

Disaster-Resilient Communication Networks: Principles and Best Practices

Andreas Mauthe*, David Hutchison*, Egemen K. Çetinkaya†, Ivan Ganchev‡, Jacek Rak§,
James P.G. Sterbenz¶*, Matthias Gunkel||, Paul Smith**, and Teresa Gomes††

*School of Computing and Communications, Lancaster University, United Kingdom

†Department of Electrical and Computer Engineering, Missouri University of Science and Technology, USA

‡Telecommunications Research Centre (TRC), University of Limerick, Ireland and

Department of Computer Systems, Plovdiv University "Paisii Hilendarski", Plovdiv, Bulgaria

§Faculty of Electronics, Telecommunications and Informatics, Gdansk University of Technology, Poland

¶Information and Telecommunication Technology Center, Department of Electrical Engineering and Computer Science
The University of Kansas, Lawrence, KS, 66045, USA

||Fixed Mobile Engineering Deutschland, Optical Packet Transport, Deutsche Telekom Technik, Germany

**Digital Safety and Security Department, AIT Austrian Institute of Technology, Austria

††Dept. of Electrical and Computer Engineering & INESC Coimbra, University of Coimbra, 3030-290 Coimbra, Portugal

Email: *{a.mauthe, d.hutchison}@lancaster.ac.uk, †cetinkayae@mst.edu, ‡ivan.ganchev@ul.ie, §jrak@pg.gda.pl,
¶jjpgs@itc.ku.edu, ||GunkelM@telekom.de, **p.smith@ait.ac.at, ††teresa@deec.uc.pt

Abstract—Communication network failures that are caused by disasters, such as hurricanes, earthquakes and cyber-attacks, can have significant economic and societal impact. To address this problem, the research community has been investigating approaches to network resilience for several years. However, aside from well-established techniques, many of these solutions have not found their way into operational environments. The RECODIS COST Action aims to address this shortcoming by providing solutions that are tailored to specific types of challenge, whilst considering the wider socio-economic issues that are associated with their deployment. To support this goal, in this paper, we present an overview of some of the foundational related work on network resilience, covering topics such as measuring resilience and resilient network architectures, amongst others. In addition, we provide insights into current operational best practices for ensuring the resilience of carrier-grade communication networks. The aim of this paper is to support the goals of the EU COST Action RECODIS and the wider research community in engineering more resilient communication networks.

Index Terms—resilience, disaster-based disruptions, software defined networks

I. INTRODUCTION

Disaster-based failures can seriously disrupt a communication network, making its services unavailable. Such disruptions may follow from natural disasters, technology-related failures, or malicious attacks. These disruptions are observably increasing in number, intensity and scale. The problem needs to be urgently addressed, due to the lack of suitable mechanisms deployed in current networks. When network services that are part of a critical infrastructure become unavailable, commercial and/or societal problems are the inevitable result.

To address this issue, network resilience – *the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to*

normal operation [1] – has received a significant amount of research attention in recent years. Research has focused on a wide-range of topics, including measuring the resilience of networks, multi-level network architectures that facilitate the engineering of resilient networks, and a multitude of technologies that address specific challenges, e.g., for wireless networks. However, for various reasons, many of these solutions have yet to see operational deployment.

In response to this, the EU COST Action *Resilient communication services protecting end-user applications from disaster-based failures (RECODIS)*¹ aims to develop solutions to provide resilient communications in the presence of disaster-based disruptions of all types for existing and future communication network architectures. To support their future operational deployment, this undertaking aims to consider the wider socio-economic environment to which the solutions will be deployed. The Action has been organised into working groups that explore solutions to specific types of challenge – large-scale natural disasters, weather-based disruptions, technology-related disasters (e.g., blackouts), and malicious human activities.

As a means of supporting the research that will be undertaken in RECODIS, and the wider research community, we present important network resilience principles. The aim is to provide a normative framework that will support our joint research activities. In this regard, in order to establish a common understanding about our overall goal, we propose a definition for communication network resilience. This definition specifies that resilience is a quantifiable measure – therefore, it is important to establish frameworks for measuring

¹http://www.cost.eu/COST_Actions/ca/CA15127

network resilience. In Sec IV, we present research findings on this important topic. In addition, to clearly understand the nature of the challenges that communication networks face, in Sec. III, we provide an overview of related work on challenge taxonomies. A number of solutions for network resilience will be investigated as part of RECODIS's activities; in order to understand how they can be used as part of a larger resilient network architecture and form a systematic resilience strategy, architectural principles are presented in Sec. V. Finally, in Sec. VI, we present best practices for network resilience that are employed by a large telecoms operator for carrier-grade networks. The aim is to identify the current state of practice of operators, including the types of technologies they use and are considering, that the solutions which will be proposed by RECODIS must interface with.

II. COMMUNICATION NETWORK RESILIENCE DEFINITION

Beyond the definition for network resilience quoted at the beginning of this paper (developed as part of the EU FP7 ResumeNet project [2] and ResiliNets initiative [3], and adopted by ENISA [4]), there are many definitions of resilience. These range from short ones (e.g., the capacity to recover quickly from difficulties, as defined in the Oxford dictionary) to quite long ones. Some definitions consider resilience as an ability, capacity or property, whereas others consider it as a process. Some pertain to a particular area or discipline; others take into account different (multidisciplinary) aspects of it. The latter approach seems more reasonable to follow, in order to elaborate one common definition of resilience for use in a particular domain.

