G E N I

Global Environment for Network Innovations

Milestone 5
GENI Real-Time Measurements: Experimental Use-Case

Document ID: GENI-MS5-ERM-Sept09

September 01, 2009

Prepared by:
C.P. Lai, F. Fidler, M. Wang, and K. Bergman
Dept. of Electrical Engineering, Columbia University New York,
500 W. 120th Street,
New York City, NY 10027
under Project Nr. 1631
“Embedding real-time measurements for cross-layer communications”

TABLE OF CONTENTS

1 Document Scope

This section describes this document's purpose, its context within the overall GENI project, the set of related documents, and this document's revision history.

1.1 Executive Summary

This technical note presents the results obtained in work package, "Milestone 5: Support the GPO in developing an experimental use-case," of Project Nr. 1631, "Embedding real-time substrate measurements for cross-layer communications."

This milestone deals with ERM supporting the GPO in developing experimental designs for use-cases based on previous work on a measurement-driven cross-layer communication system. Building on the successful simulation work with the modules developed in [erm09_1], we validate our cross-layer communication schemes in developing an experimental use-case based on proactive packet protection.

In Section 2, we provide a summary on the proposed concept of a unified measurement framework (UMF) for GENI as discussed in our previous milestones [erm09_2, erm09_3]; we also give an overview of previously proposed possible implementation of the UMF in terms of network management protocols and hardware [erm09_4]. Section 3 highlights an experimental use-case that has been designed and explored, describing simulation and network test-bed experiments that show the advantages of our cross-layer schemes.

1.2 Related Documents

The following documents are related to this document, and provide background information, requirements, etc., that are important for this document.

1.2.1 GENI Documents

Document ID	Document Title and Issue Date
GENI_QR_ERM_Apr09	1Q09 Status Report
GENI_QR_ERM_Jul09	2Q09 Status Report
GENI-INF-PRO-S1-CAT-01.3	GENI Infrastructure Substrate Catalogue
GENI-MS1-ERM-March09-v1.1	Technical Note 1, Milestone 1
GENI-MS2-ERM-March09-v1.0	Technical Note 2, Milestone 2
GENI_MS3_ERM_March09_v1-0	Technical Note 3, Milestone 3
GENI_MS4_ERM_June09_v1-0	Technical Note 4, Milestone 4

1.3 Document Revision History

The following table provides the revision history for this document, summarizing the date at which it was revised, who revised it, and a brief summary of the changes. This list is maintained in chronological order so the earliest version comes first in the list.

Revision	Date	Revised By	Summary of Changes
1.0	01 Sept 09	C.P. Lai	Initial draft

2 Introduction

To summarize the major activities accomplished in previous milestones: in order to evaluate the GENI requirements for real-time user access to measurement data, we have earlier assessed the capabilities of GENI's future infrastructure with respect to real-time measurements [erm09_2]. As discussed in [erm09_3], the issue of interfacing available performance monitors (PMONs) to the GENI control framework and to the access point of the GENI researcher in a straightforward manner may lead to the following obstacles:

- Depending on the number of available PMONs, the number of required interfaces might become very large.
- Designing the interfaces requires extremely detailed, vendor-specific knowledge about the control mechanisms of the PMONs and about how measurement data is exported. Therefore, a level of abstraction is highly desired.
- A certain amount of editing and preparation of the measured data prior to the delivery to the researcher could be desirable.
- Extensibility of the measurement framework and manageability is difficult if each PMON has to be controlled individually.

Thus, rather than interfacing each performance monitoring device within the substrate directly with the control framework and the researchers, we recommended the design and use of a unified measurement framework (UMF) [erm09_3]. The UMF represents a universal, integrated measurement platform which can be accessed by the control framework and the experimenter via a limited number of well-defined interfaces. The main tasks and functionalities of the UMF are:

- to acquire measurement data from the various performance monitoring devices within the network substrate,
- abstraction of measurement capabilities and equipment from several manufacturers and vendors,
- provide a single point of access,
- basic processing of the measurement data, e.g. to extrapolate signaling data for the researcher or topology information for the control framework
- provisioning of some storage capacity for non-time-sensitive measurements,
- interfacing to the researcher via a unified measurement interface which allows the requesting and controlling of certain measurements, as well as delivers the measured data, and
- interfacing with the control framework so that the measurement framework (or a subset of PMONs) can be allocated to a requesting GENI researcher, and reconfiguration of the slice is made possible.

