On the Importance of Resilience Engineering for Networked Systems in a Changing World

David Hutchison, Dimitrios Pezaros, Jacek Rak, and Paul Smith

The authors summarize recent and current research in resilient systems and, consequently, propose a multidisciplinary research agenda in resilience engineering for networked systems.

Abstract

Resilience is featured increasingly often in the media, usually applied to society when faced, for example, with disasters such as flooding and the enormous challenges that the Covid-19 pandemic posed. There are now many resilience-related discussion groups worldwide, and some standards initiatives devoted in particular to city resilience. However, there is relatively little explicit interest in resilience engineering for communication networks and systems, including the Internet. This is perhaps surprising, given the reliance that society now places on networks and networked systems. This article reflects on key issues and developments that may change this perspective; we summarize recent and current research in resilient systems and, consequently, propose a multidisciplinary research agenda in resilience engineering for networked systems.

Introduction

The importance of resilience in the modern world must not be underestimated. This applies not least to the communication networks and services that support almost every aspect of our life and work. It seems to be assumed that these crucial engineered artefacts work — including the Internet — whatever challenges they face, but it is dangerous to believe so. During the Covid-19 pandemic, remote working and meetings as well as family interactions by means of teleconferencing became the norm. Most people generally had a satisfactory experience: communication networks served us well enough, largely due to increased (and targeted) investments in provisioning.

However, recent changes in the world around us, including new technologies and applications such as Industry 4.0, autonomous vehicles, climate change, and pandemics, are increasingly likely to lead to new and unforeseen challenges.

This article discusses these changes and the need they create for a holistic study of resilience for networks and networked systems, combining previous and new activities, and addressing risk and inevitable trade-offs. A summary of the key aspects that may influence the resilience of networked systems is shown in Fig. 1. These aspects are explained and explored in subsequent sections of the article.

We define resilience as "the ability of an engineered system to continue to provide its designed level of service in the face of any challenges." The term "resilience engineering" is used to cover the design, implementation, management, and longer-term sustainability of resilient networked systems by the organizations and people that own and operate them.

There has been growing interest in resilient networks and a corresponding amount of international research activity during the past decade or so. A book published just before the Covid-19 pandemic summarized this research effort and reported many new results specifically on the resilience of communications and networked systems [1]. The work was carried out in a European Union (EU) COST Action, called RECODIS, involving some 200 researchers from 31 countries.

Much earlier work on Quality of Service (QoS) had led to the publication of an international standards framework that identified the importance of availability and reliability for systems QoS alongside the more familiar QoS aspects including throughput, delay, and jitter [2]. Two decades ago, there was a push toward recognizing the impending importance of resilience for computer networks by some authors in the community. The motivation built on the realization that networks were becoming critical infrastructures, supporting critical services for society at large.

That research on resilience led to — alongside other efforts — the D^2R^2+DR framework, which introduced the fundamentals and principles of resilience for designing and building resilient networked systems (see chapter 1 in [1], and Fig. 2). Subsequent international collaborations led to the evaluation and further development of resilience strategies and mechanisms, including the EU ResumeNet project. This project and others led directly to the multinational RECODIS initiative and consortium.

Meanwhile, there are important new national initiatives aimed specifically at communication systems resilience, notably the recent US Cybersecurity & Infrastructure Security Agency (CISA)'s 5G Security and Resilience program [3].

The aims of this article are as follows. Reflecting on the outcomes of recent research, we try to capture the state of the art in resilience engineering for networked systems and characterize the main features of the changing environment for networked systems. We reiterate the importance of resilience ten years on from the introduction of the D^2R^2+DR framework and discuss what has

ISSN: 0163-6804
Digital Object Identifier: 10.1109/MCOM.001.2300057

David Hutchison is with Lancaster University, UK; Dimitrios Pezaros is with the University of Glasgow, UK; Jacek Rak is with Gdańsk University of Technology, Poland; Paul Smith is with Austrian Institute of Technology, Austria.

changed or is changing in the world of communications and networks. We analyze in detail the key factors from Fig. 1 that influence resilience strategies and debate their utility and validity. Next, we highlight selected outcomes from [1] aimed at reducing the impact of large-scale failure scenarios on the resilience of communications and networked systems. Finally, we outline the important resilience research topics and activities that still need to be tackled in order to build networked systems that are fit for purpose.

