Progress and Challenges in Large-Scale Future Internet Experimentation using the GpENI Programmable Testbed

Justin P. Rohrer, Egemen K. Çetinkaya, and James P.G. Sterbenz
Information and Telecommunication Technology Center
Department of Electrical Engineering and Computer Science
The University of Kansas, Lawrence, KS 66045, USA
{rohrej|ekc|jpgs}@ittc.ku.edu
http://www.ittc.ku.edu/resilinets, http://www.gpeni.net

ABSTRACT

GpENI 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. Addressing this need at scale introduces a number of challenges both in deployment and in collecting results that can be directly compared to simulation results for cross-verification purposes. In this short paper we present the scope, design goals, challenges, and current status of the GpENI programmable testbed, as well as an overview and examples of the types of experiments we are beginning to run.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design — *Distributed networks*

General Terms

Experimentation, Performance, Verification

Keywords

Programmable testbed; Future Internet design; experimental simulation methodology; GpENI, GENI, FIND, FIRE; resilient, survivable, and dependable networks; multipath geographically diverse routing; network topology; end-to-end transport

1. INTRODUCTION AND MOTIVATION

Testbeds play an important role in evaluating new protocols, and GpENI (Great Plains Environment for Network Innovation) [21] is a Future Internet research testbed that provides worldwide scalability to researchers to conduct their experiments. In this short paper we present the GpENI deployment status and experiments for which it will be used, with emphasis on network infrastructure resilience. GpENI is part of the GENI [4] and FIRE [3] programs.

2. GpENI PROGRAMMABLE TESTBED

To be useful in performing future-Internet experiments, a number of features must be present in a testbed.

2.1 Multi-Layer 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, programmability is required to control the layer 2 topology, particularly with respect to redundancy and geographic diversity, in order to enable experimentation with network topologies that attempt to maintain connectivity even when network components fail or are destroyed. In GpENI this control of layer-2 connectivity is provided in GpENI by DCN (Dynamic Circuit Network) [1].

Table 1: GpENI programmability layers

GpENI Layer		Programmability
	experiment	Gush, Raven
7	application	PlanetLab
4	end-to-end	
3	router	Quagga, XORP, Click
	topology	VINI
2	VLAN	DCN
	lightpath	
1	photonics	site-specific

At the next higher level, programmable routing functionality is enabled in GpENI using Quagga and XORP integrated into the GENIwrapper version of VINI running on dedicated nodes in each site cluster nodes [12].

At the highest levels, the ability to deploy novel transport protocols (such as our path diversification mechanism described in Section 3.3) and applications on a significant number of end systems is necessary to experiment at large scale. This is partially enabled in GpENI itself with approximately 80 PlanetLab nodes throughout 40 sites in the US, Europe, and Asia, and the ability to tie-in many more hosts from federated GENI aggregates and G-Lab [6] in Germany (which maintains a GpENI node cluster).