The UMF presents a uniform, integrated view and an abstraction of the measurement capabilities within the network substrate and makes them accessible/sliceable to the control framework.

The hardware implementation aspects of the UMF were discussed in [erm09_3] and software architectures to enable the efficient exchange of information to a UMF were investigated in [erm09_4]. Several data exchange formats and networking protocol languages were evaluated with respect to their applicability within the proposed unified measurement framework. Ultimately, we suggest a specific manifestation that may achieve the proposed UMF (Figure 2-1). Standard network management protocols such as SNMP, TL1, and SCPI could be used to access individual performance monitoring devices and to transmit measurement information to a hardware implementation of the UMF. For example, this may consist of one or more NetFPGA cards hosted in a server as suggested in [erm09_3]. The NetFPGA pre-processes the measurement information so that measurement data can be stored in databases (e.g. SQL), and accessed or exported to other services/software frameworks (e.g. SILO [silo09_1], perfSONAR [perfsonar09_1], etc.) via XML based protocols. From the UMF, information about the resources, network topology, and measurement capabilities should be sent to the GENI control frameworks by means of XML based/encoded data structures such as NDL or RSpecs.

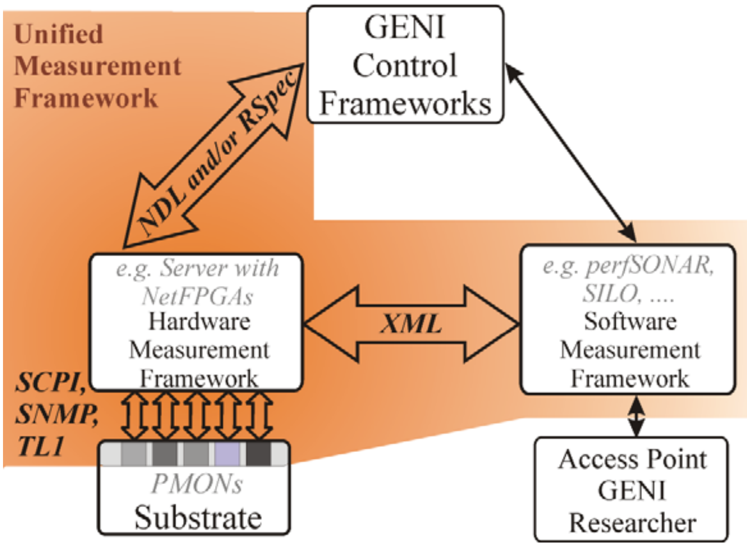

Figure 2-1: Example of UMF implementation concept showing hardware and software measurement frameworks, network management protocols, and exchange data formats and structures.

3 Experimental Use-Case

3.1 Overview

Within the scope of this milestone, the designed experimental use-case is based on the concept that detailed information about the physical layer infrastructure may be leveraged in corresponding cross-layer based protection and routing protocols, where information about physically disjoint paths is required [erm08_1]. The cross-layer scheme allows for introspective access to the physical layer to extract PMON measurement data, which can then be used to achieve significant network performance gains. The bidirectional information exchange creates a gateway for the PMON measurement data to be transmitted to the UMF, then to the researcher and/or to GENI software measurement frameworks in order to reconfigure and optimize packet routing.

Within the scope of accessing real-time measurements for a cross-layer network communications experiment, the proposed experimental use-case is based on a proactive packet protection switching mechanism [gerstel08_1]. In this quality-of-service (QoS) aware scheme, a degradation of packets can be proactively detected at the network receiving end using the embedded PMON measurements. The transmitting node can then allow the data stream to be switched and rerouted to an alternate, disjoint protection path to ensure no packet loss [ecoc09_1, ecoc09_2]. Our vision is for the network's packet routing decisions to be optimized dynamically with respect to varying QoS requirements and real-time substrate-embedded measurements on a packet-by-packet basis.

As the following subsections describe, the proactive packet protection mechanism allows for lower packet loss rates and overall enhanced network performance. This devised experimental use-case leverages the physical layer performance monitoring measurement capabilities embedded in GENI-related network test-beds and networking equipment, such as ORCA/BEN's Polatis fiber switch (which monitors optical power), ORCA/BEN's Infinera DTN ROADM (which monitors bit-error rates), and MAX's Adva transponder (which monitors bit-error rates), among others [erm09_2]. In this way, the proposed UMF will capitalize on the embedded network monitoring functionalities within the GENI substrate.