The Changing Environment

There are many challenges to computer networks and networked systems caused by a variety of externalities — some of them are familiar, and some have more recently emerged. The familiar ones are faults/failures and cyber attacks, while newer ones include extreme events arising from the climate crisis and, recently, the Covid-19 pandemic. These challenges can significantly impact the operation and resilience of computer networks and the services they support, e.g., loss of connectivity, unexpected and unpredictable load profiles, and disruptions caused by attack-induced failures. Moreover, the challenges introduce uncertainty into the environment in which computer networks are deployed, thus planning and risk management are much more demanding.

There are increasingly significant effects of climate change on the quality of people's lives. In particular, the number, intensity, and scale of natural disasters such as hurricanes, tornadoes, floods, and fires demonstrate their devastating power on affected components of communication network infrastructure [1]. These challenges often lead to simultaneous failures of multiple network components (termed *massive failures*), thereby significantly degrading the performance of the communication infrastructure or even making communication impossible at the times when people need it most. The immediate cause of such a scenario might be a partial degradation of the effective link capacity (see Chapters 13 and 21 of [1]), e.g., due to a decreased optical signal to noise ratio (OSNR), for instance of wireless optical links in adverse weather conditions. Requirements on OSNR also typically depend on the nominal capacity of optical transmission systems.

The inability to communicate or to receive rescue messages in many cases can even lead to loss of life. As communication networks are undoubtedly an integral part of contemporary critical infrastructure, the availability of their services during natural disasters should be regarded as a crucial concern for network operators, regulators, owner and stakeholder organizations, and user communities.

Indeed, despite utilization of various resilience mechanisms, current networked systems continue to fail due to new and growing problems. Taking weather-related disasters, we note that the 2017 hurricane Maria in Latin America was responsible for the lack of Internet access and mobile communications in Dominicana. The 2018 Attica fires in Greece made communications in affected areas (including for rescue teams operations) barely possible (Chapter 1 of [1]). The same reference also mentions hardware and software failures being a frequent cause for the unavailability

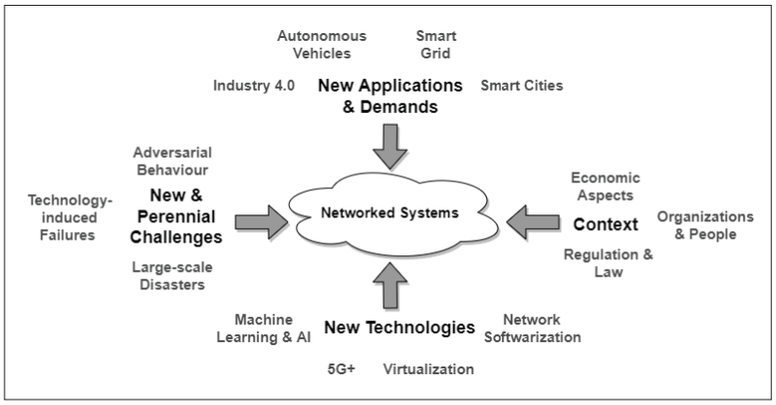

FIGURE 1. Networked systems in context: resilience influences.

of services at a global scale, e.g., reported failures of Amazon services from 2017, and Google and Microsoft services from 2019. Considerably more attention to the effectiveness of resilience mechanisms is, therefore, much needed.

The need for resilience is particularly acute in cyber-physical, industrial control environments that are increasingly interconnected with advanced Information and Communications Technology (ICT) capability for automation and sustainability. Such environments often do not benefit from well-understood ICT protection mechanisms, and they remain exposed to adversarial events manifest through their digital infrastructure. For example, Iran's Khuzestan steel manufacturing plant suffered catastrophic failure within the space of only a few minutes in June 2022, as widely reported in the media at the time, when its process control system was compromised by remote adversaries. And as reported in Security magazine, there was a 38 percent increase in cyber attacks in 2022, which is surely a major indicator of the urgent need for increased protection via resilience engineering.

The recent Covid-19 pandemic made the role of the Internet as a major critical (inter)national infrastructure clearer than ever. Internet communications now support even more of the services economy. From the online delivery of education to the remote treatment of health service patients and remote business working, it is becoming clear that new working modalities have been emerging that will persist into the future. In this environment, reliance on the Internet is much increased, therefore its ability to detect and react to challenges is absolutely crucial.

Sustainability and Contextual Barriers

Several largely non-technical inhibitors or obstacles make it challenging to implement networked systems resilience.

Cost and Sustainability

An obvious resilience inhibitor is its financial cost, in terms of capital investment and continued operational costs; it can be difficult to motivate enterprise investment in network resilience, especially for low-probability, high-impact events. This is problematic as many stakeholders in the provisioning of computer networks are private enterprises whose primary responsibility is to their shareholders.

