

Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability, and Disruption Tolerance

invited paper

James P.G. Sterbenz · David Hutchison ·
Egemen K. Çetinkaya · Abdul Jabbar · Justin P. Rohrer ·
Marcus Schöller · Paul Smith

This research was supported in part by the National Science Foundation FIND (Future Internet Design) Program under grant CNS-0626918 (Postmodern Internet Architecture), by NSF grant CNS-1050226 (Multilayer Network Resilience Analysis and Experimentation on GENI), and the European Commission FIRE (Future Internet Research and Experimentation Programme) under grant FP7-224619 (ResumeNet).

James P.G. Sterbenz
Electrical Engineering and Computer Science,
Information and Telecommunication Technology Center,
The University of Kansas, Lawrence, Kansas, USA
<http://www.ittc.ku.edu/resilinet>
jpgs@ittc.ku.edu, +1 508 944 3067
School of Computing and Communications, InfoLab21,
Lancaster University, Lancaster, UK
jpgs@comp.lancs.ac.uk

David Hutchison
School of Computing and Communications, InfoLab21,
Lancaster University, Lancaster, UK
dh@comp.lancs.ac.uk

Egemen K. Çetinkaya
Electrical Engineering and Computer Science,
Information and Telecommunication Technology Center,
The University of Kansas, Lawrence, Kansas, USA
ekc@ittc.ku.edu

Abdul Jabbar*
General Electric Global Research,
Niskayuna, NY, USA
jabbar@ge.com

Justin P. Rohrer*
Computer Science Department,
Naval Postgraduate School Monterey, CA, USA
jprohrer@nps.edu

Marcus Schöller*
NEC Laboratories Europe,
Heidelberg, Germany
marcus.schoeller@neclab.eu

Paul Smith*
Safety and Security Department,
AIT Austrian Institute of Technology,
Seibersdorf, Austria

Abstract Communication networks are constructed as a multilevel stack of infrastructure, protocols, and mechanisms: links and nodes, topology, routing paths, interconnected realms (ASs), end-to-end transport, and application interaction. The resilience of each one of these levels provides a foundation for the next level to achieve an overall goal of a resilient, survivable, disruption-tolerant, and dependable Future Internet. This paper concentrates on three critical resilience disciplines and the corresponding mechanisms to achieve multilevel resilience: redundancy for fault tolerance, diversity for survivability, and connectivity for disruption tolerance. Cross-layering and the mechanisms at each level are described, including richly connected topologies, multipath diverse routing, and disruption-tolerant end-to-end transport.

Keywords resilient, survivable, disruption-tolerant Future Internet · dependability, reliability, availability, performability · redundancy, diversity, eventual connectivity · cross-layer optimisation · multilevel network analysis

1 Introduction and Motivation

The increasing importance of the Global Internet has led to it becoming one of the critical infrastructures [2] on which almost every aspect of our lives depend. Thus it is essential that the Internet be *resilient*, which we define as 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 [62,

paul.smith@ait.ac.at

*Research performed while authors were at The University of Kansas and Lancaster University

63]. It is generally recognised that the current Internet is not as resilient, survivable, dependable, and secure as needed given its increasingly central role in society [1,3–5,21,58]. Thus, we need to ensure that *resilience is a fundamental design property of the Future Internet*, and seek ways to increase the resilience of the current and future Internet. This requires an understanding of vulnerabilities of the current Internet, as well as a methodology to test alternative proposals to increase resilience. In particular, we are interested in understanding, modelling, and analysing the properties of *dependability* that quantifies the reliance that can be placed on the service delivered including reliability and availability [7,31] and *performability* that quantifies the level of performance [36–39] when the network is challenged. This notion of resilience subsumes *survivability* that is the ability to tolerate the correlated failures that result from attacks and large-scale disasters [15,16,20,34,40,47,64] and *disruption-tolerance* that is the ability to communicate even when stable end-to-end paths may not exist due to weak channel connectivity, mobility, unpredictable delay, and energy constraints [17,29,64].

This paper presents a brief survey of techniques to achieve resilience in terms of fault-tolerance, survivability, and performability, at every level of the network and is organised as follows: Section 2 briefly reviews the ResiliNets strategy and design principles (with emphasis on redundancy, diversity, and connectivity) and discusses the importance of cross-layering. Section 3 discusses resilience at each level: physical infrastructure, network topology, path routing, inter-realm, end-to-end transport, and applications. Section 4 describes the multilevel state-space resilience metric. Section 5 concludes with a summary. This invited paper is based, in part, on a tutorial given at RNDM 2011 (IFIP/IEEE *Reliability Networks Design and Modeling*) in Budapest.

2 Strategy, Principles, Multilevel Architecture

This section presents a brief review of the ResiliNets strategy and principles and that form the foundation for resilience techniques at multiple levels of the network to provide overall fault tolerance, survivability, and disruption tolerance. The importance of interlevel translucency and a model for cross-layering are also presented.

2.1 ResiliNets Strategy

The ResiliNets D^2R^2+DR strategy (described in detail [63]) has been developed for the architecture and

design of resilient systems, consisting of nested control loops, as shown in Figure 1.

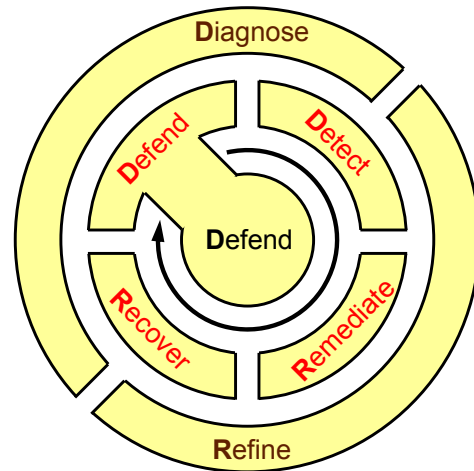


Fig. 1 ResiliNets strategy

At the center is structural **passive defence** that resists challenges to the network, such as the redundant, diverse topologies described in Section 3.2. This is surrounded by a pair of control loops.

The inner D^2R^2 real-time loop consists of four phases of operation in every network subsystem and protocol: **Active defence** resists attacks and challenges on the network using mechanisms such as filtering on known signatures. When challenges do manage to penetrate the network, context-aware **detection mechanisms** trigger adaptive **remediation** mechanisms such as dynamic rerouting and walling off compromised systems that aim to deliver the best service possible after an adverse event and during an ongoing challenge or attack. Finally, **recovery** mechanisms and infrastructure redeployment are used to bring the network back to normal operations and acceptable service for all users.

The outer **DR** background loop is used to **diagnose** the root cause of why a challenge was able to penetrate the network, and to perform an analysis of the entire inner loop operation to **refine** future behaviour for improved D^2R^2 inner-loop operation.

2.2 Resilience Principles

The ResiliNets strategy motivates a set of design principles for resilient systems [62,63], which include prerequisites (service requirements, normal behaviour understanding, threat and challenge models, metrics, heterogeneity); tradeoffs (resource tradeoffs, complexity, state management); enablers (self-protection, connec-




Fig. 2 Resilience principles

tivity, redundancy, diversity, multilevel, context awareness, translucency); and behaviour (self-organising and autonomic, adaptable, evolvable), as shown in Figure 2.

2.2.1 Redundancy, Diversity, and Connectivity

In this paper we concentrate on seven of the resilience-enabling principles (quoted in part from [63] that describes all of the principles in greater detail); beginning with the enablers to fault-tolerance, survivability, and disruption tolerance: **redundancy**, **diversity**, and **connectivity and association**.

Redundancy in space, time, and information increases resilience against faults and some challenges if defences are penetrated. Redundancy refers to the replication of entities in the network, generally to provide fault-tolerance. In the case that a fault is activated and results in an error, redundant components are able to operate and prevent a service failure. It is important to note that redundancy does not inherently prevent the redundant components from sharing the same fate, motivating the need for diversity.