The proactive packet protection mechanism comprises one potential experimental use-case that highlights the advantages of cross-layer communications schemes and is a first step in the direction of enabling real-time measurement based cross-layer experimentation within GENI. It leverages PMON measurements, for example, bit-error rate (BERs) data from the Infinera DTN. As soon as the received BER exceeds a certain BER_E, the correction threshold of the underlying network's forward-error correction (FEC), a fast-reroute mechanism switches the data stream to a protection path. Our proposed proactive protection schemes starts earlier (at a lower pre-defined threshold $BER_T<BER_E$), with the ultimate goal of near hitless protection for sufficiently slow impairment dynamics.

The validity of the cross-layer experimental use-case (i.e. the packet protection scheme) was confirmed in a network simulation environment using the ns-2 modules developed in a previous milestone [erm09_1] and will be presented at an upcoming conference [ecoc09_1]. The cross-layer scheme was also confirmed experimentally on an in-house network test-bed; this work will also be presented at an upcoming conference [ecoc09_2].

3.2 Use-case validation: ns-2 simulations

In [erm09_1], we described the discrete-event network simulations that were developed within the scope of enabling real-time measurements in future network infrastructures such as GENI. The simulations were developed in the ns-2 open source network modeling environment, which had previously lacked the ability to implement a realistic physical layer model. The implementation of such a model in a unified manner is complicated due to the variety of possible impairments. The open source nature of ns-2 provided us the ability to integrate additional new modules based on C++. Thus, we implemented the following functionalities: intra-packet BER variations, a FEC module, and a local control plane [erm09_1]. The modules can be downloaded from the ERM GENI Wiki [geni09_1] and the following results will be presented in [ecoc09_1].

Based on these aforementioned modules, we are able to incorporate general physical layer BER variations with ns-2, enabling cross-layer simulations beyond existing modules specific to the wireless channel. This functionality allows us to explore a cross-layer experimental use-case that shows the advantages of cross-layer schemes within a network infrastructure. In [ecoc09_1], we compare fast packet protection mechanisms over either a packet-switched or a circuit-switched core, using GbE clients with 1500 bytes maximum transfer unit (MTU) and a 100-Gb/s line infrastructure as an example. As discussed above, the packet protection scheme sets a lower BER threshold for packet rerouting on an alternate path, as compared to a fast-reroute mechanism.

Within ns-2, we explore a step-like increase of the BER (Figure 3-1) and simulate the number of lost packets n in the network as a function of the slope of this BER step. Proactive protection offers zero packet loss until the BER step becomes so steep that the time span t_P between BER_E and BER_T is shorter than the round-trip time (RTT), the minimum time required for the protection mechanism to kick in. For fast transitions, the number of lost packets using proactive protection converges to n=RTT/t_S, with t_S being the TDM slot time.

We subsequently study sinusoidal log-BER variations with period Δt. We found that proactive protection shows no packet loss for t_P > RTT. Assuming packet durations much smaller than RTT, the two protection mechanisms perform identically for $t_P + t_E$ = RTT, with t_E being the time duration where packet loss occurs (red area in Figure 3-1). Beyond this, $t_P + t_E \leq$ RTT, the RTT is large compared to the impairment dynamics, and proactive protection offers no advantage over fast-reroute. For both schemes, the number of lost packets decreases with increasing BER dynamics. Additionally, Figure 3-1 compares the cross-layer proactive protection and fast-reroute schemes for several combinations of RTT and BER variation speeds $1/\Delta t$. The green area under the curve $t_P + t_E$ = RTT denotes the operating region where proactive protection outperforms fast-reroute. As expected, fast impairment dynamics require short RTTs for proactive protection to provide an advantage over fast-reroute. For typical optical transport networks (4 ms < RTT < 40 ms), proactive protection is effective against quasi-static impairments, but is likely to fail for dynamic impairments such as fast amplifier power transients.

Thus, our simulation exploration confirms that for a specific size of the network and physical impairment time scales, we have a reduced (and thus improved) packet loss rate with the proposed cross-layer packet protection scheme as compared to existing FEC-based fast-rerouting methods. In this way, we implement our experimental use-case of proactive packet protection and show that this cross-layer communications system may be advantageous in the future GENI network infrastructure.