In short, the term resilience – derived from the Latin *resalire* – means to spring back. It is used in multiple disciplines, ranging from psychology to physical sciences and ecology, and finishing up with engineering, where it relates to the concept of being able to absorb and recover from hazardous events and disasters. A very good analysis of the more widely used definitions of resilience, relevant to communities, is provided in [5], based on five core concepts – attribute, continuing, adaptation, trajectory, and comparability. That report also points to different ways to classify the definitions, e.g., by contrasting “being vs becoming” or “adaptation vs resistance,” or in terms of trajectory or predictability, or by taking into consideration the temporal nature of resilience. Hybrid definitions of resilience also exist, based on different combinations, e.g., of engineering with ecology, or ecology with the behavioural science [5]. A good collection of the ‘best’ ten definitions of resilience (with respect to ecology and society) can be found in [6]. Some of those definitions are based on the descriptive concept, whereas the second group follows the hybrid approach, and the third group uses the normative concept.

Of course, no perfect definition exists. The same is true for one commonly accepted definition of resilience that could be used across all disciplines. The best way, perhaps, is to select some of the (best) existing definitions with a good applicability to a particular domain, and try to compile a comprehensive definition of resilience for that domain. An attempt to do this

for the domain of communication networks was made within the framework of the EU COST Action RECODIS. As a result, the following (long) definition was elaborated:

Resilience of a communication network is a quantitative property of a network that occurs on each level of its hierarchy, and is related to the ability to maintain the same level of functionality in the face of internal changes and external disturbances as a result of large-scale natural disasters and corresponding failures, weather-based disruptions, technology-related disasters, and malicious human activities; to withstand all these without losing the capacity to allocate resources efficiently; to maintain acceptable level of service in the face of various faults, challenges to normal operation, fluctuating environment, and human use; and to absorb recurrent disturbances so as to retain essential infrastructures and processes, with sufficient cost-efficiency and flexibility over the long term.

This could be reduced to the following short definition, for practical use:

Resilience of a communication network is its ability to maintain the same level of functionality in the face of internal changes and external disturbances as a result of large-scale natural disasters and corresponding failures, weather-based disruptions, technology-related disasters, and malicious human activities.

An important component of these definitions is the ability to quantify and measure resilience – we present an overview of frameworks for measuring resilience in Sec. IV. In addition, it is important to understand the nature of the challenges, i.e., internal changes and external disturbances, that could affect a communications network; we present a taxonomy of challenges in the following section.

III. TAXONOMIES OF CHALLENGES TO RESILIENCE

End-to-end communications faces a number of challenges that may result in unsuccessful delivery of information. A challenge is *an adverse event or condition that can cause deviation of normal network operation* [7]. The proper recognition of these network challenges and the corresponding impact on networks is crucial so that appropriate planning and measures can take place. We note that, while a taxonomy of the faults is detailed in IFIP 10.4 group documents [8], the taxonomy of challenges is the focus of this work. Challenges are the abnormal events that can trigger the faults and errors, eventually causing system failure [7]. Network challenges can be categorized based on a number of criteria including: cause (natural, human-made, or challenge-dependent), boundary (internal or external), target (direct or collateral), objective (non-malicious, selfish, or malicious), intent (non-deliberate or deliberate), capability (accidental or incompetence), dimension (hardware, software, protocols, or traffic), domain (medium, mobility, delay, or energy), scope (nodes, links, or area), significance (minor, major, or catastrophic), persistence (short-lived, long lived, or transient), as well as repetition (single, multiple, or adaptive) [7]. A taxonomy of major challenges is shown in Figure 1.




Fig. 1. Taxonomy of major challenges (based on [7], [9])

Large-scale disasters can be result of forces of nature, including: earthquakes (e.g., the 2006 Taiwan earthquake [10], the 2008 Wenchuan earthquake [11], the 2011 Japan earthquake [12], etc.). Other causes of disaster-based failures are, e.g., hurricanes (for instance Katrina [13]) responsible for remarkable disruptions of communication links/hardware (nodes). Natural disasters also can be associated with cosmological events (for instance geomagnetic storms [14]). Human-made disasters can either follow from ignoring early warnings in the operation of a system or be caused by malicious activities.

Socio-political & economical challenges refer to deliberate activities (also acts of terrorism [15]) that are prepared to disrupt normal operation of a communication network (e.g., to achieve advantage on economical markets [16], or as a response to political decisions [17]).

Dependent failures refer to challenges that may result in a cascade of failure. Canonical examples include the power grid and the Internet dependent failures [18]. Another dependent failure scenario is a Border Gateway Protocol (BGP) cascading failure (e.g., invalid BGP advertisement propagation), in which BGP routing relies on the correct announcement of these advertisement messages [19].

Human errors are deliberate or non-deliberate activities. We can mention here that misconfiguration errors are the result of human incompetence (which can even lead to catastrophic failures) [19].

Malicious attacks are with a deliberate intent designed to cause as much disruption as possible. Such activities are commonly targeted at important software/hardware elements of the network [20].

Unusual traffic affects the network traffic, and may cause the end systems to be unresponsive as in the case of a flash crowd [21]. Traffic volumes that deviate significantly from the normal expected traffic can be a problem, if the capacity of the network resources are not provisioned to handle the overload traffic. Unusual traffic can occur after the occurrence of a catastrophic event (e.g., news web servers became overwhelmed after the 9/11 terror attacks [22]). Such an event does not necessarily disrupt the network infrastructure itself, but may result in a significant increase in the number of simultaneous requests to get information.




Fig. 2. The ANSA dependability framework

Environmental challenges are greatly influenced by characteristics of the communication environment (e.g., mobility, delay issues in wireless networks [23]).

Any challenge can also be characterized based on its spatial region and time duration. The spatial and temporal features of a challenge might be different, including its *spatial and temporal impact* [7]. For instance, concerning spatial region, an attack being a challenge related to a single node may influence the performance of the entire network. Similarly, time duration of, e.g., an earthquake (seconds) differs significantly from the duration of impact (here in days).

