role in ensuring, monitoring, and incentivizing all aspects concerned with "Internet survivability". In particular, the increasingly mission-critical role that today's Internet is playing for an ever-growing number of stakeholders argues for the creation of dedicated agencies or public and/or private organizations whose sole focus is ensuring its long-term survivability. The US Cybersecurity and Infrastructure Security Agency (CISA), part of the Department of Homeland Security, is one such example, but complementary efforts that are concerned with visionary architectural frameworks and their possible realization or with more economics-driven architectural designs are needed.

5 WORKSHOP PARTICIPANTS

Organizers:

Fabián E. Bustamante (Northwestern U.)

Walter Willinger (NIKSUN Inc.)

Co-organizers:

David Alderson (Naval Postgraduate School)

Marwan Fayed (Cloudflare)

Steven Low (Caltech)

Stefan Savage (UCSD)

Henning Schulzrinne (Columbia U.)

Participants:

John Allspaw (Adaptive Capacity Labs)

Luis Amaral (Northwestern U.)

James Anderson (Columbia U.)

Todd Arnold (US Military Academy West Point)

Paul Barford (U. Wisconsin-Madison)

Pete Beckman (Argonne National Labs/Northwestern U.)

Zachary Bischof (Georgia Tech)

Alberto Dainotti (Georgia Tech)

John Doyle (Caltech)

Zakir Durumeric (Stanford U.)

Ramesh Govindan (USC)

Dominic Gross (U. of Wisconsin-Madison)

Lorin Hochstein (Coupang)

Yih-Chun Hu (UIUC)

Igor Kadota (Northwestern U.)

Z. Morley Mao (U. of Michigan)

Nikolai Matni (U. of Pennsylvania)

Deepankar Medhi (NSF)

Douglas Montgomery (NIST)

Fernando Paganini (Universidad ORT Uruguay)

Zoran Perkov (Super Stealth Startup Inc.)

Ahmed Saeed (Georgia Tech)

Aaron Schulman (UCSD)

Yixin Sun (U. of Virginia)

Joshua Taylor (NJIT)

Cecilia Testart (Georgia Tech)

Lang Tong (Cornell U.)

Ermin Wei (Northwestern U.)

David Woods (Ohio State U.)

Le Xie (TAMU)

Lixia Zhang (UCLA).

ACKNOWLEDGMENTS

We are grateful to the participants for contributing their insights at the workshop, and for feedback on this report. We are also grateful to the Department of Computer Science, Northwestern University, in particular Ms. Wynante R. Charles, for their help in the organization of this workshop. The workshop was supported by NSF CNS-2332178. The opinions expressed do not necessarily reflect views of the NSF. Any errors are the responsibility of the authors.

REFERENCES

- [1] 2020. Consolidation in the Internet Economy. (Feb 2020). <https://future.internetsociety.org/2019/>
- [2] D. L. Alderson. 2019. Overcoming barriers to greater scientific understanding of critical infrastructure resilience. *Handbook on Resilience of Socio-Technical Systems* (2019), 66.