CFI '11 June 13–15, 2011, Seoul, Korea

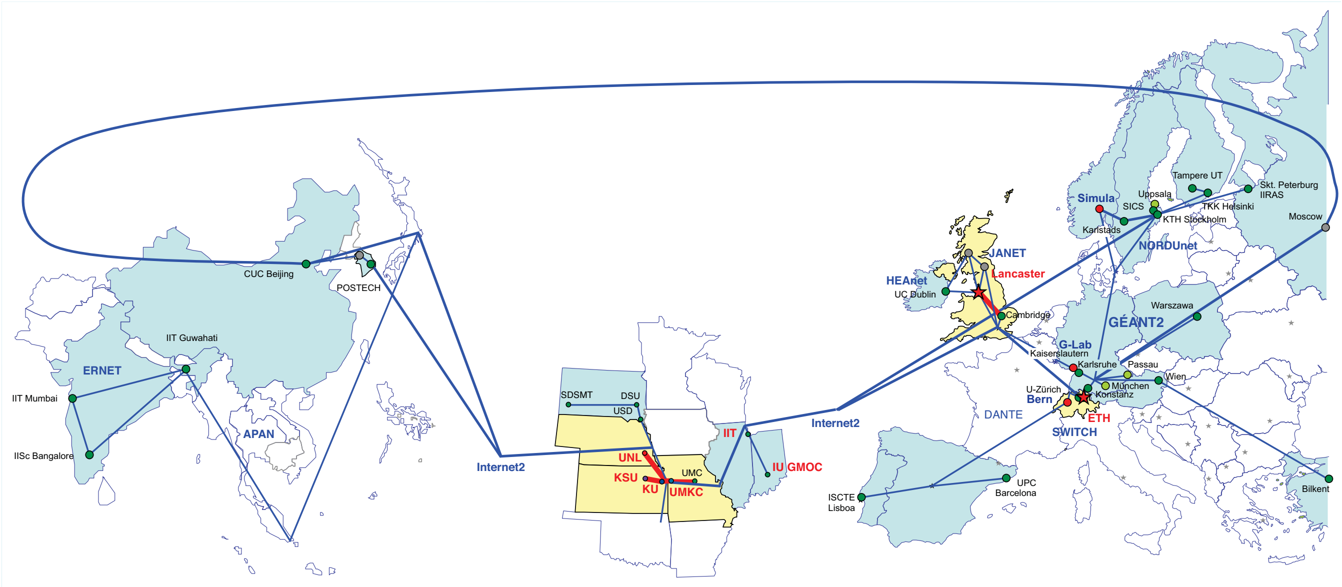

Figure 1: GpENI map

2.2 GpENI Deployment

The GpENI infrastructure [21] is in the process of expanding to 40 clusters with 200 nodes worldwide, federated with the larger GENI PlanetLab control framework and interconnected to several ProtoGENI facilities, as shown in Figure 1. This enables users 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. In our own research efforts, we are using these facilities to enable experiments that cross-verify the analytical and simulation-based resilience research currently underway at The University of Kansas [19], leveraging topology and challenge generation tools (KU-LoCGen and KU-CSM [11]) developed for this purpose, with emphasis on resilience metrics [15] and multipath multi-realm diverse transport [18, 17] developed as part of our NSF FIND research in the PostModern Internet Architecture project [9].

3. FUTURE INTERNET EXPERIMENTS

This section gives some examples of the types of research questions we expect to be able to answer through experimentation on the GpENI testbed. We also show a few simulation-based results from existing research that we are now cross-verifying experimentally. More detail on the simulations are available in the referenced works.

3.1 Realistic Topology Generation

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. KU-LocGEN is a topology generator that represents hierarchical structure of the Internet, with *geographic* node placement and cost constraints on links [14, 13, 20]. The generated topologies are then imported into ns-3 simulations or used to configure GpENI experiment topologies.

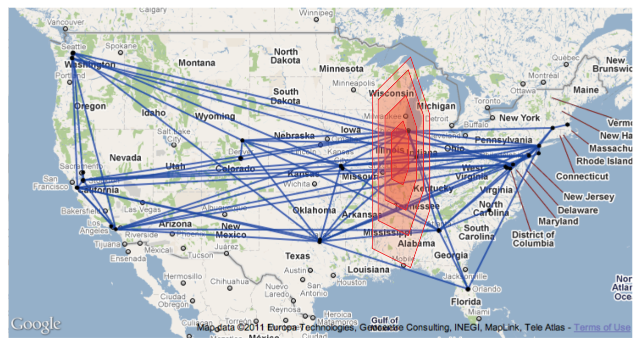

Figure 2: Sample challenge in Midwest US

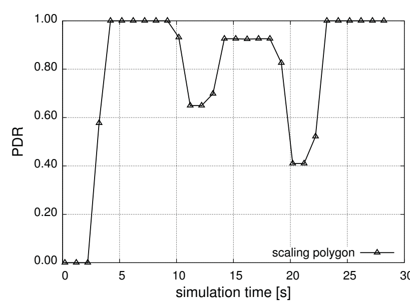

Figure 3: PDR during the challenge

3.2 Challenge Simulation

To understand the effects of challenges on the network, the KU-CSM [11, 10] challenge generator model simulates various challenges in ns-3, including random software and hardware failures, malicious attacks, and geographically correlated failures that represent a large-scale natural or human-made disaster. Figure 2 shows an example of how we apply area based challenges to the network. A set of polygons of increasing size are used to simulate a cascading power failure in the Midwestern US. As the challenge increases in size, the overall PDR (packet delivery ratio) is affected, as shown by Figure 3 (the intermediate rise in PDR as the area increases is due to route reconvergence). The topology visualization was generated using KU-TopView, our publicly available online topology viewer [8, 20]. We developed this tool to aid in manipulating, combining, and visualizing topology data from physical- and logical-layer data sets.