Diversity is closely related to redundancy, but has the key goal to avoid *fate sharing*. Diversity in space, time, medium, and mechanism increases resilience against challenges to particular choices, and consists of providing alternatives so that even when challenges impact particular alternatives, other alternatives prevent degradation from normal operations. Diverse alternatives can either be simultaneously operational, in which case they defend against challenges [53, 55], or they may be available for use as needed to remediate. Diversity is an essential technique to provide survivability.

Connectivity and association among communicating entities should be maintained when possible based on eventual stability, but information flow should still take place even when a stable end-to-end path does not exist based on the eventual connectivity model [64], using DTN (disruption-tolerant networking) techniques such as partial paths, store-and-forward with custody transfer, and store-and-haul (store-carry-forward).

Thus, each of these principles has a direct role in a particular aspect of resilience:

- redundancy for fault tolerance
- diversity for survivability
- connectivity for disruption tolerance

2.2.2 Multilevel Principles

Four additional principles capture the aspects of resilience mechanisms operating at multiple levels in the network: **multilevel resilience**, **context awareness**, **translucency**, and **resource tradeoffs**.

Multilevel resilience is needed in three orthogonal dimensions: *protocol layers* in which resilience at each layer provides a foundation for the next layer above; *planes*: data, control, and management; and *network architecture* inside-out from fault-tolerant components, through survivable subnetwork and network topologies, to the Global Internet including attached end systems and applications.

Context awareness is highly related to multilevel resilience, which allows the networked systems to sense the communication environment and **detect** when a challenge has penetrated the network, and **adaptable** and **autonomic** behaviour that permits the system to react and optimise as part of the **remediation** process of the D^2R^2+DR strategy.




Fig. 3 Cross-layer knobs and dials

Translucency permits cross-layer and plane interactions, and is critical to optimising resilience across levels. Figure 3 depicts simple cross-layering, in which the application sets *knobs* based on its service requirements to be conveyed to the network below. The network uses these to optimise its behaviour, and conveys its state up to the application via *dials*, which allow the application to optimise and further adjust its knobs [13, 65].

The multilevel aspect of resilience analysis is discussed in Section 2.3. Figure 4 depicts a multilevel cross-layering model, for simplicity showing only two levels: E2E (end-to-end) and HBH (hop-by-hop). As described

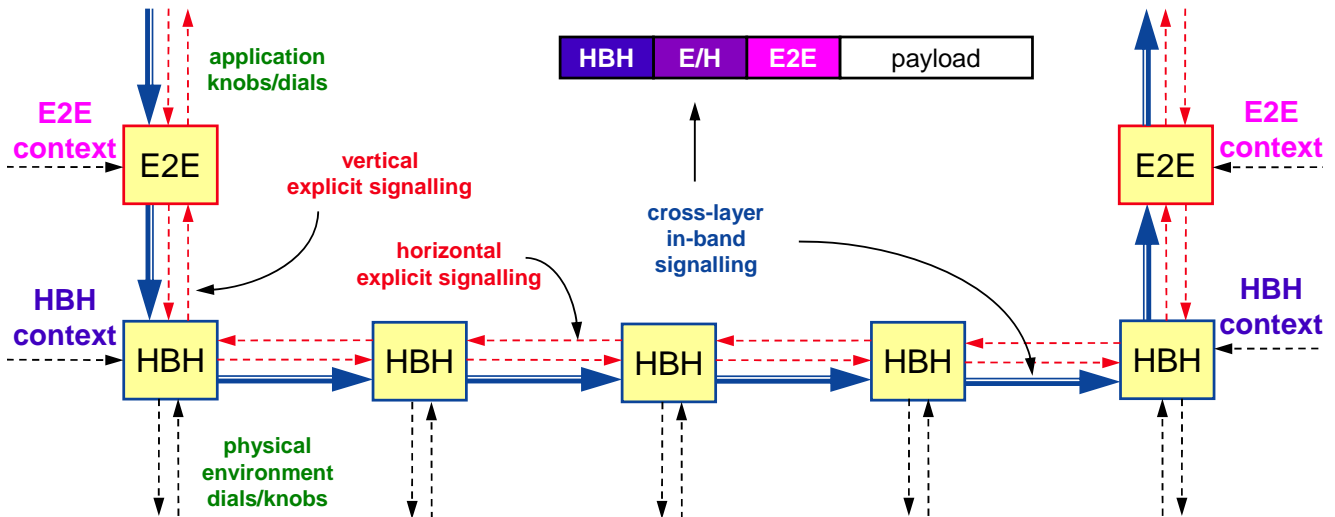


Fig. 4 E2E and HBH cross-layering

in Section 2.3, there are additional levels in the architecture, and in particular the third data-plane *inter-realm* level between HBH and E2E that interconnects heterogeneous (sub-)network realms. This model builds on the taxonomy originally introduced for ETEN (explicit transport error notification) [30].

Resource tradeoffs determine the deployment of resilience mechanisms. The relative composition and placement of these resources must be balanced to optimise resilience and cost. The maximum availability of a particular resource serves as a constraint in these optimisations. Resources to be traded against one-another include bandwidth, memory [43], processing, latency [65], energy, and monetary cost. Of particular note is that maximum resilience can be obtained with unlimited cost, but there are cost constraints that limit the use of enablers such as redundancy and diversity. In the context of multilevel resilience, resource optimisation is done across levels. While each level forms a resilient foundation on which the next level operates, the resilience of the upper level can compensate for limited resilience of the level below, helping limit the cost of both levels.

2.2.3 Cross-Layer Model

Cross-layer signalling is implemented as downward *knobs* \mathbb{K} that influence the behaviour of the level below, and as upward *dials* \mathbb{D} that instrument the characteristics of a layer. Cross-layer signals (\mathbb{K}, \mathbb{D}) are a combination of in-band controls embedded in PDU (protocol data unit) headers (\mathbf{k}, \mathbf{d}) and out-of-band signalling messages (\mathbf{K}, \mathbf{D}) , thus $(\mathbb{K}, \mathbb{D}) = (\mathbf{K} \cup \mathbf{k}, \mathbf{D} \cup \mathbf{d})$. Data flows vertically between layers and horizontally within lay-

ers as shown by the thick arrows (for clarity only unidirectionally) in Figure 4. In-band signalling (\mathbf{k}, \mathbf{d}) is shown by the narrow line associated with the data arrows; explicit out-of-band signals (\mathbf{K}, \mathbf{D}) are shown by dashed arrows in both the forward and reverse direction of data flow. Every node collects context information c_n about its environment at layer L_n (for simplicity in the figure not all HBH context gathering is shown). At a given layer, context and cross-layer fields are computed as they flow through the nodes. For example, if the HBH dial to E2E transport is BER error rate, then as a packet flows through the network the BER is recomputed to account for each hop [30].

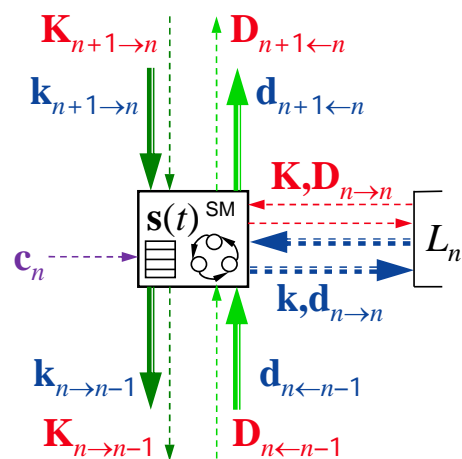


Fig. 5 Cross-layer model

Figure 5 shows a single node containing a protocol state machine and its state at level n . Vertical controls

represented as $(\mathbb{K}_{n+1 \rightarrow n}, \mathbb{D}_{n \leftarrow n-1})$ indicate knobs down from layer $n+1$ and dials up from layer $n-1$. Horizontal signals are represented as $(\mathbb{K}_{n \rightarrow n}, \mathbb{D}_{n \leftarrow n})$.