Finally, we note that as new challenges are explored, this taxonomy can be adapted to include new aspects. The investigation of challenges in emerging technologies (e.g., Software-defined Networks (SDNs), Network Functions Virtualization (NFV), and Cloud) is a research direction to understand and develop mechanisms against challenges [24].

IV. MEASURING NETWORK RESILIENCE

There have been a number of frameworks proposed for measuring resilience, survivability, and dependability. Dependability [25] measures, such as reliability (the probability that a system will remain operational for a specified period of time) and availability (the probability that a system is up at a particular point in time), as well as performability [26] that measures the degraded performance of a complex system such as the Internet, can be used to characterize the resilience (and survivability) of communication networks.

ANSA [27] was a distributed systems project that included a dependability framework, as shown in Fig. 2. Occurrences of value tuples (with only 1 value dimension shown here) are measured over time. *Expectations* can include a value at a given time, a value over time e_1 , a range of values over time e_2 , and a range of values that changes over time e_3 . *Occurrences* are measured with expectations; correct operations consists of occurrences corresponding within time and value of expectations (o_1, o_2, o_3). Failures are measured when occurrences do not meet expectations, for example o_u .

A couple of notable frameworks to measure survivability against correlated failures include a state-based approach [28], in which various events move the system between operational states, and a two-dimensional Markov chain [29], in which one dimension measures the number of failures (and repair)

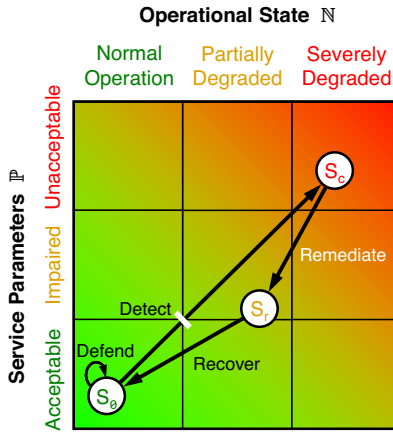


Fig. 3. ResiliNets metrics framework




Fig. 4. T1A1 and ResiliNets temporal metrics




Fig. 5. The ANSI/ATIS T1 service outages model [32]

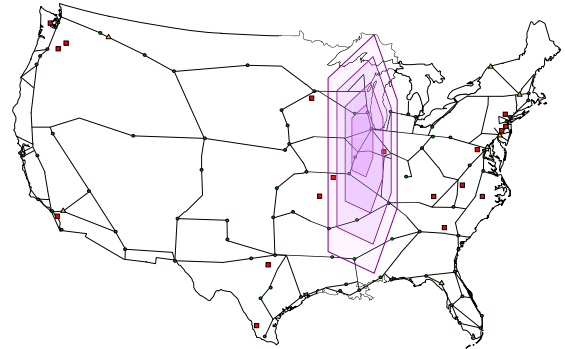


Fig. 6. Midwest US power blackout model

and the other dimension measures conventional arrivals and service.

ResiliNets uses a two-dimensional state space to measure resilience, as shown in Fig. 3 [30], [31]. The horizontal axis measures an objective function of the network operational state; a resilient infrastructure maintains normal operation in the face of challenges. The vertical axis is an objective function of service metrics; a resilient service remains acceptable *even when* the network degrades. The ResiliNets $D^2R^2 + DR$ (Defend, Detect, Remediate, Recover + Diagnose, Refine) resilience strategy [1] is overlaid. As long as the system remains in normal operation with acceptable service S_0 , *defences* have held. If either network or service monitoring *detects* a deviation to a challenged state S_c , *remediation* mechanisms are invoked, driving to a better state S_r . Eventually, when infrastructure has been repaired or replaced, the system can *recover* to its original state. Resilience is measured as $1 - (\text{area under state trajectory})$. The smaller the area, the higher the resilience, either because the triangle is narrow (infrastructure resilience) or shallow (service resilience). After *diagnosis* and *refinement* evolves the network and improves D^2R^2 , the loop will be tighter giving a smaller area and better resilience.

ANSI/ATIS T1 has modelled service outages as a triple (U, D, E) [32] where U is unservability, D is duration, and E is extent. Events are quantified as *minor*, *major*, or

catastrophic based on the magnitude, as shown in Fig. 5. A large-scale disaster would be catastrophic in this classification, and have a large surface when plotted in three dimensions of (U, D, E) . Furthermore, a temporal view (based on [32] with $D^2R^2 + DR$ overlaid) is shown in Fig. 4. In normal operation, performability $P(t)$ is one. A challenge reduces performability by P_u to P_a . After detection and remediation t_r , performability is increased by P_r . Eventually, as infrastructure is repaired and replaced after t_R , performability returns to 1.

To understand the effect of large scale disasters [31] as well as to facilitate identification of critical region vulnerability [33], ResiliNets uses an ns-3 simulation-based methodology, as well as graph-theoretic and optimisations to model various challenges. These can be circular (as would be the case for CME – solar coronal mass ejections and EMP – electromagnetic pulse weapons), polygonal (as would be the case for power blackouts as exemplified in Fig. 6), and may move (as would be the case for hurricanes and typhoons).

An arbitrary challenge can be applied to any real or proposed network topology, which can then be modelled to understand the impact. This is a *multilevel model*, in which the physical infrastructure failures are propagated up to the network topology level, propagated up to the routing level, propagate up to the end-to-end transport level, and finally up to the application level [34]. Building on such multi-level resilience metrics, an analysis can be carried out that assess the

resilience of a network across different layers. This analysis for instance helps to highlight structural vulnerabilities. In the design phase of a network they can be used to evaluate alternative architectures or configurations. For instance, using metrics such as topology diversity and connectivity, path routing diversity and connectivity, etc. can be analysed to demonstrate the effect failures of individual components or links have on the overall network resilience.

