The University of Kansas

INFORMATION & TELECOMMUNICATION TECHNOLOGY CENTER
The University of Kansas

Technical Report

# Multilayer Network Resilience Analysis and Experimentation on GENI

James P.G. Sterbenz,
Justin P. Rohrer, and Egemen Çetinkaya

ITTC-FY2011-TR-61349-01

July 2010

Project Sponsor:
The National Science Foundation

# Project Summary

## EAGER: Multilayer Network Resilience Analysis and Experimentation on GENI

As society increasingly depends on networks for all aspects of its functioning, the consequences of its disruption become increasingly severe, as does its attractiveness to crackers and terrorists who wish to disrupt society. It is therefore essential to understand and evaluate the resilience of current networks, and to evaluate alternatives for future Internet architecture, mechanisms, and protocols. We define *resilience* as the ability of the network to deliver desired service when subject to a variety of challenges, including large-scale disasters and attacks.

GENI is evolving to provide a promising environment in which to do experimental research in the resilience and survivability of future networks, by allowing programmable control over topology and mechanism, while providing the scale and global reach needed to conduct network experiments far beyond the capabilities of a conventional testbed. We propose to use GENI to perform resilience and survivability experiments at at scale, both in terms of node count and with the geographic scope needed to emulate area-based challenges such as large-scale disasters. The goal is to advance the state-of-the-art in resilient network design based on experimental research, and to better understand the application of large-scale programmable infrastructure to performing network research, and its relationship to analytical and simulation-based techniques.

**INTELLECTUAL MERIT.** The proposed approach to this exploratory research in future Internet resilience is to leverage analytical and simulation-based research done under the NSF FIND and EU FIRE programs as a basis for experimental research using GENI. In particular, the PlanetLab and ProtoGENI control frameworks will be used for wide-area and access network emulation, respectively, with heavy use of the GpENI PlanetLab-based substrate that provides an international footprint and is also federated with the public PlanetLab and directly connected to the ProtoGENI Internet2 wavelength.

We will use the results of the earlier research to drive and cross-verify with the GENI-based experimental research, with emphasis on geographic and mechanism diversity as the basis for survivability against attacks and large-scale disasters such as hurricanes and power blackouts. We will rely on the large scale, geographic scope, and programmability of GENI to permit experiments that evaluate the performance of the network when subject to correlated failures. We will concentrate our research on metrics to measure the resilience of alternative topologies and mechanisms, and the performance of multipath routing and end-to-end transport to provide resilience. We will work with other program participants, particularly for measurement and monitoring of our experiments.

**BROADER IMPACT.** The proposed research aims to gain new understanding on how to design and deploy resilient networks based on experimentation using GENI. This supports not only the long-term vision of using GENI to support Future Internet Design, but specific techniques and mechanisms may be applied to improve the resilience of the current Internet and its parts to better serve society. This project will also contribute to the education of graduate students in KU EECS 983: Resilient and Survivable Networks, whose students will be encouraged to use GENI for class projects, including female graduate students in the ResiliNets research group.

**Key Words:** GENI, GpENI, PlanetLab, VINI, ProtoGENI, OpenFlow, SPP, Future Internet Design (FIND); resilient, survivable, and dependable networks; network topology analysis; multipath geographically diverse routing; end-to-end transport, experimental simulation methodology

# Project Description

## 1 Introduction and Motivation

As society increasingly depends on networks for all aspects of its functioning, the consequences of its disruption become increasingly severe, as does its attractiveness to crackers and terrorists who wish to disrupt society. We define *resilience* as the ability of the network to deliver desired service when subject to a variety of challenges, including large-scale disasters and attacks [1]. *Survivability* is an important aspect of resilience that related to correlated failures, such as attacks against key parts of the infrastructure and large-scale disasters such as hurricanes and power blackouts [2, 3]. The Internet and all of its interconnected subnetworks (including wireless sensor networks and SCADA – supervisory control and acquisition) is one of our critical national infrastructures, and is interdependent on other critical infrastructures, such as the power grid. It is therefore essential to understand and evaluate the resilience of current networks, and to evaluate alternatives for future Internet architecture, mechanisms, and protocols.