- [3] James Anderson, John C Doyle, Steven H Low, and Nikolai Matni. 2019. System level synthesis. *Annual Reviews in Control* 47 (2019), 364–393.
- [4] A.M. Annaswamy, K.H. Johansson, and G.J. Pappas. 2023. Control for Societal-scale Challenges: Road Map 2030. *IEEE Control Systems Society Publication* (2023).
- [5] Fabián E. Bustamante and Walter Willinger. 2023. NSF Workshop: Towards Re-architecting Today’s Internet for Survivability. (November 2023). <https://aqualab.cs.northwestern.edu/nsfworkshop23-internetsurvivability/> This web page includes the Workshop Agenda, presentation slides and other documents.
- [6] Jean M Carlson and John Doyle. 2000. Highly optimized tolerance: Robustness and design in complex systems. *Physical review letters* 84, 11 (2000), 2529.
- [7] Richard I. Cook. 2020. Above the line, below the line. *Commun. ACM* 63, 3 (feb 2020), 43–46. <https://doi.org/10.1145/3379510>
- [8] Richard I. Cook and Beth Adele Long. 2021. Building and revising adaptive capacity sharing for technical incident response: A case of resilience engineering. *Applied Ergonomics* 90 (2021), 103240. <https://doi.org/10.1016/j.apergo.2020.103240>
- [9] Marwan Fayed, Lorenz Bauer, Vasileios Giotsas, Sami Kerola, Marek Majkowski, Pavel Odintsov, Jakub Sitnicki, Taejoong Chung, Dave Levin, Alan Misllove, Christopher A. Wood, and Nick Sullivan. 2021. The ties that un-bind: decoupling IP from web services and sockets for robust addressing agility at CDN-scale. In *Proc. of ACM SIGCOMM*. 433–446. <https://doi.org/10.1145/3452296.3472922>
- [10] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven years in the life of Hypergiants’ off-nets. In *Proc. of ACM SIGCOMM*.
- [11] Dan Goodin. 2024. What we know about the xz Utils backdoor that almost infected the world. (2024). Published online 3/31/2024. <https://arstechnica.com/security/2024/04/what-we-know-about-the-xz-utils-backdoor-that-almost-infected-the-world/>.
- [12] Yotam Harchol, Dirk Bergemann, Nick Feamster, Eric Friedman, Arvind Krishnamurthy, Aurojit Panda, Sylvia Ratnasamy, Micheal Schapira, and Scott Shenker. 2020. A Public Option for the Core. In *Proc. of ACM SIGCOMM*.
- [13] Alex Heath. 2021. Locked out and totally down: Facebook is scrambling to fix massive outage. *The Verge* (October 4 2021).
- [14] E Hollnagel (Ed.). 2013. *Resilience engineering in practice: A guidebook*. Ashgate Publishing, Ltd.
- [15] E Hollnagel, DD Woods, and N Leveson (Eds.). 2006. *Resilience Engineering: Concepts and Precepts*. Ashgate Press, Aldershot, UK.
- [16] Geoff Huston. 2019. DNS Resolver Centrality. APNIC Blog. (September 2019). <https://labs.apnic.net/?p=1260>
- [17] Aqsa Kashaf, Vyas Sekar, and Yuvraj Agarwal. 2020. Analyzing third party service dependencies in modern web services: Have we learned from the miradyn incident?. In *Proc. of IMC*.
- [18] Rashna Kumar, Sana Asif, Elise Lee, and Fabian E. Bustamante. 2023. Each at Its Own Pace: Third-Party Dependency and Centralization Around the World. *Proc. ACM Meas. Anal. Comput. Syst.* (2023).
- [19] Learning From Incidents (LFI). 2019. (2019). <https://www.learningfromincidents.io/>
- [20] Nikolai Matni, Aaron D Ames, and John C Doyle. 2024. Towards a Theory of Control Architecture: A quantitative framework for layered multi-rate control. *arXiv preprint arXiv:2401.15185* (2024).
- [21] James McCauley, Yotam Harchol, Aurojit Panda, Barath Raghavan, and Scott Shenker. 2019. Enabling a Permanent Revolution in Internet Architecture. In *Proc. of ACM SIGCOMM*.
- [22] Yorie Nakahira, Quanying Liu, Terrence J. Sejnowski, and John C. Doyle. 2019. Diversity-enabled sweet spots in layered architectures and speed-accuracy trade-offs in sensorimotor control. *Proc. of the National Academy of Sciences (PNAS)* 118, 22 (2019).