How this can be used in practice to assess the structural resilience of a specific network is demonstrated in [30], [1] using the GEANT2 topology. This analysis also shows what impact a power-grid failures on the resilience of an ISP topology at different levels can have. In order to measure resilience a two dimensional state space is used in which resilience is expressed in a range from (0,1) where 0 is representing no resilience, and 1 is representing infinite resilience. The two dimensions according to which the resilience is assessed are *operational state* (ranging from normal to severely degraded) and *service parameters* (ranging from acceptable to unacceptable). Each of the dimensions is associated with objective functions that allow to measure resilience. A challenge can cause a change in the resilience state that is then measurable in a degraded of the overall resilience value. In the context of multi-level resilience the service parameters at the layer boundary becomes the operational metric of the layer above, i.e. the impact of a resilience challenge at one level is reflected in the degraded operational state provided to the service user.

A formal multi-level graph model and framework for multi-level graph evaluation is developed in [34], which is used to analyse the resiliency of single and multi-level graphs, using the flow robustness metric (the fraction of node pairs that remain connected after a number of removals [35]). Then Çetinkaya et al. [34] define the concept of multi-provider graph to represent the inter-provider AS (autonomous system) topology, and analyse the flow robustness in three different multi-provider graphs under different types of challenges (random and targeted).

Another example is presented in [36], where the robustness of a network is being analysed by considering multiple network levels using a graph theoretical approach. A computational framework for network resilient is being developed through which the impact of challenges can be assessed. In order to do so the network topology is captured in form of a graph and then assessed using a graph explorer approach considering multi-level network metrics. With this approach the entire state space of a communication system is analysed and a risk map can be developed. This risk map captures different system levels by, for instance, looking at how application level performance is affected through specific challenges to various network components.

V. ARCHITECTURAL CONCERNS FOR RESILIENCE

In computer networks, resilience is an *infrastructure and management* property. Therefore, network architectures have, on the one hand, to display a certain degree of structural resilience that enables them to withstand and compensate

for any potential malfunction in the case of attacks, outages or other challenges. On the other hand, future network architectures also have to have active resilience management components that enable autonomic detection and protective actions, once challenges have been identified. Hence, the resilience of networks has to be ensured at the *structural* and the *operational* level.

At a structural level, resilience is achieved through means such as replication and diversity, which can be implemented in a complementary manner at different system levels. For instance, at the network layer, a number of alternative routes exist through which packets can be forwarded, and (ideally) routers and equipment will be from different vendors running alternative implementations of routing protocols. This strategy allows a network or system, in the event of failures, to utilise alternative resources in order to maintain the service.

Meanwhile, at the operational level, resilience is achieved through active detection, remediation and recovery actions (see the ResiliNets $D^2R^2 + DR$ resilience strategy [1]). Anomaly detection can be performed throughout the system at the application layer, edge network as well as in the network core itself. Appropriate remediation strategies then help to ensure continued operation, even during situations where there are challenges. Thus, a network architecture that is designed according to resilience principles has to provide both – *structural* and *operational* resilience.

In keeping with the traditional separation of concerns (as expressed through layering principles) network resilience architectures also have to consider structural and operational resilience across the different communication layers. It is recognized that cross-layer analysis and information flows are required, in order to ensure the effectiveness of resilience operation. The provisions of so-called Multi-Level Resilience (MLR) as a property of network architectures has two main aspects: *Structural Multi-Level Resilience* (SMLR), as mainly expressed through frameworks and architectural models and the assessment of the resilience level they offer (e.g., [30]), and *Operational Multi-Level Resilience* (OMLR), i.e., providing coordinated resilience mechanism across system layers and even system boundaries.

It is important to note that a resilient network architecture is based on the concept of autonomic components that have large degrees of self-organisation. These autonomic components are not only adaptable to the environment, but can also evolve. Collaboration between them largely happens through information exchange and the adoption of joint policies and rule sets.

A. Structural Network Resilience

In order to deal with challenges in a more systematic manner, a number of resilience principles have been defined as part of a resilience framework [1]. These principles help to systematically address the different areas and concepts that are related to the resilience domain, considering the *prerequisites*, *trade-offs*, *enablers* and *behaviour*. An overview of these principles for resilience are presented in Fig. 7.




Fig. 7. An overview of the ResiliNets resilience principles [1]

Regarding the key enablers, *redundancy* refers to the replication of components, i.e., hardware (such as switches and routers), as well as communication links (e.g., the availability of multiple alternative paths between source and destination). The general idea is that redundant components can pick up tasks of a failing entity. This principle can be applied at all layers of the communication stack. *Diversity*, in addition to redundancy, helps to avoid challenges of the same kind affecting (homogeneous) system components in the same way. For instance, hardware components from different suppliers are less likely to suffer the same faults, or software that is developed by different software engineers should also not contain the same programming errors. Hence, diversity helps to improve resilience, if the diverse components provide an equivalent level of service, while being structurally and operationally different. *Connectivity* and *association* refer to the continued ability for information exchange between system entities, even in the case of disruptions.

More specifically, for example, in [37], [38] the Spine concept is introduced as a subnetwork structure that is embedded at the physical layer with higher availability. The spine concept allows the creation of heterogeneous availability subnetworks at the physical layer, laying the foundation for service differentiation at upper layers. In [39] it is used to provide differentiated services over multilayer communications networks. Multiple logical networks with diverse availability levels are defined via a cross-layer mapping. Results in [39] showed the proposed model can create a wider range of availability levels compared to existing techniques, although in some cases the spine requires slightly more resources.

B. Operational Network Resilience Support

At the operational level, resilience has to be provided in conjunction with network management functions. There are two main elements that need to be part of this resilience management process – challenge detection and remediation (alongside recovery). For the first task, network state is assessed by anomaly detection components that are distributed throughout an interconnected network infrastructure. Detection components can sit alongside the routing infrastructure or be part of it.