3.3 End-to-End Multipath

At the end-to-end layer we are developing a multipath selection algorithm *Path Diversification* [17], 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 [16], to form a multipath routing protocol, or at the end-to-end layer to form our multipath transport protocol ResTP [18]. Implementing the mechanism and experimenting in the GpENI testbed allows us to examine the advantages and disadvantages of each of these scenarios.

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

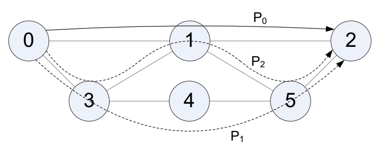

Figure 4: Diverse paths example

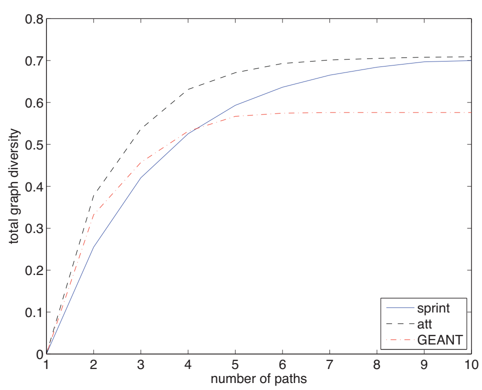

Figure 5: Graph diversity comparison

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 4. 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 4 shows the shortest path, P_0, along with the alternate paths P_1 and P_2 given a failure on node 1, both P_0 and P_2 will fail, and $D(P_2) = \frac{2}{3}$ reflecting this vulnerability. P_1 on the other hand has a diversity of 1, not sharing any common point of failure with P_0. We can then combine the available diversity between all endpoints in the graph to calculate the *total graph diversity* (TGD) to compare multiple topologies (Figure 5). We are currently extending this metric to incorporate the geographic *distance* between paths [20].

3.4 Methodology and Cross-Verification

The need for cross-verification brings up the question of *what with*? Two open-source simulators stand out, ns-2 and ns-3. Unfortunately our attempts to use the former clearly demonstrated that many of the core models produce invalid results and thus are not suitable for cross-verification. Furthermore, heterogeneous wired/wireless simulation is not supported, and models frequently require source code modifications that are then tied to a particular release. Ns-3 is taking a more rigorous and modular approach, however it is much less established, lacking a large library of common models. Examples of this include the DSDV and DSR models, an HTTP model, and a simple TDMA MAC, all of which our group has developed and contributed to the ns-3 releases. We are hopeful that if others take the same approach and contribute a few basic models, the ns-3 platform will soon be a suitable cross-verification tool for testbed experimentation.

Resilient topologies generated by KU-LoCGen and analyzed by KU-CSM are used to generate layer-2 topologies that configure the topology of GpENI experiments. We evaluate performance when slice topologies 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 also characterize the packet-loss probability of wireless links at the Utah Emulab, and the capabilities for emulating jamming and misbehaving nodes within the Emulab-federated CMU wireless emulator [2]. Workflow infrastructure is provided by Raven [5] to deploy experiments on these aggregates in an automated and repeatable manner.

3.5 Large-Scale Deployment

In order to provide the ability to experiment with non-IP network layers, the GENI federation has converged on ethernet VLANs as the common denominator across all testbeds. There are several resulting implications and technical challenges, for example, no matter the scale of the testbed (global in GpENI's case), it is one giant broadcast domain given the capabilities of commodity ethernet switches, and the usable L2 topology is restricted to a tree. The cost of native layer-2 interconnection on a global scale is also high. To address these challenges GpENI is deploying a number of emerging technologies both to manage the testbed itself as well as addressing the needs of experimenters. DCN (previously mentioned) is one such tool which establishes VLAN circuits across the testbed to manage broadcast traffic and provided a layer-2 point-to-point abstraction for experimenters. We have used L2TPv3 tunnels over IP research networks to mitigate the cost of long-distance layer-2 connectivity, however this still is limited to a tree topology. The *tinc* [7] project goes a step further, allowing the creation of a full mesh of VPN L2 tunnels while preventing broadcast storms and is a promising solution to these challenges.