Communication networks must evolve to

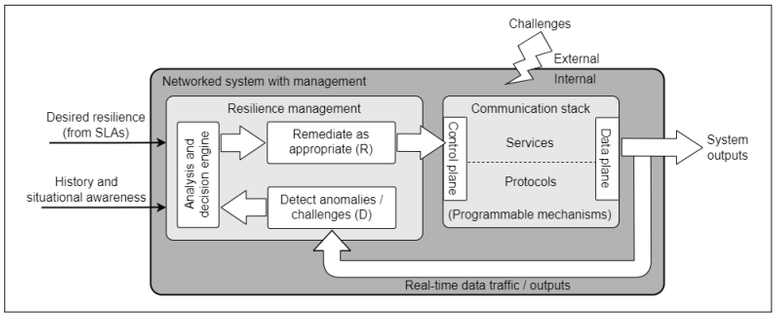

FIGURE 2. A resilience control loop: derived from the real-time component of the D^2R^2+DR resilience strategy (see Chapter 1 in [1]).

remain resilient but also economically viable while meeting new demands and providing new services. For instance, during the Covid-19 pandemic, it was reported that telecommunications companies (operators and service providers) were increasing the resources in their networks to cope with significantly increased demands coming from remote meetings in both work and social settings. The increased costs were essential to prevent any possibility of overloading and consequent loss of resilience.

Resilience objectives seem generally in conflict with energy efficiency, as duplicated resources consume extra energy. However, among resilience strategies that have been proposed for recent communication architectures, some concepts are more energy-efficient than expected. Examples include the use of all-optical transport solutions to avoid energy-inefficient signal conversions between the optical and electrical domains, traffic shaping over fewer paths, using renewable energy sources, adaptation of the transmission to dynamic characteristics of traffic, the use of sleep mode for backup paths, and reducing energy consumption by providing differentiated protection levels for different service classes [4].

Progress in the accuracy and efficiency of Artificial Intelligence and Machine Learning models also means that certain operational patterns can be predicted with higher confidence and in adequately short timescales (although there are significant challenges which are discussed in the following sections), hence reactive systems do not need to rely exclusively on expensive resources and hardware redundancy.

Closely related to these issues is a tension between building resilience into networked systems and ensuring their sustainability, which requires continued effort and resources to maintain and develop them over time. However, resilience investment should ultimately pay for itself — failures of an unprotected system can be so expensive in financial and societal terms that even one severe instance could be far more costly than the cost required to provide resilience protection. Thorough risk assessment at the design stage can estimate these respective costs to identify appropriate resilience configurations.

Human, Organizational and Other Factors

An increasingly essential, but little-understood, contextual challenge is the role that non-technical issues can play in both subverting and supporting resilience in complex systems. Real-world computing systems are located within the context of an organization, in which there are people who play a variety of key roles including policy makers, managers, technical staff, and operators. Any or all of the organizational or human elements can play their part in introducing vulnerabilities into the operational system. Operators can, conversely, be sources of support for resilient operation. In the design and operation of resilient systems, it is essential to consider all three elements in combination, the so-called OTI (Organizational, Technological, Individual) approach (see Chapter 32 in [1]). Designers and engineers of complex and critical systems must be trained to adopt an OTI approach. Sustainable systems resilience also needs further multidisciplinary research in sociology, psychology, organizational and management science as well as in engineering, computer science and new and emerging technologies.

Potential Technical Enablers For Improving Networked Systems Resilience

The emergence of new communication technologies and networked system models and approaches brings opportunities and also challenges for how resilience can be realized. In this section, we describe four examples that change the way we can and should approach the realization of network resilience.

Programmability and Virtualization

Software-Defined Networking (SDN) and Network Function Virtualization (NFV) have enabled the decoupling of network services from their physical hosting platform. Both technologies have had a tremendous impact on accelerating research and development for in-network service deployment and functional composition, using for example Service Function Chaining (SFC). The past decade has seen an explosion of research on novel network services and applications focusing on monitoring and reacting to network events, mainly enhancing performance but also assisting resilience. Advances in programmable dataplane technology have enabled the implementation of In-band Network Telemetry (INT) to diagnose network performance through in-band querying and capturing of switch-internal state [5], while protocol-independent programmable switch fabrics have demonstrated high-performance Distributed Denial of Service (DDoS) detection using information-theoretic and statistical analysis for packet classification [6].

However, NFV and SDN technologies introduce additional complexity and can give rise to failure and attack semantics that may be impossible to prevent at design time. For example, the fundamental space/time complexity of the Tuple Space Search (TSS) scheme used by popular packet classification software in hypervisor switches (e.g., Open vSwitch) is subject to low packet rate explosion attacks that have been shown to degrade switch performance by 88 percent [7].