2.2.4 Cross-Layer Composition

Multiple layers can be composed either by concatenation, or by opaquely bypassing layers, as shown in Figure 6.

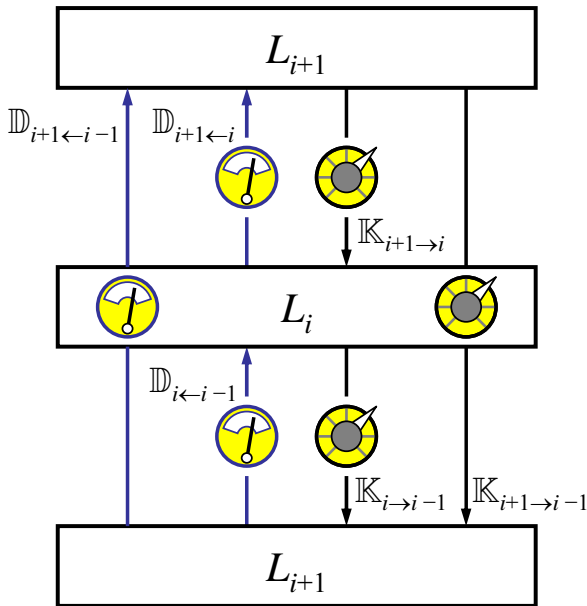


Fig. 6 Cross-layer composition

In the case of concatenation, a multilayer knob, dial pair $(\mathbb{K}_{n+1 \rightarrow n-1}, \mathbb{D}_{n+1 \leftarrow n-1})$ is formed by the concatenation of individual single-layer knobs and dials:

$$\mathbb{K}_{n+1 \rightarrow n-1} = \mathbb{K}_{n+1 \rightarrow n} \parallel \mathbb{K}_{n \rightarrow n-1} \text{ and}$$

$\mathbb{D}_{n+1 \leftarrow n-1} = \mathbb{D}_{n+1 \leftarrow n} \parallel \mathbb{D}_{n \leftarrow n-1}$. Bypassing layers runs the risk that the bypassed layer is not able to alter its behaviour based on the knobs and dials opaque to it. If multiple control loops bypass layers and overlap one another (are not strictly nested) the risk of unpredictable feature interaction is increases.

2.3 Multilevel Network Architecture

We now describe the organisation of the network architecture into multiple levels, for the purpose of understanding, analysing, and reasoning about network resilience.

The three dimensions of multilevel resilience and relationship to the ResiliNets strategy dimension is shown

in Figure 7 (not corresponding directly to the three dimensions of the cube). The three multilevel resilience dimensions are protocol layer, protocol plane, and network engineering. In each case, we can view the resilience of a level providing a foundation for the level above. While this bottom-up view is useful, it is important to understand that this is constrained by **resource tradeoffs**, that is the lack of resilience of a given level due to cost constraints can be compensated by the level above. For example, a fully resilient network layer would consist of a full mesh of strongly connected links, but this is infeasible for almost all networks due to the cost of n^2 interconnections. Thus, a resilient transport layer is designed to deal with an imperfect network layer.

The $\mathbf{D}^2\mathbf{R}^2 + \mathbf{DR}$ strategy dimension is also projected onto this cube. The inner $\mathbf{D}^2\mathbf{R}^2$ real-time loop exists primarily in the control plane in the mechanisms that implement the inner loop, but also in the data plane to the degree that these mechanisms are part of data transfer (for example embedding of FEC). The \mathbf{DR} outer loop is related to the management plane in the long-term analysis and evolution of network resilience architecture and mechanisms.

2.3.1 Protocol Layer Dimension

The protocol layer dimension is the conventional layered network abstraction model, showing some of the alternatives that might be chosen at each level; in each case these alternatives are dynamically **adaptive** based on **context awareness**, particularly for **remediation**:

- physical layer using robust coding techniques
- medium access control layer alternatives (e.g. CSMA vs. TDMA vs. CDMA)
- hop-by-hop link layer with alternative error control (FEC – forward error correction, ARQ – automatic repeat request, and hybrid)
- network layer with multipath diverse multipath routing and eventual connectivity
- internetwork layer between realms of diverse network technology, policy, and trust [8]
- end-to-end disruption tolerant transport layer with adaptive error control (FEC – forward error correction, ARQ – automatic repeat request, and hybrid) traded against HBH error control
- application layer that is adaptive between service requirements and E2E transport service (e.g. adapting frame rate and resolution)

Cross-layer optimisations are shown by the cross-layer loops in Figure 7. More detail on the techniques at each layer are provided in Section 3.

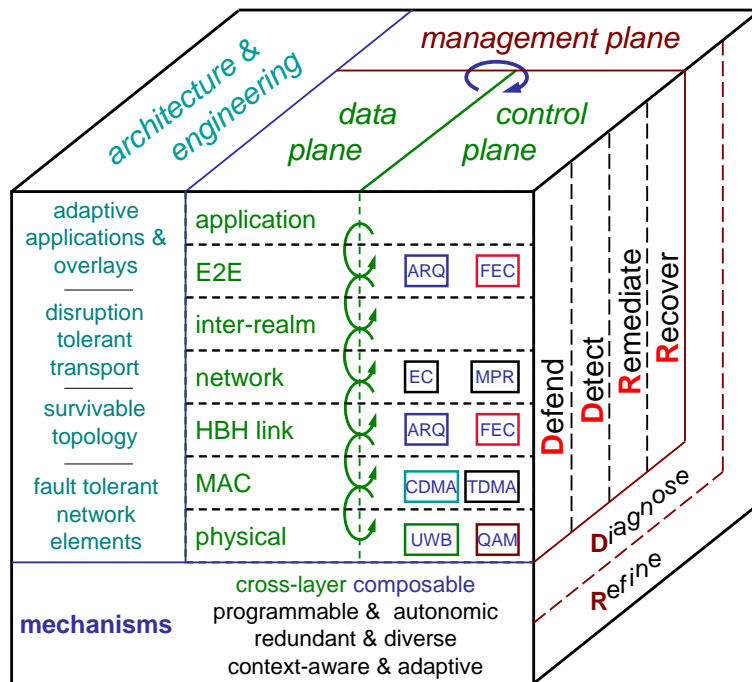


Fig. 7 Multilevel mechanisms and ResiliNets strategy

2.3.2 Protocol Plane Dimension

The protocol plane dimension is the division into data, control, and management planes. It is critical that each of these is resilient, and that the resilience mechanisms across planes are **translucent** to interact with one-another, as shown by the cross-plane loops in Figure 7.

2.3.3 Network Engineering Dimension

The network engineering dimension is related to the layered view, but reflects the structure of the network from an architecture, design, and engineering perspective. At the bottom, individual components such as routers should be fault-tolerant for dependable operation. The network topology based on these components should be survivable to the failure of multiple components. End-to-end transport should be disruption tolerant to challenged topologies. Finally, the Global Internet including the distributed applications should provide a resilient service to users.

3 Multilevel Resilience

For purposes of multilevel resilient network architecture, we concentrate on a set of six levels shown in Table 1. These levels capture aspects of the protocol layer, plane, and network engineering levels described above. Note, in particular, that the conventional layer 3

is divided into two sub-layers: topology (3t) and path routing (3r).

These levels are the basis for the description of multilevel resilience in the rest of this paper, in each case considering redundancy for fault tolerance, diversity for survivability, and connectivity for disruption tolerance.

3.1 Physical Infrastructure Level

The physical infrastructure level consists of the components that constitute the network: links (wired or wireless associations) and nodes (switches, routers, servers, end systems). The network infrastructure is also interdependent with the power grid: when the grid fails network components fail (and vice versa) [14,32]. This relationship will be used as an example in Section 3.2.2.