Large scale resilience experiments are run over interconnected aggregates using DCN [1] (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 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 selectively disable node slivers and links to emulate correlated network failures and attacks. In the future we will also use the wireless emulator under the ProtoGENI framework to emulate jamming attacks to wireless access networks. Each challenge set is classified as a single scenario, and each scenario is run multiple times to establish reasonable confidence in the results.

4. CONCLUSIONS

Testbeds are essential for evaluating new protocols and the performance of new network architectures. GpENI is a Future Internet programmable research testbed that provides worldwide scalability to researchers to conduct their experiments. In this short paper we presented the GpENI programmable testbed deployment status and an overview of the type of experimentation it supports.

5. ACKNOWLEDGMENTS

This research was supported in part by NSF 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 by the EU FP7 FIRE programme ResumeNet project (grant agreement no. 224619).

6. REFERENCES

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

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

[3] FIRE: Future internet research & experimentation. `http://cordis.europa.eu/fp7/ict/fire/`, December 2009.

[4] GENI: Global environment for network innovations. `http://www.geni.net/`, December 2009.

[5] Raven provisioning service. `http://raven.cs.arizona.edu/`, December 2009.

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

[7] tinc wiki. `http://www.tinc-vpn.org/`, 2010.

[8] Resilinets topology map viewer. `http://www.ittc.ku.edu/resilinets/maps/`, January 2011.

[9] B. Bhattacharjee, K. Calvert, J. Griffioen, N. Spring, and J. P. G. Sterbenz. Postmodern internetwork architecture. Technical Report ITTC-FY2006-TR-45030-01, Information and Telecommunication Center, 2335 Irving Hill Road, Lawrence, KS 66045-7612, February 2006.

[10] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. Sterbenz. Modelling communication network challenges for future internet resilience, survivability, and disruption tolerance: A simulation-based approach. *Springer Telecommunication Systems*, 2011. (accepted March 2011).

[11] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. G. Sterbenz. 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)*, pages 538–544, Moscow, Russia, October 18–20 2010.

[12] R. Cherukuri, X. Liu, A. Bavier, J. Sterbenz, and D. Medhi. Network virtualization in GpENI: Framework, implementation & integration experience. In *IEEE/IFIP ManFI*, Dublin, Ireland, May 2011.

[13] M. A. Hameed, A. Jabbar, E. K. Çetinkaya, and J. P. Sterbenz. Deriving network topologies from real world constraints. In *Proceedings of IEEE GLOBECOM Workshop on Complex and Communication Networks (CCNet)*, pages 415–419, Miami, FL, December 2010.

[14] A. Jabbar, Q. Shi, E. Çetinkaya, and J. 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.

[15] A. J. Mohammad, D. Hutchison, and J. 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.

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

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

[18] J. P. Rohrer, R. Naidu, and J. P. G. Sterbenz. 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)*, pages 1–7, St. Petersburg, Russia, October 2009.

[19] J. P. Sterbenz, E. K. Çetinkaya, M. A. Hameed, A. Jabbar, and J. P. Rohrer. Modelling and analysis of network resilience (invited paper). In *Proceedings of the Third IEEE International Conference on Communication Systems and Networks (COMSNETS)*, pages 1–10, Bangalore, India, January 2011.

[20] 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). *Springer Telecommunication Systems*, 2011. (accepted March 2011).

[21] J. P. G. Sterbenz, D. Medhi, B. Ramamurthy, C. Scoglio, D. Hutchison, B. Plattner, T. Anjali, A. Scott, C. Buffington, G. E. Monaco, D. Gruenbacher, R. McMullen, J. P. Rohrer, J. Sherrell, P. Angu, R. Cherukuri, H. Qian, and N. 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)*, pages 428–441, Berlin, Germany, May 18–20 2010.