Artificial Intelligence and Machine Learning

A major technology trend in recent years has been the (re-)emergence of Artificial Intelligence (AI) and Machine Learning (ML), which introduces new opportunities (e.g., to enable self-* properties, autonomic operation, challenge detection and

classification tasks) to realize computer network resilience. However, this technology also introduces challenges. For example, ML-based approaches have been used for some time to deal with the prevention of previously unseen cyber incidents and attacks, a known problem faced by signature-based detectors and misuse-based classifiers. But, the inherent low tolerance for false positives, high cost of errors, lack of extensive training data, and non-stationary behavior of Internet traffic, make intrusion detection distinct from standard ML tasks and thus an extremely challenging domain.

Adversarial ML is a subset of evasion attacks and comprises techniques aiming to mislead an ML model to produce incorrect output for a label [8]. An attacker exploits knowledge of an ML model they have direct (read) access to, by subtly altering an input to produce a mislabelling. At the same time, the alteration remains unnoticeable to a human operator. Adversarial ML has been recently gaining research traction with defence mechanisms being regularly proposed only to be then defeated by new ML models [9]. The true extent of the operational risk of this adversarial behavior is arguably not fully understood — further research is required to gain insights into this issue. In fact, adversarial ML requires the development of new approaches that can strengthen models against adversarial attacks. While some progress has been made, a well-established approach including best practises for computer networks protection is still missing.

Reinforcement Learning (RL) brings a different perspective to certain classes of resilience challenges, such as volumetric DDoS attacks, where RL agents that are distributed across a network follow a traffic shaping policy. Recent work [10] leverages an SDN architecture to update RL agents from multiple traces per timestep and perform fine-grained, per-source throttling. This shows a significant increase in goodput of legitimate TCP traffic over highly dense host environments; it also remains protocol-agnostic to offer future-proofing against the rollout of protocols such as Quick UDP Internet Connections (QUIC).

Data and Instrumentation

As shown in Fig. 2, resilience management consists of monitoring (the Detect phase) and control (the use of Remediation mechanisms). The primary sources of monitored data are network traffic, outputs of key networked services, and contextual information such as Cyber-Threat Intelligence (CTI) feeds. Instrumentation and data harnessing mechanisms that will operate as a native part of networked systems are crucial in essential domains such as telecommunications and energy, and in major endeavors such as Smart Cities. Big (indeed vast) data gathering is already feasible in such instances.

Increased Automation

More network traffic, and a greater variety, means operational management should be more highly automated. For example, the increased traffic share that is attributed to machine-to-machine communications produces communication patterns that are different to the user-initiated packet traffic that many networks have been engineered to handle. Sustaining multi-gigabit throughput while handling long bursts of minimum-sized packets is challenging, and resilient network operation will need to be assured by dynamic and adaptive resource provisioning mechanisms [1].

New types of network, including Internet of Things (IoT), edge computing, and space communications, which offer the prospect of improved performance in specific application domains, may present new challenges to realizing system resilience. Moreover, networks that underpin new and highly dynamic applications such as autonomic vehicles in cities and on highways, or mobile robots in factory settings, may need predictive routing in tandem with machine learning for failure modes and resilience mechanisms – combined with optimizing other QoS aspects. These developments also point to the need for autonomic network and services management [11].

Intent-based networking is being developed to allow applications or end-users to request a service and a quality level using an abstracted or policy-level language. This will be mapped into a set of actions in the network, reducing the deployment time and allowing the management system to adapt dynamically to traffic demands. The quality level should be extended to include a resilience parameter or metric [12].

New Application Areas and New Demands

There are new application areas for computer networks that introduce strict service requirements. The criticality of these application areas invites new regulation and law. Together, these act as drivers for network resilience.

Interdependent and Digitalized Critical Infrastructures

Networked systems are becoming increasingly complex and interconnected — and interdependent. Such systems need to be co-designed, underpinned by further research, to avoid cascading effects. This is particularly urgent now that Smart Grids are being developed alongside the introduction of renewable forms of power supply, such as wind and solar farms. In Smart Cities, rules or at least guidelines address the protection of the critical infrastructure in city regions. The resilient (and smart) city has been the subject of many recent global activities, though relatively little effort has been devoted to technological issues, including communications resilience.