## 2 Proposed Research

GENI is evolving to provide a promising environment in which to do experimental research in the resilience and survivability of future networks, by allowing programmable control over topology and mechanism, while providing the scale and global reach needed to conduct network experiments far beyond the capabilities of a conventional testbed. We propose to use GENI in general, and the GpENI infrastructure [4,5] (expanding to 40 clusters with 200 nodes worldwide), federated with the larger GENI PlanetLab control framework and interconnected to several ProtoGENI [6] facilities to perform resilience and survivability experiments at at scale, both in terms of node count and with the geographic scope needed to emulate area-based challenges such as large-scale disasters. Furthermore, we propose to use these experiments to cross-verify with analytical and simulation-based resilience research currently underway at The University of Kansas leveraging topology and challenge generation tools (KU-LoCGen and KU-CSM) developed for this purpose, with emphasis on resilience metrics [7, 8] and multi-path multi-realm diverse transport developed as part of our NSF FIND research [9].

Since the proposed research will be conducted on the evolving GENI infrastructure, there is risk involved, mitigated by the direct access to the GpENI project at KU over which the investigator has direct control. Furthermore there is no current NSF program directly supporting experimental research on GENI; the GENI program itself only supports design and deployment of the infrastructure. These two factors support this project as an EAGER.

### 2.1 Exploratory Research

This section describes the research aspects of this proposed projects, with a few example simulation-based results to give an idea of the *experimental* research we propose to apply to GENI. Considerably more details on the simulation results are in the referenced works.

### 2.1.1 Realistic Topology Generation and Challenge Simulation

Realistic topology generators are crucial to networking research in terms of design, optimization, and analysis, however currently available topology models do not take into account constraints on geographic node locations and the accompanying effects on link costs. To remedy this we are building KU-LocGEN, a topology generator that represents the tiered architecture of the current Internet, while constrained with geographical node coordinates to reflect the realism behind the evolution of a network [10–13].

Figure 1: Sample challenge in Midwest US

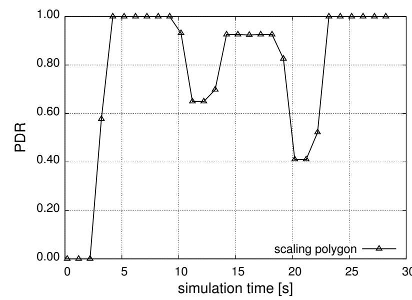


Figure 2: PDR during the challenge

To understand the effects of challenges on the network, we are building a challenge generator model KU-CSM [14–16] in ns-3 to simulate various challenges, including random software and hardware failures, malicious attacks, and geographically correlated failures that represent a large-scale natural or human-made disaster. Figure 1. shows an example of how we apply area based challenges to the network. As the challenge increases in size, the overall PDR (packet delivery ratio is affected, as shown by Figure 2 (the intermediate rise in PDR as the area increases is due to route reconvergence). We propose to apply these techniques experimentally using GENI .

### 2.1.2 End-to-End Multipath

We are exploring a number mechanisms to increase resilience and survivability, with emphasis on *diversity* in geographic path and mechanism, such as geographically diverse paths, one of which is wired and one of which is wireless. We are developing a multipath selection algorithm *Path Diversification* [9], that uses maximally-disjoint paths based on the degree of diversity required for a particular application, while meeting selectable constraints such as path stretch. Path diversification is flexible enough to be used at the network layer, in conjunction with a topology discovery mechanism such as OSPF link-state advertisements [17], to form a multipath routing protocol. It may also be implemented at the end-to-end layer in a source-routed environment, to form a multipath transport protocol [18]. We are currently examining the advantages and disadvantages of each of these scenarios in simulations.

To briefly summarize our diversity metric and its use, let the shortest path between a given ($s$, $d$) pair be $P_0$. Then, for any other path $P_k$ between the same source and destination, we define the diversity function $D(x)$ with respect to $P_0$ as shown in Figure 3. The path diversity has a value of 1 if $P_k$ and $P_0$ are completely disjoint and a value of 0 if $P_k$ and $P_0$ are identical. Figure 3 shows the shortest path, $P_0$, along with the alternate paths $P_1$ and $P_2$ both of which have a (link) novelty