3.1.1 Redundancy

Redundancy in the design of network components is a well established technique for fault tolerance. Fault-tolerant components contain redundant sub-components such that random failures do not prevent the component from continuing operation. The canonical example of this is using triple-modular redundancy [33] as a way to significantly increase system reliability.

Table 1 Multilevel resilience metrics and mechanisms

#	Level	Metrics	Mechanisms
7	application	latency, throughput, QoE	proxy servers, caching, content replication, adaptivity
4	E2E transport	path latency, goodput, PDR (packet delivery ratio)	erasure coding, adaptive error control, disruption tolerance
3.5	inter-realm	realm connectivity and transit	cross-layering, heterogeneity
3r	path routing	routing delay, overhead, path dependability	store-and-haul, multipath, eventual connectivity
3t	topology	graph metrics (clustering coefficient, largest component size, betweenness, etc)	k -connectivity, p -cycles, diversity paths
2	physical	link quality (FER, BER), availability, reliability, Pr[component failure]	triple-modular redundancy, redundant links, robust coding, HARQ

3.1.2 Diversity

Diversity in network components provides alternatives, such that if a particular type of component is attacked not all of them will fail. With respect to hardware vendor and software system, this means avoiding monocultures. For example, a large scale coordinated attack against either Microsoft Windows or Cisco IOS would have disastrous consequences on the Global Internet, given the dominant market share of these end-system and router platforms.

With respect to mechanism, diversity means a mix of wired and wireless links such that if a wired link is cut, the wireless can be used, and if a wireless link is attenuated or jammed the wired can be used. This diversity significantly raises the difficulty of effective attacks against the network infrastructure.

3.1.3 Connectivity

Connectivity in the context of infrastructure refers to mechanisms that permit information transfer across a link even in challenged environments. This is typically due to wireless channels with weak, intermittent, or asymmetric connectivity. Techniques at the physical layer include robust coding, and at the link layer adaptive hybrid error control.

3.2 Network Topology Level

The network topology level uses the individual nodes and links in the physical infrastructure level to construct a network graph to connect the end systems and servers with intermediate systems (switches and routers). The goal of this level is to use the foundation of resilient components to create a resilient network topology.

3.2.1 Topology Redundancy

Redundancy in the topology means that the graph is rich enough to be fault tolerant. At the very least a bi-connected topology is required such that all node pairs remain connected when any single link fails, and all remaining nodes can communicate when a single node fails. Note that dual ring networks such as SDH/SONET provide bi-connectivity. Additional redundancy to multiple failures can be provided with k -connectivity [45]. Techniques such as p -cycles [22] can provide this redundancy in mesh networks.

3.2.2 Topology Diversity

Diversity in topology is needed to provide survivability against correlated failures, including attacks against the infrastructure by intelligent adversaries with knowledge of the network structure and its vulnerabilities, as well as correlated area-based failures from natural disasters such as hurricanes and coronal mass ejections. Furthermore, topological diversity becomes an essential mechanism with the increasing interdependencies between critical infrastructures. A sample scenario representing an increasing power grid failure area in the central region of the US is shown in Figure 8. In this case, geographically diverse links in the ISP topology can alleviate the impact of link failures as long as alternative link capacities handle the offered load [11,12].

Diverse topology design is an essential mechanism that should be considered to build resilient networks. In the case of the Baltimore tunnel fire [10,67], the redundancy of having different service providers was useless since different service providers lay their fibres through the same geographic location. Therefore, not only logical topology, but also underlying physical topologies should be considered carefully when designing networks. However, increased geographical diversity increases the build-up and operational costs of networks. Therefore, optimising the network resilience and cost is non-trivial.




Fig. 8 Topology diversity example

3.2.3 Topology Connectivity

Connectivity in the context of the topology level consists of creating topologies that are richly interconnected enough and are aware of weakly-connected components, such that end-to-end stable paths are available whenever possible. The goal is first to maintain a connected topology whenever possible, or practical given cost constraints, such that routing can converge; this is eventual stability [64]. Only when stable end-to-end paths are not available does the path routing level need deal with eventual connectivity, described in Section 3.3.

3.3 Path Routing Level

The path routing level uses the topology level over which to route end-to-end paths with desired resilience properties. Resilient routing is able to route paths even when the underlying topology is not very-well or stably connected. In the case of dynamic networks such as MANETs (mobile ad hoc networks), resilient routing is able to discover and dynamically reroute paths as the network topology changes.

3.3.1 Path Routing Redundancy

Redundancy in path routing means that multiple paths are available between a pair of endpoints. This can be in the form of alternate paths that can be quickly switched (e.g. fast IP reroute [35] and SPGC [46]), or multipath routing available to the transport layer so that a particular flow can be spread over multiple paths such that the disruption of any single path does not disrupt the end-to-end flow.

3.3.2 Path Routing Diversity

Diversity in path routing exploits diversity in the topology level to create multiple diverse paths that can be

used by a transport flow. Several measures of diversity quantify the degree to which alternate paths share the same nodes and links: EPD (effective path diversity), TGD (total graph diversity), and cTGD (compensated graph diversity) [54, 53, 55, 56]. The path diversity measures provide a single value that can evaluate the topology and utility of added path diversity. Furthermore, it is important to measure diversity in terms of *physical distances*, not only node and link disjointness. The previous path diversity measures consider the sharing of components, but do not capture the geographic characteristics necessary for area-based challenges such as large-scale disasters or to prevent the geographic fate sharing of distinct links in the same conduit, as in the Baltimore tunnel fire. The diversity measures can be augmented with a minimum distance between any pair of nodes along alternate paths, and as the area inside a polygon or set of polygons, the borders of which are defined by a pair of alternate paths, as shown in Figure 9.




Fig. 9 Geographical distance and area diversity

3.3.3 Path Routing Connectivity

Connectivity in the path routing level consists of discovering paths through weakly connected topologies, and using the *eventual connectivity* model to communicate even when stable end-to-end paths do not exist [64]. Traditional routing protocols require that a complete path exists from source to destination before communication is initiated. This is true even for MANET protocols that are able to frequently reroute in mobile networks, either proactively maintaining full topology or reactively discovering routes on demand.

This *eventual stability* model of routing assumes that routing converges eventually after partitioning. While this is a traditional operating assumption, it does not hold in networks challenged by weak, intermittent, and asymmetric connectivity, or by high mobility. In these networks, routing algorithms may rarely or never converge to stable end-to-end paths. The eventual connectivity model [6] relaxes the traditional assumptions so that communication can proceed along partial segments [23] of paths between communicating nodes. Information progresses as far as possible, along whatever

paths possible, until it reaches its destination. This extends the concept of store-and-forward, and requires modifying the typical forwarding behavior of dropping packets if an outgoing link to the next node becomes temporarily unavailable.

Communication over intermittent links is depicted in Figure 10, in which all links are intermittent such that there is never a complete path, but there are times in which partial segments of a path are available. Data can be moved as shown by the solid arrow. Then, when the next two intermittent links become temporarily available, data progresses along the path of the dashed arrows (and at this point the first link may become unavailable). This model of store-and-forward communication was adopted by the DTN (delay- and disruption-tolerant networking) community [17], in which *bundles* of information are forwarded through the network with *custody-transfer* supplying nearly-reliable transfer [59].




Fig. 10 Communication over intermittent links

Furthermore, it is possible to *exploit* mobility to communicate when otherwise impossible. In the worst case, eventual connectivity routing will store data until a promising outgoing link becomes available. Proactive control can be used in two ways to expedite the transfer of data [64]. Movement control can be used to exert control on other nodes to move them into range such that a path toward the destination exists. Alternatively, mobile node can *store-and-haul* [64] packets toward their destination by physically transporting the data, as shown in Figure 11. This is now commonly called *store-carry-forward* and *message ferrying* [71].




Fig. 11 Store-and-haul data forwarding

High mobility often poses challenges to conventional MANET routing protocols especially after they reach their reactive limit. In this case it is necessary to use knowledge of the location and trajectories of nodes to predict future location without requiring rapid convergence of routing algorithms. Trajectory routing [69], as done in AeroRP [41, 44, 49, 50] uses geolocation and velocity vectors to compute destination node locations.