In cases where there is autonomic operation, for example in self-driving vehicles, it is essential to develop fast anomaly detection and adequately fast self-healing mechanisms — indeed, for any potentially compromised system where human intervention would be too slow. There are also more stringent QoS requirements — notably latency — on communication networks coming from new application domains including autonomous vehicles, but also remote surgery and medical services, the fast-developing IoT and Industry 4.0, and telecontrol of power plants such as Small Modular Reactors. These applications place increasing demands on the resilience and safety of computer networks.

It is clear that we need new engineering approaches that take into account both the cyber and physical aspects of interdependent systems. The Smart Grid, for example, would benefit from engineering tools and methods using real-time simulators and cybersecurity testbeds (e.g., cyber ranges) to realize secure and resilient engineering

The inherent low tolerance for false positives, high cost of errors, lack of extensive training data, and non-stationary behavior of Internet traffic, make intrusion detection distinct from standard ML tasks and thus an extremely challenging domain.

The RECODIS COST Action formulated three main types of challenge: massive failures driven by natural and environmental disruptions; technology-induced failures; and adversary activity.

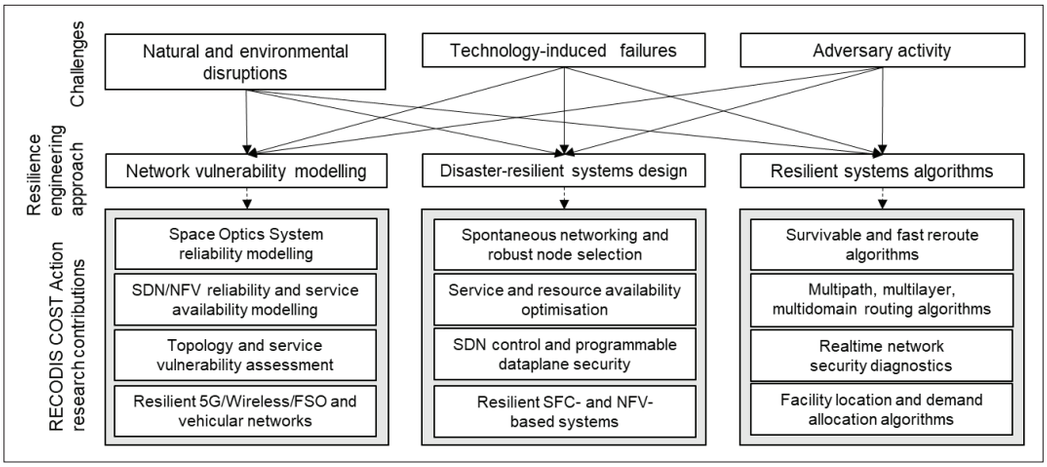

FIGURE 3. Major challenges addressed by the RECODIS COST Action [1].

of distributed control systems and energy services.

New Laws and Regulations

A consequence of new applications will be the emergence of new laws, regulations, and changes in social and operational norms. In some cases, these can be a catalyst to change the resilience properties of computer networks and systems, e.g., to make them compliant with new legal requirements or best practices. This can be achieved with organizational changes and may result in the introduction of new services, e.g., for resilience monitoring, toward ensuring and demonstrating compliance. There have also been changes in best practice and legal requirements with respect to cybersecurity breaches and failures, and the need to disclose information to data subjects and authorities, with the Network and Information Systems (NIS) Directive and the General Data Protection Regulation (GDPR) in the European Union (EU). Their introduction has resulted in significant investment in technology and organizational changes to meet legal requirements.

A Contemporary Viewpoint on Communications Resilience

The RECODIS COST Action developed a taxonomy of environmental and technology-related challenges to communications. In the Action, many teams of researchers pursued solutions to those challenges, and the findings of the project provide a substantial and significant viewpoint on the state of the art in resilience research for communications and networked systems [1]. However, there are inevitably gaps as well as new challenges that have arisen since the project results were published in 2020.

Three main types of challenge were formulated: massive failures driven by natural and environmental disruptions; technology-induced failures; and adversary activities. The main resilience engineering approaches that have been adopted to tackle different challenges together with original research contributions from RECODIS COST are summarized in Fig. 3. Even though significant progress has been made in each of these areas [1], they require more joined-up thinking and an integrated approach. These points are further taken up below.

In particular, in the aftermath of cyber incidents — as well as after natural disasters — it is essential to assure availability of information and the accessibility of network resources, either through algorithmic modelling or system design. Also, the evolution of communication technologies (5G and beyond) requires still greater attention into the effects of atmospheric disturbances at higher frequencies, and further work is needed on assuring QoS during short-term weather disruptions that are becoming more prevalent due to global warming effects.