$$D(P_k) = 1 - \frac{|P_k \cap P_0|}{|P_0|} \tag{1}$$

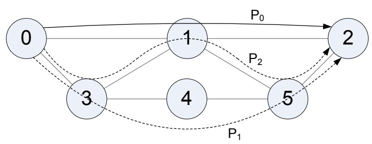


Figure 3: Diverse paths example

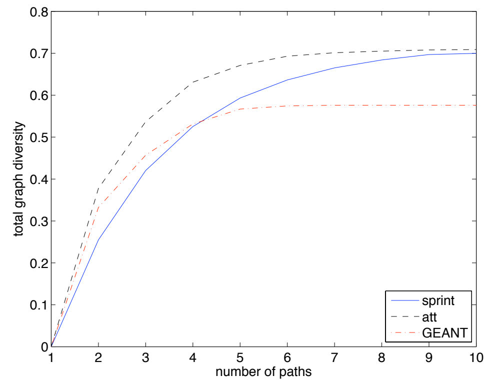


Figure 4: Graph diversity comparison

of 1 [19]. Given a failure on node 1, both $P_0$ and $P_2$ will fail. In our approach, $D(P_2) = \frac{2}{3}$, which reflects this vulnerability. $P_1$ on the other hand has a diversity of 1, and does not share any common point of failure with $P_0$. We then apply the path diversity measure to find sets of diverse paths between all node pairs in a network graph, using various constraints (number of paths, stretch) to bound the set of paths selected, and call the result *total graph diversity* (TGD). This allows us to compare the diversity available in multiple topologies (Figure 4). As part of this work we intend to extend the diversity metric with the concept of distance, so that paths may be more or less diverse based on *how far* apart they are, not merely the fact that they are different.

## 2.2 Details of Experiments

The thrust of this proposal is to use GENI to perform experimental research in resilience and survivability, at a scale previously not possible. We also seek to better understand the methodology that ties analytical and simulation-based research, by using tools such as our topology generator to drive the experiments, and to cross-verify experimental with simulation results.

### 2.2.1 Methodology and cross-verification

Resilient topologies generated by KU-LoCGen and analysed by KU-CSM will be used to generate layer-2 topologies that configure the topology of GENI (GpENI and ProtoGENI) experiments. We will evaluate performance when GENI slices are challenged by correlated failures of nodes and links, measuring connectivity, packet delivery ratio, goodput, and delay, when subject to CBR, bulk data transfer, and transactional (HTTP) traffic. We will also characterize the packet-loss probability of wireless links at the Utah Emulab [20], and the capabilities for emulating jamming and misbehaving nodes within the Emulab-federated CMU wireless emulator. In parallel with understanding the behavior of the GENI PlanetLab and ProtoGENI aggregates when programmed with challenges, we will also use the workflow infrastructure (using GUSH and Raven) to deploy experiments on these aggregates in an automated and repeatable manner. We will use monitoring tools developed by the GENI community as well as our own instrumentation to measure performance.

### 2.2.2 Large-scale experimentation

We will then run large scale resilience experiments over interconnected aggregates using DCN [21] (within GpENI) and OpenFlow and configured paths, with VINI/Planetlab layer-3 topologies, to emulate both existing ISP and synthetic topologies. Over these topologies we will run our multipath-aware transport protocol ResTP to evaluate its performance under varying application and traffic loads. Based on the output of our challenge generation simulations, we will selectively disable node slivers and links to emulate correlated network failures, and attacks. We also plan to use the wireless emulator under the ProtoGENI framework to emulate jamming attacks to wireless access networks. Each set of challenges will be classified as a single scenario, and each scenario will be run multiple times so that we can establish reasonable confidence in the results.

## 2.3 GENI Resources Required

To perform these experiments, several major parts of GENI are required: the PlanetLab and ProtoGENI control frameworks, along with the GpENI and federated substrate infrastructures. GpENI [4, 5] is an international programmable testbed centered on a regional optical network in the Midwest US funded in part by the NSF GENI program, expanding to approximately 40 institutions in Europe (funded in part by the EU FIRE Future Internet Research and Experimentation programme), Asia, and Canada. GpENI is programmable from layers 1 to 7 with DCN (dynamic circuit network) and VINI integrated with the PlanetLab control framework in GENI Cluster B. Additionally, we may be able to exploit OpenFlow capabilities being deployed in GENI. Finally, we will depend on monitoring tools under development by the community to enhance our own experiment instrumentation. This section describes the required and desirable GENI resources required to execute the proposed experimental research.

### 2.3.1 Network Programmability

To perform experiments in which new network topologies, mechanisms, and protocols are proposed to enhance resilience and survivability, it is essential to have programmable control of these aspects. At the lowest level, we need to control the layer 2 topology, particularly with respect to redundancy and geographic diversity, so we can experiment with network topologies that attempt to maintain connectivity even when network components fail or are destroyed. We therefore need the ability to arbitrarily interconnect nodes provided by ProtoGENI [6] in the local clusters and by VLAN configuration in the wide-area GENI, supplemented by the dynamic control of layer-2 connectivity provided in GpENI by DCN [21] and layer-3 topologies provided by the GENIwrapper version of VINI [22]. We hope that the planned unification of PlanetLab and VINI will give us access to larger-scale topologies throughout the larger federated PlanetLabs (including OneLab [23] in Europe).