3.4 Inter-Realm Level

The inter-realm level is similar to the path routing level, except that it is the interconnection between *realms* in an internetwork. Realms provide the ability to internetwork disparate networking technologies (that may not be IP-compatible), and provide trust and policy boundaries. In the case of the current Internet, realms are equivalent to ASs (autonomous system domains); in the case of the Postmodern Internet [8], an inter-realm level (layer 3.5 in conventional notation) provides internetworking between realms using different addressing, forwarding, signalling, and intra-realm routing paradigms.

3.4.1 Inter-Realm Redundancy

Redundancy in the inter-realm level simply means that there are redundant realms and inter-realm links available for transit, such that the failure of a realm or its attachment point does not affect nodes and users outside the realm.




Fig. 12 Diversity example

3.4.2 Inter-Realm Diversity

Inter-realm diversity requires diversity in geography and mechanism, just as for the path routing level, as shown in Figure 12. This is naturally provided to the degree that redundant realms use different internal network-layer paradigms and mechanisms. Geographic diversity requires that realm service providers expose geolocation information about their nodes and links such that SRLGs (shared risk link groups) [66] are avoided, to

prevent the sorts of failures that occurred during the Baltimore tunnel fire [10,67], during which users that were dual-homed across different service providers still lost connectivity because their fibre links burned at the same time. Service providers are currently unable and unwilling to expose the internal network structure needed to achieve inter-realm diversity.

3.4.3 Inter-Realm Connectivity

Inter-realm connectivity refers to providing disruption-tolerant forwarding through and between realms even when the realm or inter-realm link is dynamically, intermittently, or weakly connected; this is the inter-realm equivalent of the intra-realm path routing connectivity described in Section 3.3.

3.5 End-to-End Transport Level

The end-to-end transport level uses the paths that the routing level has created to transfer data end-to-end through the network between applications and users. Resilient transport is able to adapt its mechanisms and reliability services based on the application needs (knobs from the application level above $\mathbb{K}_{7 \rightarrow 4}$) and on the end-to-end path characteristics (dials from the routing level below $\mathbb{D}_{4 \leftarrow 3t}$). The ResTP resilient transport protocol [48] (and its aeronautical subset AeroTP [50]) uses adaptive and composable mechanisms with cross-layering to achieve resilience.

Figure 13 shows the possibilities for error control at the physical infrastructure HBH (hop-by-hop) level vs. E2E (end-to-end) transport level: None, Open loop FEC, Closed-loop ARQ, and Hybrid (represented by $H = O \cup C$). The set of feasible choices is governed by application service requirements. Reliable transfer requires E2E ARQ (shown as the green or darkly-shaded oval), due to the end-to-end arguments [57]. Quasi-reliable transfer can be achieved by a variety of HBH or E2E FEC or ARQ based on optimisations of a particular scenario (shown as the yellow or lightly-shaded oval).

3.5.1 Transport Redundancy

Redundancy in end-to-end transport exploits multipath routing to increase fault-tolerance against the failure of individual paths. These paths can be available as hot-standbys, permitting remediation by rapid fail-over to the alternate paths. Alternatively, the transport flow can spread across k paths using erasure coding such that information can be recovered even if one or multiple paths fail, depending on the strength of the code;

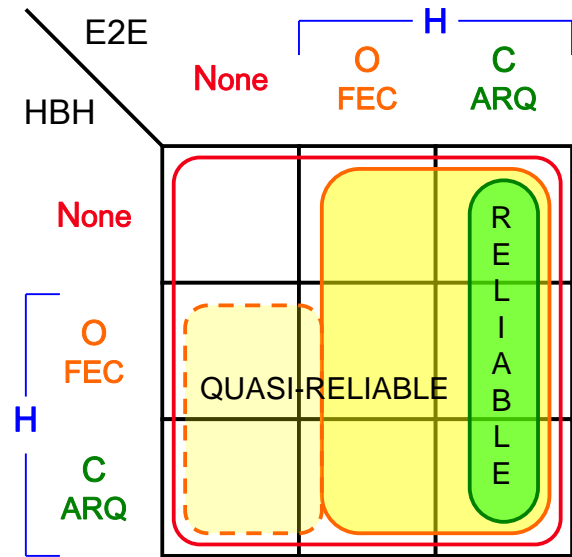


Fig. 13 Composable error control

this is desirable for real-time service at the cost of additional bandwidth. ResTP uses the desired service knobs from the application level in combination with the path characteristics dial from the path routing level to determine the degree k and whether to failover or spread across them.

3.5.2 Transport Diversity

Diversity in end-to-end transport exploits diversity available from the path routing level to increase survivability, communicated by cross-layer signals. The principal types of diversity are geographic and medium.

Geographic diversity consists of specifying not only the degree k , described in Section 3.3, but also the geographic distance d desired to meet a particular threat model and service specification. For example, a real-time service that is resilient to area based challenges of diameter d would request spreading over k paths such that no paths pass through node pairs closer than d apart.

Medium diversity consists of choosing alternatives, typically wired and wireless, such that challenges to either do not affect end-to-end communication. For example, a fiber cut is survived by wireless links; jamming or weather-based attenuation [26–28] to the wireless link is survived by the fiber link [63], as shown in Figure 12.

3.5.3 Transport Connectivity

Connectivity at the end-to-end transport level exploits the benefits of path routing and inter-realm eventual




Fig. 14 End-to-end path composition

connectivity, but must deal with the possibility of individual paths that are composed of hop-by-hop links of significantly different strength and symmetry, as shown in Figure 14; perhaps resulting in composed forward and reverse paths of significantly different characteristics.

A transport protocol such as ResTP (and its disruption-tolerant subset AeroTP [50–52]) needs to follow the general principle of DTN protocols [9, 17, 18, 29]. These include avoiding chatty round-trip signalling with opportunistic data transfer. ResTP provides explicit support for cross-layering (as described in Section 2.2.2) with composable mechanisms [19] that are dynamically adaptable to the instrumentation provided by dials from the path routing and inter-realm levels.

3.6 Application Level

The application level is the interface to the user, and relies on the end-to-end transport level to provide those associations.

3.6.1 Application Redundancy

Redundancy in applications refers to multiple instances of a particular application being available to a user, for example access to email on a variety of platforms in case one fails.

3.6.2 Application Diversity

Diversity in applications refers to providing alternatives in the applications users choose. Just as it is essential to avoid monocultures in network components (Section 3.1), they should be avoided in end system vendor, operating system, and application.

3.6.3 Application Connectivity

Connectivity in the context of applications means that they are adaptive and resilient to imperfect connectivity at the end-to-end transport level, and can optimise service performance with cross-layer optimisations and user-directed feedback. This can happen within an application, such as adapting frame rate and resolution of the available bandwidth and allowing the user to choose the tradeoff (e.g. resolution for talking heads and frame rate for action [68]). This adaptation can also occur across applications, for example degrading from video-conference to voice-conference to chat to email as available bandwidth degrades. Furthermore, applications can provide feedback to users to more intelligently direct their operation, for example in choosing links to follow based on Web browser estimates on URL (uniform resource locator) response times [61].

4 Multilevel Analysis

A resilience metric is essential for understanding the resilience of current networks, and to evaluate alternative Future Internet architecture and infrastructure, whether evolutionary or revolutionary. Ideally, this is represented as a single number \mathfrak{R} in the range of (0,1), where 0 indicates no resilience and 1 indicates infinite resilience. This is clearly a difficult problem given the complexity of network architectures and protocols at a number of levels, the variety of challenges they must tolerate, the range of application scenarios, traffic demand, and many parameters to characterise dependability [7, 31, 42, 70] and performability [36–39].