Meanwhile, technology-induced failures of multiple network elements have been increasing in number, intensity, and scale. Examples include the interdependence between communication and other networks, including power, finance, and transportation. Even a single, significant, failure in one of these networks can cascade, thus resulting in the collapse of many interconnected systems. Another technology challenge tackled within RECODIS was assuring the dependability of virtualized and programmable networked systems; network virtualization can increase resource management complexity but, with proper care, NFV allied with SFC can offer an excellent basis for resilience-by-design of shared networked infrastructures that can be reconfigured and provisioned on-the-fly, in response to adversarial and other events.

The third issue addressed in Fig. 3 refers to malicious human activities causing severe losses at the network and systems level, for instance by disrupting major network links or nodes. It is of course vital to implement resilience mechanisms that will enable fast and effective recovery from cyber attacks at the network or service level while also defending against possible attacks against the physical (hardware) layer. RECODIS made excellent progress in these areas, but much more needs to be done. Specifically, tailored strategies need to be developed for protecting critical infrastructures and systems, using risk and cost-benefit assessments.

Summary of Issues and a Research Agenda

In this section, we discuss open issues and give suggestions for a resilient networked systems research agenda, based on the drivers, barriers and technical enablers (Fig. 4) that have been dis-

cussed in previous sections in this article. Figure 4 also summarizes the essential activities needed to address the open issues in resilience research today, as introduced next.

New enabling technologies, such as virtualization and ML, bring benefits but potential downsides in terms of increased system complexity. These must be investigated along with the management challenges they bring to resilient networked systems [13]. It is important to apply the principles of networked systems resilience (see for example Chapter 1 in [1]) to an increasingly virtualized service space.

Alongside this, further effort is needed to establish usable service level metrics, i.e., measures of how well users are served by the system. Work reported by ENISA [12] still forms a good basis for this research; in that work, a two-dimensional network state space graph shows how a network's operational state will deteriorate in the face of challenges and can then be moved toward a recovery state when appropriate resilience mechanisms are applied. However, research needs to be extended to consider the harder problem of deriving a service-space graph (along with the definition of suitable service metrics).

Also in the service space, an inter-connect of sub-systems could be structured for resilience — including the co-design of interdependent networks — in which operational resilience would be assured via lightweight monitoring and control (i.e., management). Also, a formal model needs to be derived for assurance purposes; this would be an entirely new and important line of research.

Many disciplines and communities have different views of resilience, and the 2021 book edited by Michael Ungar provides a comprehensive coverage of multisystemic resilience [14]. In nature, resilience can occur as a result of evolution, and researchers study such systems to gain insights into how this resilience is achieved. Perhaps the most under-explored research issue is the role of human factors in systems resilience. In human studies, resilience can be observed as an acquired behavior, often as a result of incidents that shape future responses — for example, a child's development will often be shaped in the light of both good and bad life events. In social systems such as cities, resilience to particular events may be shaped by awareness and training such as readiness for evacuation following floods or other serious incidents.

By contrast — in the engineering domain — when designing and building synthetic systems (and specifically networked systems) to offer a continuous service, resilience is a quantitative property. However, it is not only a technological concern; alongside this, designers and implementers must consider the organizational and human elements of these systems. The environment in which an engineered system operates is a crucial part of the design space, where many of the system requirements arise. This outlook is covered in greater detail in Chapter 34 of the multisystemic resilience book [14].

Reinforcing this point, the early pioneering "soft systems" work of Peter Checkland [15] has led to an understanding of the need to model the complex socio-technical aspects of systems engineering, including human and organizational factors. The absence of a systems approach to network resilience can result in solutions that are not well-suited for their deployment context and, therefore, fail to address fully their critical objectives.

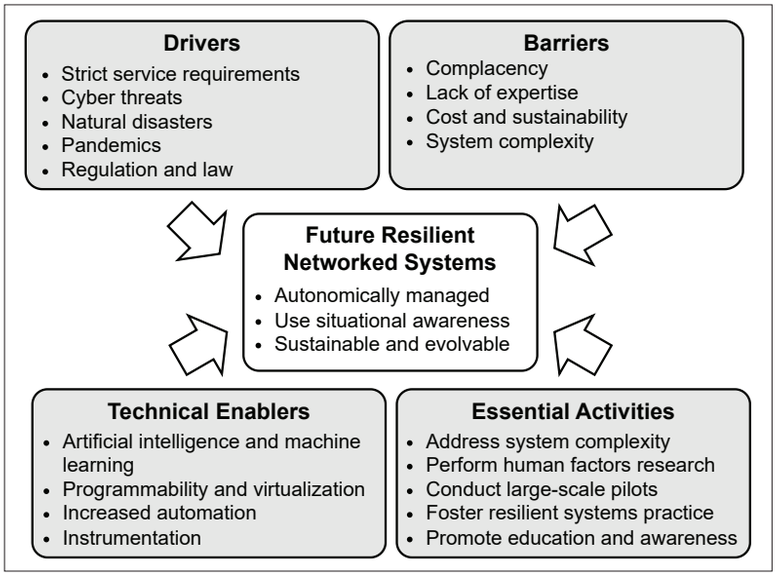

FIGURE 4. An overview of a research agenda for the realization of future resilient networked systems.