At the next level, we need the ability to program routing functionally to evaluate geographically-diverse multipath routing, and its interaction with resilient transport with the topology is challenged. In the local area, we will instantiate our resilient routing algorithms on ProtoGENI clusters, and in the wide area use the programmable routing functionality in GpENI (Quagga and XORP) as well as OpenFlow [24] as it becomes available.

At the top level, we need the ability to deploy our resilient multipath transport protocol and path diversification mechanisms (described in Section 2.1.2) on a significant number of end systems and run traffic-generating applications, to experiment at large scale. We will need to dedicate some ProtoGENI nodes as end systems, as well as PlanetLab client machines throughout GENI, including GpENI (approximately 80 end-nodes throughout 40 sites in the US, Europe, and Asia) and G-Lab [25] in Germany (which maintains a GpENI node cluster).

### 2.3.2 Topology and Connectivity

To perform realistic resilience and survivability experiments, infrastructure is required that provides the necessary scale and geographic scope. To emulate area-based challenges such as large-scale natural disasters and power blackouts, the wide geographic footprint provided by PlanetLab [26] and GpENI is needed. To provide rich access network topologies, scenarios, and application traffic, ProtoGENI clusters on the PlanetLab/GpENI wide area network are needed, supplemented by the G-Lab cluster in Germany. Topologies will be driven by the KU-LoCGen tool previously described to generate real-world scenarios with the desired resilience properties (degree of connectivity, betweenness, path diversity, etc.).

## 2.4 Why GENI is Required

While the previous section describes the GENI resources required, this section emphasizes why extant testbeds are not sufficient for large-scale resilience and survivability research, and how the evolving GENI infrastructure serves as an enabler. GENI provides several key capabilities that are critical to large-scale network experiments that conventional research testbeds such as PlanetLab and Emulab lack. In particular, four characteristics provided by GENI are critical to the proposed experimental research: scope, scale, network programability, and hierarchy. To understand the effects of area-based challenges such as wide-area power blackouts and large-scale natural disasters, the national and global *scope* of GENI is needed, consisting of GpENI and federated GENI PlanetLab nodes, including G-Lab. While GENI is still in the early stages of construction, sufficient *scale* is currently being deployed to test both wide-area tier-1 backbone resilience (using GENI PlanetLab and GpENI), as well access network resilience using ProtoGENI clusters. GENI provides the *network programmability* not available in the traditional PlanetLab, permitting experiments that manipulate the topology to evaluate diversity and redundancy structures, allowing nodes and links to be taken down based on their geographic location, and permitting experiments with new routing algorithms and cross-layering with end-to-end transport. This is particularly important for our path-diverse resilient transport protocol research. The hierarchy provided by the wide-area PlanetLab control framework (coupled with topology control provided by VINI and programmable router support in GpENI, SPP [27] in the US GENI backbone, and OpenFlow) above local-area ProtoGENI clusters provides the ability to experiment with the hierarchical structures present in real networks. This hierarchy is not only crucial to our experiments, but we expect that this work will serve as a demonstration of the need for both the PlanetLab and ProtoGENI control frameworks. The physical interconnection of the GpENI midwest optical backbone to ProtoGENI and SPP in the Kansas City Internet2 PoP and recent work done by ProtoGENI, PlanetLab, and the GPO to provide a common interface to these control frameworks serve as a key enabler.

# References Cited

[1] James P. G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks: Special Issue on Resilient and Survivable Networks (COMNET)*, 54(8):1245–1265, June 2010.

[2] R.J. Ellison, D. Fisher, R.C. Linger, H.F. Lipson, T. Longstaff, and N.R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University, 1997.

[3] James P. G. Sterbenz, Rajesh Krishnan, Regina Rosales Hain, Alden W. Jackson, David Levin, Ram Ramanathan, and John Zao. Survivable mobile wireless networks: issues, challenges, and research directions. In *WiSE '02: Proceedings of the 3rd ACM workshop on Wireless security*, pages 31–40, New York, NY, USA, 2002. ACM Press.

[4] James P. G. Sterbenz, Deep Medhi, Byrav Ramamurthy, Caterina Scoglio, David Hutchison, Bernhard Plattner, Tricha Anjali, Andrew Scott, Cort Buffington, Gregory E. Monaco, Don Gruenbacher, Rick McMullen, Justin P. Rohrer, John Sherrell, Pragatheeswaran Angu, Ramkumar Cherukuri, Haiyang Qian, and Nidhi Tare. The Great plains Environment for Network Innovation (GpENI): A programmable testbed for future internet architecture research. In *Proceedings of the 6th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (TridentCom)*, Berlin, Germany, May 18–20 2010.