In addition to the coordinated detection of challenges, other aspects of resilience have to be carried out in a coordinated manner throughout the interconnected network infrastructure. To realise this, an approach that makes use of policies, in conjunction with remediation mechanisms that can be progres-




Fig. 8. A multi-stage approach to network resilience [40]

sively refined, has been proposed [41], [42]. The idea is that an initial indication of an anomaly triggers a staged network management and resilience process, in which the diagnosis is refined in order to gain more certainty about the challenge. Based on this, remedial actions are also refined, in order to become more targeted. This concept is depicted in Fig. 8.

For instance, during the onset of a Distributed Denial of Service (DDoS) attack, the first indication will be an increase in traffic on a specific network link. Once a certain threshold has been reached, remediation actions have to be taken in order to protect the network resources. At this stage, the cause of the anomaly might not be clear and hence further analysis helps to establish the exact cause (e.g., a DDoS attack or flash crowd), and also to instantiate at each stage appropriate countermeasures (ranging from initial rate limiting the link to selectively dropping of packets from offending sources). How this scheme can be realised in the context of Software-defined Networks (SDN) is shown in [43], using a policy pattern based control and management scheme.

VI. NETWORK RESILIENCE: OPERATOR BEST PRACTICES

In this section, we present current best practices for resilience in carrier networks. The aim of the EU COST Action RECODIS is to be enhance these capabilities, in order to improve the resilience of such networks (and others) to large-scale disasters. In general, today’s operated optical transport networks provide unprotected or protected services from 1 Gbit/s up to 100 Gbit/s to destinations across the world².

When we refer to protected services, resilience usually follows a 1+1 dedicated path or link protection over topologically disjoint parts of the network. Failure impact is quasi-instantaneously remediated. The reaction can be carried out directly on the data plane, without involvement of the control plane. Typically, protection-based reactions are preconfigured in order to allow for a fast failover, typically within the 50 milliseconds range.

The restoration process, in contrast, takes a comparably long time. Indeed, the recovery mechanism itself starts immediately after failure detection. However, the process might include communication with a controller instance. This controller re-computes the best lightpath guidance through the network. Traffic engineering and optimization routines might decide on modifications of existing end-to-end paths. Furthermore,

²See Deutsche Telekom’s international carriers’ carrier DWDM network as an example: <http://www.telekom-icss.com/lambdaconnect>



Fig. 9. Deutsche Telekom's international high-quality DWDM network

the controller defines the sequential recovery order of all affected lightpaths. Accordingly, each path is signaled from the source to the destination node, and then the bandwidth is re-provisioned to fulfil the given service level agreements. Usually, there is no accurate or obligatory time limit for such a recovery mechanism, besides that it needs to be done as fast as possible.

A. Multi-layer Resilience with a Flexible Optical Layer

In a multi-layer framework, recovery needs to be accomplished on multiple involved layers. The most relevant work [44], [45] exclusively targets two layers – the optical and the IP packet layer – which have to react to a failure in a coordinated way.

At the optical layer, today's long-haul 100Gbit/s transceivers usually operate with Dual-Polarization Quaternary Phase-Shift Keying (DP-QPSK) modulation, and have a reach limitation of about 2500km. This is long enough for most applications in national European core networks and enables even a transparent backup lightpath connection. Recent progress in optics and electronics has enabled a significant increase in symbol rates and modulation technology. Therefore, near-future high-speed transceivers can also operate at DP-M-QAM formats. Such a flexrate transceiver is reported which supports net bit rates from 100 to 400Gbit/s in steps of 25Gbit/s and generates modulation formats from DP-QPSK to DP-64QAM [46]. Flexrate transceivers offer many advantages, such as less stock holding costs and improved network usability.

However, this comes along with a fundamentally reduced transmission reach. In case of a cost-efficient optical restoration solution, the recovered lightpath length usually exceeds the original reach which might overburden the previous settings of the flexrate transceivers at both sides of the lightpath. Then, a transport network operator has two options: either to shut-down the optical link entirely or to squeeze out the capabilities of affected transceivers in the best possible way, i.e., to reduce their transport capacity down to a value that is applicable for the restored link conditions.

This best-effort optical capacity poses new questions from the packet layer perspective. Traditional IP/MPLS networks

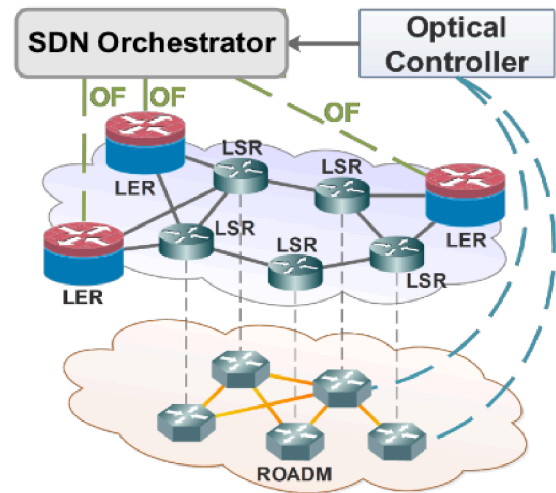


Fig. 10. A high-level view of a hierarchical SDN control architecture, including the main components

rely on fixed capacities provided from the optical layer. IP networks prefer an entire interface shut-down with no capacity to a partial recovered capacity. Existing load balancing mechanisms such as Equal-Cost Multi-Path (ECMP) are not capacity aware, i.e., they split the IP traffic load onto all involved interfaces likewise. As ECMP with unequal optical capacities may lead to a degradation of the Quality of Service, operators prefer to switch off a partially recovered capacity. Of course, from an overall network perspective this is counterproductive.