Figure 5 summarizes important open issues, including some that apply to deployments of Internet technologies, which may be of interest to the Internet Advisory Board (IAB) and Internet Research Task Force (IRTF). The figure is not intended to be complete — rather, this is a starting point aimed towards a resilience research and future study agenda.

It is timely and important to launch practical projects (preferably large-scale pilots) that implement resilience engineering and management using, for example, the D^2R^2+DR or a similar framework at a suitable scale — and in the field rather than in a research laboratory. Only then can we assess the soundness, cost-effectiveness, and utility of the resilience principles and small-scale experiments that have thus far been presented in many research papers. It is very challenging to understand the many systems aspects without suitable pilots that deploy resilience technologies in context. Such pilot projects would be the most effective way of promoting the importance and the value of resilience thinking and practice.

Conclusion and Reflections

Given society's increasing dependence on networked systems, it is crucial that they remain resilient in an ever-changing environment — whether in response to degraded performance due to resource demands or cyber challenges to the infrastructure. To maintain adequate levels of service at all times, network systems designers need to think about resilience and devise strategies for seamless self-adaptation of services to changing operational conditions.

Applications and traffic instrumentation with always-on measurement capability will need to be developed to implement adaptive resource allocation with support from the network infrastructure. Metrics that are able to capture complex operational properties of networked systems

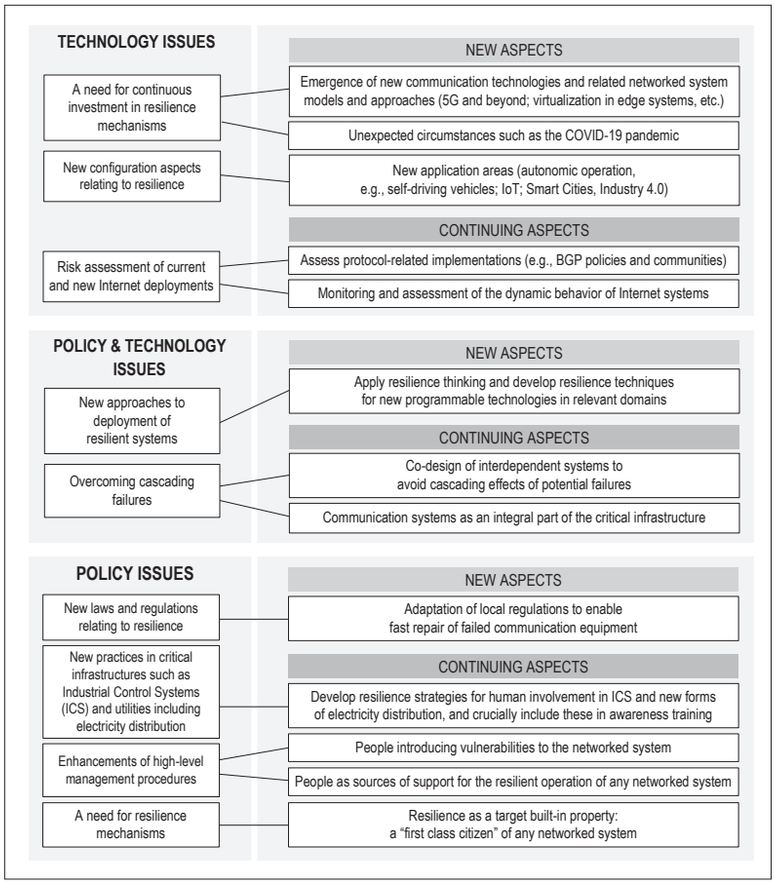

FIGURE 5. Resilience engineering: open issues.

will need to be standardized and subsequently incorporated in service level agreements to assure deterministic or statistical levels of resilience for networked systems and services. The complexity of an ever-changing environment also points to the need for dynamically verifiable software-defined systems that remain trustworthy despite increasingly autonomous operation, and these systems must not themselves become easy targets for cyber attacks.

Resilient systems engineering is absolutely essential in the modern world. Therefore, it makes sense for resilience to become an integral part of a system design brief. Finally, multidisciplinary research is clearly needed in this vital area.