Fig. 15 Resilience across multiple levels

4.1 Resilience Evaluation \mathbb{R}

To gain tractability we formulate the resilience metric as a two dimensional state space, as shown in Figure 15 [24, 25, 60, 63]. The horizontal dimension is the operational state \mathbb{N} of the network, described as *normal operation* (for which the network is designed), through *partially degraded* to *severely degraded*. A resilient network infrastructure is one that resists degrading even when challenged. The vertical dimension is the service provided \mathbb{P} , described as *acceptable* (based on meeting a service specification), through *impaired* to *unacceptable*. A resilient service is one that resists impairment even when network operation is degraded.

Each of the axes (\mathbb{N}, \mathbb{P}) is an objective function of a set of parameters; these may be a boolean, linear, or other combination. We measure the resilience \mathbb{R} as the area under the trajectory from the initial state, generally acceptable service under normal operations, to a challenged state $S_0 \rightarrow S_c$.

The relationship of the the state-space formulation to the ResiliNets strategy described in Section 2.1 is also depicted in Figure 15. The inner D^2R^2 loop trajectory is shown. Defence prevents the system from leaving its initial state S_0 . If a challenge causes the state to change significantly, this is detected by a change in the operational or service parameters when the state goes to a challenged state S_c . Remediation improves the situation to S_r , and recovery finally returns the system to its original state S_0 .




Fig. 16 Resilience across multiple levels

4.2 Multilevel Multiscenario Resilience \mathbb{R}

In the multilevel analysis, as shown in Figure 16, the service parameters at the level boundary B_{ij} become the operation metrics at boundary $B_{i+1, j+1}$. In other words, the service provided by a given layer becomes the operational state of the layer above, which has a new set of service parameters characterizing its service to the layer above. Note that the operational and service metrics $\mathbb{N} \cup \mathbb{P}$ may directly correspond to the cross-layer knob and dial parameters $\mathbb{K} \cup \mathbb{D}$ described in Section 2.2.

By beginning at the bottom level and progressing up the service layers, an overall multilevel resilience value can be computed [24], and by composing these across all scenarios of interest for a given network architecture, it may be possible to derive a single resilience value \mathbb{R} .

5 Summary

Resilience is an essential property of the Future Internet, including fault-tolerance, survivability, and disruption tolerance. Achieving overall network resilience requires decomposing the network into levels such that the resilience of each level provides a foundation for the next.

This paper described three major aspects of resilience: redundancy for fault-tolerance, diversity for survivability, and connectivity for disruption tolerance. After describing a model for cross-layering, techniques to achieve this were discussed at each level: physical infrastructure, network topology, path routing, inter-realm, end-to-end transport, and applications. Finally, the composition of a multilevel state-space resilience metric was described.

Acknowledgements The authors would like to thank members of the ResiliNets research group at the University of Kansas and Lancaster University, as well as members of the EU ResumeNet project for discussions on, and contributions to aspects of this work. We would like to thank Jacek Rak for inviting this paper based on the RNDM 2011 tutorial given by the primary author of this paper.

References

1. Information Security: Computer Hacker Information Available on the Internet. Tech. Rep. T-AIMD-96-108, United States General Accounting Office (1996). URL <http://www.fas.org/irp/gao/aimd-96-108.htm>
2. Protecting America's infrastructures. Report, President's Commission on Critical Infrastructure Protection (1997)
3. A roadmap for cybersecurity research. Technical report, Department of Homeland Security (DHS) (2009)
4. European information society. http://ec.europa.eu/information_society/policy/nis/strategy/activities/ciip/index_en.htm (2010)
5. UK resilience homepage. <http://www.cabinetoffice.gov.uk/ukresilience.aspx> (2010)
6. Afek, Y., Gafni, E.: End-to-end communication in unreliable networks. In: Proceedings of the 7th Annual ACM Symposium on the Principles of Distributed Computing, Toronto, Canada, pp. 131–148 (1988)
7. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* **1**(1), 11–33 (2004)
8. Bhattacharjee, B., Calvert, K., Griffioen, J., Spring, N., Sterbenz, J.P.G.: Postmodern Internetwork Architecture. Technical Report ITTC-FY2006-TR-45030-01, The University of Kansas, Lawrence, KS (2006)
9. Burleigh, S., Hooke, A., Torgerson, L., Fall, K., Cerf, V., Durst, B., Scott, K., Weiss, H.: Delay-tolerant networking: an approach to interplanetary internet. *IEEE Communications Magazine* **41**(6), 128–136 (2003)
10. Carter, M.R., Howard, M.P., Owens, N., Register, D., Kennedy, J., Pecheux, K., Newton, A.: Effects of catastrophic events on transportation system management and operations, Howard Street tunnel fire, Baltimore City, Maryland – July 18, 2001. Tech. rep., U.S. Department of Transportation, ITS Joint Program Office, Washington DC (2002)
11. Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.: Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach. *Springer Telecommunication Systems* pp. 1–16 (2011). DOI 10.1007/s11235-011-9575-4. Published online: 21 September 2011
12. Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.G.: A comprehensive framework to simulate network attacks and challenges. In: Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 538–544. Moscow (2010)
13. Clark, D.D.: Protocol design and performance. tutorial notes, IEEE INFOCOM (1995)
14. Cowie, J.H., Ogielski, A.T., Premore, B., Smith, E.A., Underwood, T.: Impact of the 2003 Blackouts on Internet Communications. Preliminary report, Renesys Corporation (2003). (updated March 1, 2004)
15. Demeester, P., Gryseels, M., Autenrieth, A., Brianza, C., Castagna, L., Signorelli, G., Clemence, R., Ravera, M., Jajszyk, A., Janukowicz, D., Doorselaere, K.V., Harada, Y.: Resilience in multilayer networks. *IEEE Communications Magazine* **37**(8), 70–76 (1999)
16. Ellison, R.J., Fisher, D.A., Linger, R.C., Lipson, H.F., Longstaff, T., Mead, N.R.: Survivable network systems: An emerging discipline. Tech. Rep. CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University, PA (1999)
17. Fall, K.: A delay-tolerant network architecture for challenged internets. In: SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 27–34. ACM, New York, NY, USA (2003)
18. Fall, K., Farrell, S.: DTN: An architectural retrospective. *IEEE Journal on Selected Areas in Communications (JSAC)* **26**(5), 828–836 (2008)
19. Feldmeier, D.: An overview of the TP++ transport protocol project. In: A.N. Tantawy (ed.) *High Performance Networks: Frontiers and Experience, Kluwer International Series in Engineering and Computer Science*, vol. 238, chap. 8. Kluwer Academic Publishers, Boston, MA, USA (1993)
20. Frank, H., Frisch, I.: Analysis and Design of Survivable Networks. *IEEE Transactions on Communication Technology* **18**(5), 501–519 (1970)
21. Goodman, S., Lin, H.: *Toward a Safer and More Secure Cyberspace*. National Academies Press (2007)
22. Grover, W.D., Stamatelakis, D.: Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity for self-planning network restoration. In: Proceeding of the IEEE International Conference on Communications (ICC), vol. 1, pp. 537–543 (1998)
23. Heimlicher, S., Karaliopoulos, M., Levy, H., Spyropoulos, T.: On leveraging partial paths in partially-connected networks. In: Proceeding of the 28th IEEE Conference on Computer Communications (INFOCOM). Rio de Janeiro, Brazil (2009)
24. Jabbar, A.: A framework to quantify network resilience and survivability. Ph.D. thesis, The University of Kansas, Lawrence, KS (2010)
25. Jabbar, A., Narra, H., Sterbenz, J.P.G.: An approach to quantifying resilience in mobile ad hoc networks. In: Proceedings of the 8th IEEE International Workshop on the Design of Reliable Communication Networks (DRCN), pp. 140–147. Krakow, Poland (2011)
26. Jabbar, A., Raman, B., Frost, V.S., Sterbenz, J.P.G.: Weather disruption-tolerant self-optimising millimeter mesh networks. In: Proceedings of IWSOS: Third International IFIP/IEEE Workshop on Self-Organizing Systems, *Lecture Notes in Computer Science*, vol. 5343, pp. 242–255. Springer (2008)
27. Jabbar, A., Rohrer, J.P., Frost, V.S., Sterbenz, J.P.G.: Survivable millimeter-wave mesh networks. *Computer Communications* **34**(16), 1942–1955 (2011)
28. Jabbar, A., Rohrer, J.P., Oberthaler, A., Çetinkaya, E.K., Frost, V., Sterbenz, J.P.G.: Performance comparison of weather disruption-tolerant cross-layer routing algorithms. In: Proc. IEEE INFOCOM 2009. The 28th Conference on Computer Communications, pp. 1143–1151 (2009)
29. Khabbaz, M.J., Assi, C.M., Fawaz, W.F.: *Disruption-Tolerant Networking: A Comprehensive Survey on Recent Developments and Persisting Challenges*. IEEE Communications Surveys & Tutorials (2011). DOI 10.1109/SURV.2011.041911.00093
30. Krishnan, R., Sterbenz, J.P.G., Eddy, W.M., Partridge, C., Allman, M.: Explicit transport error notification (ETEN) for error-prone wireless and satellite networks. *Computer Networks* **46**(3), 343–362 (2004)
31. Laprie, J.C.: Dependability: Basic concepts and terminology. Draft, IFIP Working Group 10.4 – Dependable Computing and Fault Tolerance (1994)
32. Liscouski, B., Elliot, W.J.: Final Report on the August 14, 2003 Blackout in the United States and Canada:

- Causes and Recommendations. Tech. rep., U.S. – Canada Power System Outage Task Force (2004)
33. Lyons, R., Vanderkulk, W.: The use of triple-modular redundancy to improve computer reliability. *IBM Journal of Research and Development* **6**(2), 200–209 (1962)
 34. Medhi, D., Tipper, D.: Multi-layered network survivability-models, analysis, architecture, framework and implementation: An overview. In: Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX), vol. 1, pp. 173–186 (2000)
 35. Menth, M., Hartmann, M., Martin, R., Cicic, T., Kvalbein, A.: Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute? *Computer Networks* **54**(8), 1300–1315 (2010)
 36. Meyer, J.F.: On Evaluating the Performability of Degradable Computing Systems. *IEEE Transactions on Computers* **100**(29), 720–731 (1980)
 37. Meyer, J.F.: Performability: a retrospective and some pointers to the future. *Performance Evaluation* **14**(3-4), 139–156 (1992)
 38. Meyer, J.F.: Performability evaluation: Where it is and what lies ahead. In: Proceedings of the IEEE International Computer Performance and Dependability Symposium (IPDS), pp. 334–343 (1995)
 39. Meyer, J.F.: Defining and evaluating resilience: A performability perspective. In: Proceedings of the International Workshop on Performability Modeling of Computer and Communication Systems (PMCCS). Eger, Hungary (2009)
 40. Molisz, W., Rak, J.: End-to-end service survivability under attacks on networks. *Journal of Telecommunications and Information Technology* **3**, 19–26 (2006)
 41. Narra, H., Çetinkaya, E.K., Sterbenz, J.P.: Performance Analysis of AeroRP with Ground Station Advertisements. In: Proceedings of the ACM MobiHoc Workshop on Airborne Networks and Communications. Hilton Head Island, SC (2012). (to appear)
 42. Nicol, D.M., Sanders, W.H., Trivedi, K.S.: Model-based evaluation: From dependability to security. *IEEE Transactions on Dependable and Secure Computing* **01**(1), 48–65 (2004)
 43. Nussbaumer, J., Patel, B.V., Schaffa, F., Sterbenz, J.P.G.: Networking requirements for interactive video on demand. *IEEE Journal on Selected Areas in Communications* **13**, 779–787 (1995)
 44. Peters, K., Jabbar, A., Çetinkaya, E.K., Sterbenz, J.P.: A geographical routing protocol for highly-dynamic aeronautical networks. In: Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), pp. 492–497. Cancun, Mexico (2011)
 45. Rak, J.: k -penalty: A novel approach to find k -disjoint paths with differentiated path costs. *IEEE Communications Letters* **14**(4), 354–356 (2010)
 46. Rak, J.: Fast Service Recovery Under Shared Protection in WDM Networks. *IEEE/OSA Journal of Lightwave Technology* **30**(1), 84–95 (2012)
 47. Rak, J., Walkowiak, K.: Reliable anycast and unicast routing: protection against attacks. *Springer Telecommunication Systems* pp. 1–18 (2011). DOI 10.1007/s11235-011-9583-4. Published online: 1 September 2011
 48. Rohrer, J.P.: End-to-end resilience mechanisms for network transport protocols. Ph.D. thesis, The University of Kansas, Lawrence, KS (2011)
 49. Rohrer, J.P., Çetinkaya, E.K., Narra, H., Broyles, D., Peters, K., Sterbenz, J.P.G.: AeroRP Performance in Highly-Dynamic Airborne Networks using 3D Gauss-Markov Mobility Model. In: Proceedings of the IEEE Military Communications Conference (MILCOM), pp. 834–841. Baltimore, MD (2011)
 50. Rohrer, J.P., Jabbar, A., Çetinkaya, E.K., Perrins, E., Sterbenz, J.P.: Highly-dynamic cross-layered aeronautical network architecture. *IEEE Transactions on Aerospace and Electronic Systems* **47**(4), 2742–2765 (2011)
 51. Rohrer, J.P., Jabbar, A., Çetinkaya, E.K., Sterbenz, J.P.: Airborne telemetry networks: Challenges and solutions in the ANTP suite. In: Proceedings of the IEEE Military Communications Conference (MILCOM), pp. 74–79. San Jose, CA (2010)
 52. Rohrer, J.P., Jabbar, A., Perrins, E., Sterbenz, J.P.G.: Cross-layer architectural framework for highly-mobile multihop airborne telemetry networks. In: Proceedings of the IEEE Military Communications Conference (MILCOM), pp. 1–9. San Diego, CA, USA (2008)
 53. Rohrer, J.P., Jabbar, A., Sterbenz, J.P.G.: Path diversification. *Springer Telecommunication Systems* (to appear in 2012)
 54. Rohrer, J.P., Jabbar, A., Sterbenz, J.P.G.: Path diversification: A multipath resilience mechanism. In: Proceedings of the IEEE 7th International Workshop on the Design of Reliable Communication Networks (DRCN), pp. 343–351. Washington, DC (2009)
 55. Rohrer, J.P., Naidu, R., Sterbenz, J.P.G.: Multipath at the transport layer: An end-to-end resilience mechanism. In: Proceedings of the IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 1–7. St. Petersburg, Russia (2009)
 56. Rohrer, J.P., Sterbenz, J.P.G.: Predicting topology survivability using path diversity. In: Proceedings of the IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 95–101. Budapest (2011)
 57. Saltzer, J.H., Reed, D.P., Clark, D.D.: End-to-end arguments in system design. *ACM Transactions Computer Systems* **2**(4), 277–288 (1984)
 58. Schneider, F.: Trust in Cyberspace. National Academies Press (1999)
 59. Scott, K., Burleigh, S.: Bundle Protocol Specification. RFC 5050 (Experimental) (2007). URL <http://www.ietf.org/rfc/rfc5050.txt>
 60. Sterbenz, J.P., Çetinkaya, E.K., Hameed, M.A., Jabbar, A., Shi, Q., Rohrer, J.P.: Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation (invited paper). *Springer Telecommunication Systems* pp. 1–32 (2011). DOI 10.1007/s11235-011-9573-6. Published online: 7 December 2011
 61. Sterbenz, J.P., Saxena, T., Krishnan, R.: Latency-Aware Information Access with User-Directed Fetch Behaviour for Weakly-Connected Mobile Wireless Clients. Technical Report 8340, BBN Technologies, Cambridge, MA (2002)
 62. Sterbenz, J.P.G., Hutchison, D.: Resilinet: Multilevel resilient and survivable networking initiative wiki. <http://wiki.ittc.ku.edu/resilinet> (2006)
 63. Sterbenz, J.P.G., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks* **54**(8), 1245–1265 (2010)
 64. Sterbenz, J.P.G., Krishnan, R., Hain, R.R., Jackson, A.W., Levin, D., Ramanathan, R., Zao, J.: Survivable