B. Resilience Management with Software-defined Networks

Motivated by this shortcoming, a new SDN-based solution was investigated in [45] for controlling the routers such that volatile interface capacities can still be utilized and the overall multi-layer network cost can be reduced.

The control architecture is depicted in Fig. 10. Here, the logically centralized SDN orchestrator communicates through an OpenFlow interface directly with the Label Edge Routers (LER). This way, it monitors the traffic flows as well as the overall network state with respect to the different internal routing protocols. Furthermore, it reads the optical network topology from the optical controller. When no failure exists in the network, the SDN orchestrator is passive and does not interfere with the routing decision of the IP/MPLS layer.

When an optical link has failed and is already restored, its potentially reduced capacity is signaled by the optical controller to the SDN orchestrator. Using this information as well as the state of the core network, the orchestrator is able to calculate globally optimized multi-layer decisions. It might selectively override the routing decisions made by the MPLS layer to account for the capacity reduction of the recovered optical connections.

Figure 11 shows a small sample network for illustration. The operation principle is visualized by an OpenFlow rule, though in real deployments orchestration might be based by any other appropriate protocol, such as Netconf/YANG, Restconf




Fig. 11. Symmetric load balancing (done by, e.g., ECMP) in case of normal network operation




Fig. 12. Asymmetric load balancing in case of a network failure

and YANG or REST. The ECMP mechanism distributes the 500Gbit/s outgoing traffic of LSR 51 evenly on its three outgoing 200Gbit/s links to LSR 52, 53, and 54.

Due to a link failure, the bandwidth on the link between LSR 51 and 52 is assumed to be reduced to 100Gbit/s. The bottleneck is detected by the SDN orchestrator and a new path is calculated, see Fig. 12. The OpenFlow group of type select is extended to three buckets, one for each outgoing link on LSR 51. Those type select OpenFlow groups assign matching packets to one of the buckets by hashing selected header fields. Each bucket applies a node label stack to passing packets that forces them through a different outgoing link of LSR 51. By assigning appropriate weight values to all buckets, traffic can be distributed just matching the remaining optical capacities.

After this procedure is applied on each of the selected flows, the network is back in a balanced state and the SDN orchestrator signals that the optical links should be activated again. The packet layer restoration is completed. When the optical link has been repaired, the updated link capacity is signaled to the SDN orchestrator, which in turn removes the redirection rules. At this point in time, the network is reverted back to the operation mode and is controlled by the proven MPLS protocol again. This resiliency solution overcomes the limitations of current ECMP-based mechanisms by an asymmetric multipath routing, enabled by a customized OpenFlow-based Segment Routing approach. No further network equipment renewal is required, except an appropriate software interface and its implementation under an SDN orchestration environment is

needed.

VII. CONCLUDING REMARKS

The resilience of communication networks to disasters is of the highest societal and economic importance. This is becoming increasingly the case as Information and Communication Technologies (ICT) play a more central role in the operation of our critical infrastructures, such as transportation and power systems. The international research community has investigated a wide variety of approaches to improving network resilience. However, a significant amount of this research has not made it into operational environments. There are various reasons for this, but perhaps a major factor relates to a lack of consideration of the broader socio-economic context that solutions are being placed into. Specifically, there is a trade-off that needs to be made with regard to potential improvements in resilience (to *black swan* events) and the cost of implementation of solutions, such as technologies, organisational processes, and (regulatory) frameworks. The RECODIS COST Action aims to propose solutions to this problem, focusing on the resilience of communication networks to large-scale disasters, such as earthquakes and cyber-attacks.

In this paper, we have presented a number of principles of resilience, which can inform the ongoing research in the wider community and, more specifically, the activities of the Action. In the first instance, it is necessary to come to a common understanding about what we, as a community, mean when we talk about resilience for communication networks. After all, this is our goal, which should be clearly defined. The subject of resilience has been studied in a number of disciplines, including biology and economics. We draw inspiration from these works to define, in Sec. II, a longer and shorter-form definition of resilience. There are two important aspects of these definitions: (i) an explicit mention is made of cost-effectiveness; and (ii) we highlight the need for resilience to be quantifiable. Throughout the course of the Action, we will return to this definition, for example, to examine its relationship to related concepts such as sustainability.

An important aspect of resilience is preparedness – this is embodied in the “Defend” stage of the $D^2R^2 + DR$ resilience strategy. To be prepared, one has to have an understanding of the nature of the challenges that need to be addressed. To this end, we present an overview of related work on taxonomies of challenges to communication networks. It can be seen from this presentation that challenges to communication networks are wide-ranging, e.g., from terrorist attacks through to tornadoes. Therefore, it is important to analyse the risks associated with challenges, in a given context; a topic that we do not cover in this paper. With an understanding of risk, it is then possible for an operator to prioritise the implementation of resilience measures. Arguably, the presented taxonomies do not pay sufficient attention to the “internal changes” – a form of challenge – that are mentioned in our resilience definition. In particular, organisational changes, which have significant impact on resilience, are not well-explored. This could be an area for further consideration.

As mentioned earlier, resilience should be measurable – without the capacity to measure how resilient a communication infrastructure is, it is not possible to determine the effectiveness of a given set of resilience solutions. Here, we present an overview of related work that has proposed approaches to measuring the resilience of communication networks. An interesting aspect of these approaches is how they consider (model) the interdependency of resilience measures at different layers (e.g., physical and network layers). As communication networks increasingly support future cyber-physical systems, e.g., the smart grid, it could be of value to develop frameworks for identifying and measuring similar dependencies between properties in the cyber (network) and physical domains.