References

[1] J. Rak and D. Hutchison, Eds., *Guide to Disaster-Resilient Communication Networks*, Springer, Cham, 2020.

[2] ISO/IEC 13236:1998 Information Technology — Quality of Service: Framework, ISO, 1998.

[3] CISA 5G Security and Resilience, 2022; available: https://www.cisa.gov/topics/risk-management/5g-security-and-resilience.

[4] Y. Ye *et al.*, "Energy-Efficient Resilient Optical Networks: Challenges and Trade-Offs," *IEEE Commun. Mag.*, vol. 53, no. 2, 2015, pp. 144–50.

[5] C. Kim *et al.*, "In-Band Network Telemetry via Programmable Data-Planes," *Proc. 2015 ACM Conf. SIGCOMM*, ser. SIGCOMM'15, 2015.

[6] A. Ilha *et al.*, "Euclid: A Fully In-Network, P4-Based Approach for Realtime DDoS Attack Detection and Mitigation," *IEEE Trans. Network and Service Management*, vol. 18, no. 3, 2021, pp. 3121–39.

[7] L. Csikor *et al.*, "Tuple Space Explosion: A Denial-of-Service Attack Against a Software Packet Classifier," *Proc. 15th Int'l Conf., CoNEXT '19*, 2019, pp. 292–304.

[8] N. Papernot *et al.*, "SoK: Security and Privacy in Machine Learning," *Proc. 2018 IEEE European Symposium on Security and Privacy*, 2018, pp. 399–414.

[9] F. Tramèr *et al.*, "Stealing Machine Learning Models via Prediction APIs," *Proc. 25th USENIX Security Symposium*, 2016, pp. 601–18.

[10] K. Simpson *et al.*, "Per-Host DDoS Mitigation by Direct-Control Reinforcement Learning," *IEEE Trans. Netw. Serv. Manag.*, vol. 17, no. 1, 2020, pp. 103–17.

[11] White Paper No. 16: Reference Model for Autonomic Networking, Cognitive Networking and Self-Management of Networks and Services. ETSI, 2016.

[12] "Resilience Metrics and Measurements: ENISA Technical Report," 2011; available: https://www.enisa.europa.eu/publications/metricstech-report.

[13] R. Boutaba *et al.*, "Managing Virtualized Networks and Services With Machine Learning," *Commun. Networks and Service Management in the Era of Artificial Intelligence and Machine Learning*, 2021, pp. 33–68.

[14] M. Ungar, Ed., *Multisystemic Resilience: Adaptation and Transformation in Changing Contexts*, Oxford University Press, 2021.

[15] P. Checkland, *Systems Thinking, Systems Practice: Includes a 30-Year Retrospective*, Chichester: John Wiley and Sons Ltd, 1999.

Biographies

DAVID HUTCHISON is an Emeritus Professor of Computing at Lancaster University, UK, and the Founding Director of InfoLab21 in the School of Computing and Communications. His work is well known internationally for contributions in a range of areas including Quality of Service, active and programmable networking, multimedia and content distribution networks, and testbed activities. His current research focuses on the resilience of networked computer systems, and the protection of critical infrastructures and services. He was Vice Chair of the RECODIS COST Action.

DIMITRIOS PEZAROS [SM] is (full) Professor of Computer Networks in the School of Computing Science at the University of Glasgow, where he currently holds the Royal Academy of Engineering Research Chair in Digital Resilience for Critical National Infrastructure. He has published widely in the areas of computer communications, network and service management, and network resilience. He holds BSc and PhD degrees in Computer Science from Lancaster University. He is a Chartered Engineer, a fellow of the BCS and the IET, and a Senior Member of the ACM.

JACEK RAK [SM] is a professor and the head of the Department of Computer Communications at Gdańsk University of Technology, Poland. From 2016 to 2020, he was leading the EU FP7 RECODIS COST Action focusing on the disaster-resilience of communication infrastructures involving over 170 members from 31 European countries. His research interests include the resilience of communication networks and networked systems. He is the founder of RNDM — Workshop on Resilient Networks Design and Modeling.

PAUL SMITH is a Senior Scientist with the Center for Digital Safety and Security at the AIT Austrian Institute of Technology and a Visiting Researcher at Lancaster University, UK. He received his PhD in Computing from Lancaster University in September 2003. Dr Smith's research interests focus on the security and resilience of cyber-physical networked systems. Dr Smith is the co-chair of the OCG Working Group on Network Intelligence (NET-IT), represents the AIT at the European Energy ISAC, and is a member of the ACM.