- mobile wireless networks: issues, challenges, and research directions. In: Proceedings of the 3rd ACM workshop on Wireless Security (WiSE), pp. 31–40. Atlanta, GA (2002)
65. Sterbenz, J.P.G., Touch, J.D.: High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication, 1st edn. Wiley (2001)
 66. Strand, J., Chiu, A., Tkach, R.: Issues for routing in the optical layer. *IEEE Communications Magazine* **39**(2), 81–87 (2001)
 67. Styron, H.C.: CSX tunnel fire: Baltimore, MD. US Fire Administration Technical Report USFA-TR-140, Federal Emergency Management Administration, Emmitsburg, MD (2001)
 68. Swartz, M., Wallace, D.: Effects of Frame Rate and Resolution Reduction on Human Performance. In: IS&T's 46th Annual Conference. Munich (1993)
 69. Tchakountio, F., Ramanathan, R.: Anticipatory routing for highly mobile endpoints. In: Proceedings of the 6th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA), pp. 94–101. Washington, DC (2004)
 70. Trivedi, K., Kim, D., Roy, A., Medhi, D.: Dependability and security models. In: Proceedings of the International Workshop of Design of Reliable Communication Networks (DRCN), pp. 11–20. IEEE (2009)
 71. Zhao, W., Ammar, M., Zegura, E.: A message ferrying approach for data delivery in sparse mobile ad hoc networks. In: Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), pp. 187–198. Tokyo (2004)

Bios

Dr. James P.G. Sterbenz: is Associate Professor of Electrical Engineering & Computer Science and on staff at the Information & Telecommunication Technology Center at The University of Kansas, and is a Visiting Professor of Computing in InfoLab 21 at Lancaster University in the UK. He received a doctorate in computer science from Washington University in St. Louis in 1991, with undergraduate degrees in electrical engineering, computer science, and economics. He is director of the ResiliNets research group at KU, PI for the NSF-funded FIND Postmodern Internet Architecture project, PI for the NSF Multilayer Network Resilience Analysis and Experimentation on GENI project, lead PI for the GpENI (Great Plains Environment for Network Innovation) international GENI and FIRE testbed, co-I in the EU-funded FIRE ResumeNet project, and PI for the US DoD-funded highly-mobile airborne networking project. He has previously held senior staff and research management positions at BBN Technologies, GTE Laboratories, and IBM Research, where he has lead DARPA- and internally-funded research in mobile, wireless, active, and high-speed networks. He has been program chair for IEEE GI, GBN, and HotI; IFIP IWSOS, PfHNS, and IWAN; and is on the editorial board of IEEE Network. He has been active in Science and Engineering Fair organisation and judging in

Massachusetts and Kansas for middle and high-school students. He is principal author of the book High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication. He is a member of the IEEE, ACM, IET/IEE, and IEICE. His research interests include resilient, survivable, and disruption tolerant networking, future Internet architectures, active and programmable networks, and high-speed networking and systems.

Dr. David Hutchison: is Director of InfoLab21 and Professor of Computing at Lancaster University and has worked in the areas of computer communications and networking for more than 25 years. He has recently focused his research efforts towards network resilience. He has completed many UK, European and industry-funded research contracts and published many papers as well as writing and editing books on these and related areas. He has been an expert evaluator and member or chair of various advisory boards and committees in the UK (EPSRC, DTI, OFTEL, e-Science, UKLight, UKCRC, JISC, DC-KTN) and within the EU through several Framework Programmes. Also, he has served as member or chair of numerous TPCs (including the flagship ACM SIGCOMM and IEEE Infocom), and of journal editorial boards. He is an editor of the renowned Lecture Notes in Computer Science and of the Wiley CNDIS book series.

Egemen K. Çetinkaya: is a Ph.D. candidate in the department of Electrical Engineering and Computer Science at The University of Kansas. He received the B.S. degree in Electronics Engineering from Uludağ University (Bursa, Turkey) in 1999 and the M.S. degree in Electrical Engineering from University of Missouri–Rolla in 2001. He held various positions at Sprint as a support, system, design engineer from 2001 until 2008. He is a graduate research assistant in the ResiliNets research group at the KU Information & Telecommunication Technology Center (ITTC). His research interests are in resilient networks. He is a member of the IEEE Communications Society, ACM SIGCOMM, and Sigma Xi.

Dr. Abdul Jabbar: is currently a Research Engineer in the Advanced Communication Systems Lab at GE Global Research in Nikayuna, NY. He also holds the position of Adjunct Research Associate with the University of Kansas. He received his Ph.D in Electrical Engineering from The University of Kansas in 2010 with honors. He received his M.S. degree in Electrical Engineering from KU in 2004 and B.S. degree in Electrical Engineering from Osmania University, India in 2001. His interests include resilience and survivability, network algorithms, design and analysis of network architectures, topologies, and protocols, highly dynamic

networks, wireless access, and future networks. Abdul is the recipient of Moore award for best M.S. thesis and is a member of IEEE Communications Society, IEEE Computer Society, and the ACM Special Interest Group on Data Communications.

Dr. Justin P. Rohrer: is currently a Research Associate of Computer Science at the Naval Postgraduate School (NPS) and an Adjunct Assistant Professor of Electrical Engineering and Computer Science at the KU Information & Telecommunication Technology Center (ITTC). He received his Ph.D in Electrical Engineering from the University of Kansas in 2011 with honors. He received his B.S. degree in Electrical Engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2004. From 1999 to 2004, he was with the Adirondack Area Network, Castleton, NY as a network engineer. He was also an ITTC Graduate Fellow from 2004–2006. He received the best paper award at the International Telemetry Conference in 2008 and the best graduate student paper award at the same conference in 2011. His research focus is on resilient and survivable transport and routing protocols. Interests also include highly-dynamic mobile networks, and simulating network disruptions. Previous research has included weather disruption-tolerant mesh networks and free-space optical metropolitan networks. He is a member of the IEEE Communications and Computer Societies, ACM SIGCOMM, Eta Kappa Nu, and was an officer of the Kansas City section of the IEEE Computer Society.

Dr. Marcus Schöller: is a research scientist at NEC Laboratories Europe, Germany. He received the diploma in computer science at University of Karlsruhe, Germany, in 2001 and his doctorate in engineering in 2006 on robustness and stability of programmable networks. Afterwards he held a postdoc position at Lancaster University, UK, focusing his research on autonomic networks and network resilience. Marcus is currently working on the EU FP7 projects ResumeNet, with a focus on future network architecture with resilience as a key property, and the EU FP7 BeFemto project, which investigates next generation LTE-A femtocell technologies and business opportunities. His interests also include network and system security, intrusion detection, self-organization of networks, future network architectures, mobile networks including mesh and opportunistic networks.

Dr. Paul Smith: is a Senior Scientist in the Safety and Security Department of the AIT, Austrian Institute of Technology. Previous to this appointment he was a Senior Research Associate at Lancaster University, UK. He received his Ph.D. in September 2003 and graduated in 1999 with an honours degree in Computer Science from Lancaster. Pauls research interests lie in the var-

ious ways that networked (socio-technical) systems fail to provide a desired service when under duress from various challenges, such as attacks and mis-configurations. He has participated in a number of international research projects, including the EU-funded ResumeNet project which investigated a framework for network resilience. Currently, he is primarily working on the EU-funded PRECYSE project, which is investigating the security and resilience of critical information infrastructures.