To systematically engineer resilient communication networks – and go beyond the development of point solutions for specific challenges – it is important to make use of architectural principles. Here, we have presented related work on two main forms of architectural resilience that should be considered: structural and operational. Structural resilience is concerned with techniques such as redundancy and diversity of components, whereas operational aspects consider resilience management functions that coordinate end-to-end and multi-level resilience mechanisms. A combination of both is required – a key research question relates to determining when should one approach be taken over another, and how do operational resilience approaches interface with the organisations and individuals that operate them (i.e., questioning the role of the human on the loop).

Finally, we have presented best practices for network resilience that are being employed by a major telecoms operator, and the applied research direction they are taking with respect to how resilience can be improved with the use of software-defined networks. A goal of the EU COST Action RECODIS is to examine approaches such as these, considering their suitability to mitigate large-scale disasters, and propose solutions that can improve the overall resilience provided by communication networks to end-users.

ACKNOWLEDGMENTS

This article is based upon work from COST Action CA15127 (“Resilient communication services protecting end-user applications from disaster-based failures – RECODIS”), supported by COST (European Cooperation in Science and Technology), and by the NSF grant CNS-1219028 (Resilient Network Design for Massive Failures and Attacks). The work has also been supported, in part, by the Fundação para a Ciência e a Tecnologia (FCT) under project grant UID/MULTI/00308/2013.



REFERENCES

[1] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, “Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines,”

Computer Networks: Special Issue on Resilient and Survivable Networks (COMNET), vol. 54, no. 8, pp. 1245–1265, June 2010.

[2] (2009, December) ResumeNet <http://www.resume.net.eu/project/index> wiki.

[3] J. P. G. Sterbenz and D. Hutchison. (2006, April) Resilinet: Multilevel resilient and survivable networking initiative <http://wiki.ittc.ku.edu/resilinet> wiki.

[4] European Network and Information Security Agency (ENISA), “Measurement Frameworks and Metrics for Resilient Networks and Services: Technical report,” Tech. Rep., February 2011.

[5] A Community and Regional Resilience Institute (CARRI) Report, “Definitions of community resilience: An analysis,” <http://www.resilientus.org/wp-content/uploads/2013/08/definitions-of-community-resilience.pdf>, 2013.

[6] F. S. Brand and K. Jax, “Focusing the meaning(s) of resilience: resilience as a descriptive concept and a boundary object,” *Ecology and Society*, vol. 12, no. 1, p. 23, 2007.

[7] E. K. Çetinkaya and J. P. G. Sterbenz, “A Taxonomy of Network Challenges,” in *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, Budapest, March 2013, pp. 322–330.

[8] A. Avižienis, J.-C. Laprie, and B. Randell, “Dependability and Its Threats: A Taxonomy,” in *Building the Information Society*, ser. IFIP International Federation for Information Processing, R. Jacquart, Ed. Springer Boston, 2004, vol. 156, pp. 91–120.

[9] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. G. Sterbenz, “Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach,” *Telecommunication Systems*, vol. 52, no. 2, pp. 751–766, 2013.

[10] Y. Kitamura, Y. Lee, R. Sakiyama, and K. Okamura, “Experience with restoration of asia pacific network failures from taiwan earthquake,” *IEICE Transactions on Communications*, vol. E90-B, no. 11, pp. 3095–3103, November 2007.

[11] Y. Ran, “Considerations and Suggestions on Improvement of Communication Network Disaster Countermeasures after the Wenchuan Earthquake,” *IEEE Communications Magazine*, vol. 49, no. 1, pp. 44–47, 2011.

[12] S. Urushidani, M. Aoki, K. Fukuda, S. Abe, M. Nakamura, M. Koibuchi, Y. Ji, and S. Yamada, “Highly available network design and resource management of SINET4,” *Telecommunication Systems*, vol. 56, no. 1, pp. 33–47, 2014. [Online]. Available: <http://dx.doi.org/10.1007/s11235-013-9817-8>

[13] J. Cowie, A. Popescu, and T. Underwood, “Impact of hurricane Katrina on Internet infrastructure,” Renesys, report, September 2005.

[14] J. Kappenman, “A Perfect Storm of Planetary Proportions,” *IEEE Spectrum Magazine*, vol. 49, no. 2, pp. 26–31, 2012.

[15] C. Partridge, P. Barford, D. D. Clark, S. Donelan, V. Paxson, J. Rexford, and M. K. Vernon, “The internet under crisis conditions: Learning from september 11,” National Research Council, report 10659, 2003.

[16] S. Bafna, A. Pandey, and K. Verma, “Anatomy of the Internet Peering Disputes,” *CoRR*, vol. abs/1409.6526, 2014. [Online]. Available: <http://arxiv.org/abs/1409.6526>

[17] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé, “Analysis of Country-wide Internet Outages Caused by Censorship,” *IEEE/ACM Transactions on Networking*, vol. 22, no. 6, pp. 1964–1977, December 2014. [Online]. Available: <http://dx.doi.org.libproxy.mst.edu/10.1109/TNET.2013.2291244>

[18] J. H. Cowie, A. T. Ogielski, B. Premore, E. A. Smith, and T. Underwood, “Impact of the 2003 Blackouts on Internet Communications,” Renesys Corporation, Preliminary Report, November 2003, (updated March 1, 2004).

[19] R. Mahajan, D. Wetherall, and T. Anderson, “Understanding BGP Misconfiguration,” in *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, Pittsburgh, PA, August 2002, pp. 3–16.

[20] J. Mirkovic and P. Reiher, “A taxonomy of DDoS attack and DDoS defense mechanisms,” *SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 39–53, 2004.

[21] J. Jung, B. Krishnamurthy, and M. Rabinovich, “Flash crowds and denial of service attacks: Characterization and implications for CDNs and web sites,” in *Proceedings of the 11th ACM International Conference on World Wide Web (WWW)*, Honolulu, HI, 2002, pp. 293–304.