- [23] Tommy Pauly, David Schinazi, Alex Chernyakhovskiy, Mirja Kühlewind, and Magnus Westerlund. 2023. Proxying IP in HTTP. RFC 9484. (Oct. 2023).
- [24] Resilience Engineering Association (REA). 2004. (2004). <https://www.resilience-engineering-association.org/>
- [25] Adam; McGee Jamie Rojas, Rick; Goldman. 2020. A Quiet Life, a Thunderous Death, and a Nightmare That Shook Nashville. *The New York Times* (December 27 2020).
- [26] Thomas C. Sharkey, Sarah G. Nurre Pinkley, Daniel A. Eisenberg, and David L. Alderson. 2020. In search of network resilience: An optimization-based view. *Networks: An international Journal* 77, 2 (2020). <https://doi.org/10.1002/net.21996>.
- [27] Oliver Spatscheck. 2013. Layers of Success. *IEEE Internet Computing* 17, 1 (2013).
- [28] theNetworking Channel. 2023. Lessons learned from 40+ years of the Internet. online. (October 2023). <https://networkingchannel.eu/lessons-learned-from-40-years-of-the-internet-downloads/>
- [29] theNetworking Channel. 2023. Lessons learned from 40+ years of the Internet: an Industry Perspective. online. (November 2023). <https://networkingchannel.eu/lessons-learned-from-40-years-of-the-internet-an-industry-perspective-downloads/>
- [30] Joao Tomé, Tom Strick, and Mingwei Zhang. 2022. Cloudflare’s view of the Rogers Communications outage in Canada. *The Cloudflare Blog* (July 8 2022).
- [31] US Federal Communications Commission. 2022. Report and Order and Further Notice of Proposed Rulemaking FCC 22-50. (2022).
- [32] US Federal Communications Commission. 2024. Public Notice DA 24-527. (2024).
- [33] Zack Whittaker. 2012. AT&T, T-Mobile open networks, offer free roaming for Sandy relief. *ZDNET* (November 2012). <https://www.zdnet.com/article/at-t-t-mobile-open-networks-offer-free-roaming-for-sandy-relief/>
- [34] D.D. Woods. 2017. STELLA Report from the SNAFUcatchers Workshop on Coping With Complexity. (2017). Brooklyn NY, March 14-16, 2017.
- [35] D Woods, T Licu, J Leonhardt, M Rayo, E Balkin, and R Ciponea. 2021. Patterns in How People Think and Work: Importance of Patterns Discovery for Understanding Complex Adaptive Systems. In *EUROCONTROL*.
- [36] David D Woods. 2015. Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering & System Safety* 141 (2015), 5–9.
- [37] David D Woods. 2018. The theory of graceful extensibility: basic rules that govern adaptive systems. *Environment Systems and Decisions* 38, 4 (2018), 433–457.
- [38] David D Woods and David L Alderson. 2021. Progress toward Resilient Infrastructures: Are we falling behind the pace of events and changing threats? *Journal of Critical Infrastructure Policy* 2, 2 (2021), 5–18.
- [39] David D Woods and John Allspaw. 2020. Revealing the critical role of human performance in software. *Commun. ACM* 63, 5 (2020), 64–67.
- [40] David D Woods and Matthieu Branlat. 2017. Basic patterns in how adaptive systems fail. In *Resilience engineering in practice*. CRC Press, 127–143.
- [41] David D Woods, Sidney Dekker, Richard Cook, Leila Johannesen, and Nadine Sarter. 2010. *Behind human error* (2nd ed.). CRC Press.
- [42] Xin Zhang, Hsu-Chun Hsiao, Geoffrey Hasker, Haowen Chan, Adrian Perrig, and David Andersen. 2011. SCION: Scalability, control, and isolation on next-generation networks. In *Proc. IEEE Security*.
- [43] Minyuan Zhou, Xiao Zhang, Shuai Hao, Xiaowei Yang, Jiaqi Zheng, Guihai Chen, and Wanchun Dou. 2023. Regional IP Anycast: Deployments, Performance, and Potentials. In *Proc. of ACM SIGCOMM*. 917–931. <https://doi.org/10.1145/3603269.3604846>