[5] James P. G. Sterbenz, Deep Medhi, Greg Monaco, Byrav Ramamurthy, Caterina Scoglio, Baek-Young Choi, Joseph B. Evans, Don Gruenbacher, Ronqing Hui, Wesley Kaplow, Gary Minden, and Jeff Verrant. Gpeni: Great plains environment for network innovation. `http://wiki.ittc.ku.edu/gpeni`, November 2009.

[6] ProtoGENI wiki. `http://www.protogeni.net/trac/protogeni`, 2010.

[7] Abdul Jabbar Mohammad, David Hutchison, and James P. G. Sterbenz. Towards quantifying metrics for resilient and survivable networks. In *Proceedings of the 14th IEEE International Conference on Network Protocols (ICNP)*, pages 17–18, November 2006.

[8] Abdul Jabbar. *A Framework to Quantify Network Resilience and Survivability.* PhD thesis, The University of Kansas, Lawrence, KS, May 2010.

[9] Justin P. Rohrer, Abdul Jabbar, and James P. G. Sterbenz. Path diversification: A multipath resilience mechanism. In *Proceedings of the 7th International Workshop on the Design of Reliable Communication Networks (DRCN)*, pages 343–351, Washington, DC, USA, October 2009.

[10] Abdul Jabbar, Qian Shi, Egemen Çetinkaya, and James P. G. Sterbenz. KU-LocGen: Location and cost-constrained network topology generator. ITTC Technical Report ITTC-FY2009-TR-45030-01, The University of Kansas, Lawrence, KS, December 2008.

[11] Abdul Jabbar, Mahmood Hameed, Qian Shi, and James P.G. Sterbenz. Location and cost constrained multilevel topology generation (poster). http://www.ittc.ku.edu/resilinets/posters/ITTC-IAB-poster-2010-kulocgen.pdf, June 2010.

[12] Mahmood A. Hameed, Abdul Jabbar, Egemen K. Çetinkaya, and James P. G. Sterbenz. Deriving network topologies from real world constraints. *IEEE Networks*, 2010. Submitted.

[13] Mahmood Hameed, Abdul Jabbar, Jing Han, and James P.G. Sterbenz. Network topology generation from real-world constraints (poster). http://www.ittc.ku.edu/resilinets/posters/ITTC-IAB-poster-2010-topology-backbone.pdf, June 2010.

[14] Rabat A. Mahmood. Simulating challenges to communication networks for evaluation of resilience. Master's thesis, The University of Kansas, Lawrence, KS, August 2009.

[15] Egemen K. Çetinkaya, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P. G. Sterbenz. A comprehensive framework to simulate network attacks and challenges. In *International Workshop on Reliable Networks Design and Modeling (RNDM)*, Moscow, Russia, October 18–20 2010. to appear.

[16] Egemen K. Çetinkaya, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz. Challenge simulation module for evaluating resilience (poster). http://www.ittc.ku.edu/resilinets/posters/ITTC-IAB-poster-2010-KU-CSM.pdf, June 2010.

[17] J. Moy. OSPF Version 2. RFC 2328 (Standard), April 1998. Updated by RFC 5709.

[18] Justin P. Rohrer, Ramya Naidu, and James P. G. Sterbenz. Multipath at the transport layer: An End-to-End resilience mechanism. In *International Workshop on Reliable Networks Design and Modeling (RNDM)*, pages 1–7, St. Petersburg, Russia, October 2009.

[19] Murtaza Motiwala, Megan Elmore, Nick Feamster, and Santosh Vempala. Path splicing. In *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on data communication*, pages 27–38, New York, NY, USA, August 17-22 2008. ACM.

[20] Emulab: Network emulation testbed. http://www.emulab.net/, December 2009.

[21] Dynamic resource allocation via GMPLS optical network. http://dragon.maxgigapop.net/, November 2009.

[22] VINI: A virtual network infrastructure. http://www.vini-veritas.net/, November 2009.

[23] OneLab: Future internet test beds. http://www.onelab.eu/, December 2009.

[24] The OpenFlow switch consortium. http://www.openflowswitch.org/, December 2009.

[25] German lab. http://www.german-lab.de/, April 2010.

[26] PlanetLab. http://www.planet-lab.org/, November 2009.

[27] Supercharged planetlab platform (SPP) hardware components. http://wiki.arl.wustl.edu/index.php/SPP_Hardware_Components, December 2009.