- [22] W. LeFebvre, "CNN.com: Facing a World Crisis," in *Proceedings of the 15th USENIX Conference on Systems Administration (LISA)*, San Diego, CA, December 2001, invited Talk.
- [23] J. Rak, *Resilient Routing in Communication Networks*. Switzerland: Springer International Publishing, 2015.
- [24] J. Moura and D. Hutchison, "Review and analysis of networking challenges in cloud computing," *Journal of Network and Computer Applications*, vol. 60, pp. 113–129, January 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S108480451500288X>
- [25] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.
- [26] J. F. Meyer, "On Evaluating the Performability of Degradable Computing Systems," *IEEE Transactions on Computers*, vol. 100, no. 29, pp. 720–731, 1980.
- [27] N. Edwards, "Building dependable distributed systems," ANSA, Technical report APM.1144.00.02, February 1994.
- [28] J. C. Knight, E. A. Strunk, and K. J. Sullivan, "Towards a rigorous definition of information system survivability," in *Proceedings of the DARPA Information Survivability Conference and Exposition DISCEX III*, Washington DC, April 2003, pp. 78–89.
- [29] P. E. Heegaard and K. S. Trivedi, "Network survivability modeling," *Computer Networks*, vol. 53, no. 8, pp. 1215–1234, 2009, performance Modeling of Computer Networks: Special Issue in Memory of Dr. Gunter Bolch.
- [30] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Redundancy, diversity, and connectivity to achieve multilevel network resilience, survivability, and disruption tolerance (invited paper)," *Telecommunication Systems*, vol. 56, no. 1, pp. 17–31, 2014. [Online]. Available: <http://dx.doi.org/10.1007/s11235-013-9816-9>
- [31] J. P. Sterbenz, E. K. Çetinkaya, M. A. Hameed, A. Jabbar, Q. Shi, and J. P. Rohrer, "Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation (invited paper)," *Telecommunication Systems*, vol. 52, no. 2, pp. 705–736, 2013.
- [32] T1A1.2 Working Group, "Network survivability performance," Alliance for Telecommunications Industry Solutions (ATIS), Technical Report T1A1.2/93-001R3, November 1993.
- [33] Y. Cheng and J. P. Sterbenz, "Critical region identification and geodiverse routing protocol under massive challenges," in *Reliable Networks Design and Modeling (RNDM), 2015 7th International Workshop on*, 2015.
- [34] E. K. Çetinkaya, M. J. F. Alenazi, A. M. Peck, J. P. Rohrer, and J. P. G. Sterbenz, "Multilevel resilience analysis of transportation and communication networks," *Telecommunication Systems*, vol. 60, no. 4, pp. 515–537, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s11235-015-9991-y>
- [35] J. P. Rohrer, A. Jabbar, and J. P. Sterbenz, "Path diversification for future internet end-to-end resilience and survivability," *Telecommunication Systems*, vol. 56, no. 1, pp. 49–67, 2014.
- [36] C. Doerr and J. M. Hernandez, "A computational approach to multi-level analysis of network resilience," in *Dependability (DEPEND), 2010 Third International Conference on*, July 2010, pp. 125–132.
- [37] D. Tipper, "Resilient network design: challenges and future directions," *Telecommunication Systems*, vol. 56, no. 1, pp. 5–16, 2014.
- [38] A. Alashaikh, T. Gomes, and D. Tipper, "The spine concept for improving network availability," *Computer Networks*, vol. 82, no. 0, pp. 4–19, 2015, robust and Fault-Tolerant Communication Networks.
- [39] A. Alashaikh, D. Tipper, and T. Gomes, "Supporting differentiated resilience classes in multilayer networks," in *12th International Conference on Design of Reliable Communication Networks - DRCN 2016*, Paris, France, 15-17 March 2016, pp. 31–38.
- [40] Y. Yu, M. Fry, A. Schaeffer-Filho, P. Smith, and D. Hutchison, "An adaptive approach to network resilience: Evolving challenge detection and mitigation," in *Design of Reliable Communication Networks (DRCN), 2011 8th International Workshop on the*, Oct 2011, pp. 172–179.
- [41] A. Schaeffer-Filho, P. Smith, and A. Mauthe, "Policy-driven network simulation: A resilience case study," in *Proceedings of the 2011 ACM Symposium on Applied Computing*, ser. SAC '11. New York, NY, USA: ACM, 2011, pp. 492–497. [Online]. Available: <http://doi.acm.org/10.1145/1982185.1982293>
- [42] A. Schaeffer-Filho, P. Smith, A. Mauthe, D. Hutchison, Y. Yu, and M. Fry, "A framework for the design and evaluation of network resilience management," in *2012 IEEE Network Operations and Management Symposium*, April 2012, pp. 401–408.
- [43] P. Smith, A. Schaeffer-Filho, D. Hutchison, and A. Mauthe, "Management patterns: SDN-enabled network resilience management," in *2014 IEEE Network Operations and Management Symposium (NOMS)*, May 2014, pp. 1–9.
- [44] M. Gunkel and M. Horneffer, "Multi-layer resilience in deutsche telekom's ip core and aggregation network," in *Photonic Networks; 15. ITG Symposium*, Leipzig, Germany, 2014, pp. 1–6.
- [45] M. Gunkel, F. Wissel, J. Blendin, D. Herrmann, M. Wichtlhuber, and D. Hausheer, "Efficient partial recovery of flexible-rate transceivers with sdn-based asymmetric multipath routing of ip traffic," in *Photonic Networks; 17. ITG-Symposium*, Leipzig, Germany, 2016, pp. 1–6.
- [46] D. Rafique *et al.*, "Multi-flex field trial over 762km of G.652 SSMF using programmable modulation formats up to 64QAM," in *Optical Fiber Communication Conference (OFC)*, March 2016.