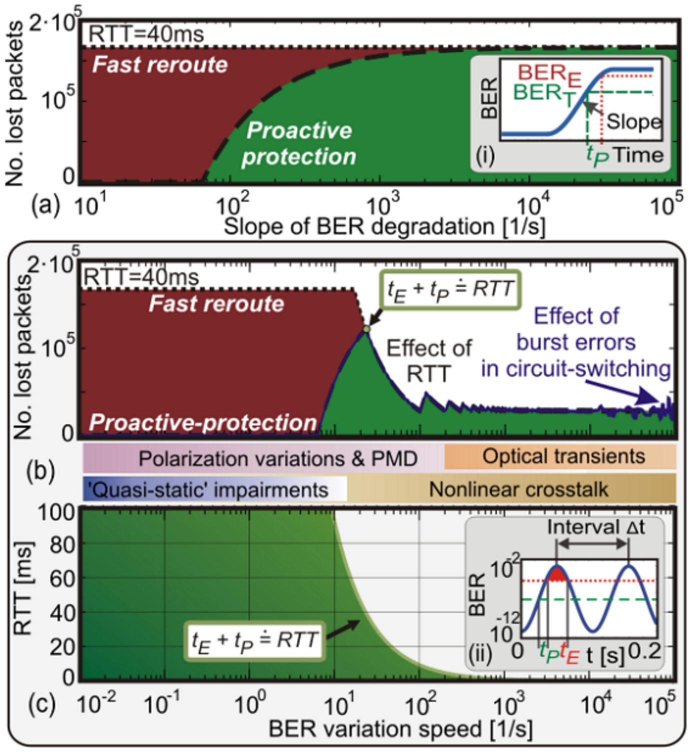

Figure 3-1: Number of lost packets (dotted: fast-reroute, dashed: proactive protection) vs. (a) slope of BER step, and (b) BER variation speed $1/\Delta t$; (c) Region where proactive protection outperforms fast-reroute [ecoc09_1].

3.3 Use-case validation: network test-bed experiments

Furthermore, to fully support the cross-layer ns-2 simulations, we demonstrated the cross-layer experimental use-case with experimental proactive protection on a small, in-house network test-bed. The optical packet switching fabric test-bed provides us with the flexibility to implement and test various networking functionalities in order to observe the realistic nature of our approach. In [ecoc09_2], we report on this experimental demonstration of the cross-layer packet protection scheme, whereby a degradation of high-priority packets is proactively detected at the receiving end using the read-out from a PMON. This allows the packet routing decision to be optimized with respect to varying QoS restrictions and signal quality degradation (e.g. BER). Experimentally, the proactive protection scheme uses a customized receiver in which optical packets are monitored and, depending on their QoS requirement (high/low priority) and BER, rerouted on a packet granularity.

In this experimental realization of the proactive packet protection scheme, contending messages within the network test-bed are dropped. Depending on their QoS class, packets with high BER are intentionally discarded after reception and rerouting is triggered. The mechanism leverages a potential signal introspection as provided by a 10-Gb/s FEC. These measurements can feedback to higher network layers to provide packet rerouting along an alternate, disjoint network protection path. Proactive protection detects a degrading BER and uses a predefined BER threshold above which packet rerouting is triggered to avoid packet loss by FEC. The loss of a degraded message is mitigated by a cross-layer control signal and subsequent protection path transmission on a parallel network route.

The system uses a modified receiver where optical packets may be monitored, proactively discarded if the signal is degraded, or forwarded to the destination port. The proactive switching mechanism is triggered on per-packet QoS and BER metrics. Data streams with high-priority/high-BER optical messages are rerouted on an alternate protection path, while low-priority (regardless of BER) and high-priority/low-BER messages are forwarded. Low BER denotes a signal quality below the predefined threshold for proactive protection switching, and high BER indicates a quality above the predefined limit.

A block diagram representation of the fabric test-bed with the cross-layer scheme can be seen in Figure 3-2, and the experiment details can be found in [ecoc09_2]. This work comprises an experimental demonstration of the cross-layer experimental use-case, whereby a proactive packet protection mechanism can be implemented in an in-house network test-bed. This helps to validate that advanced measurement-driven QoS-aware cross-layer network control can be achieved in the GENI infrastructure, within other GENI-related physical layer network substrates.

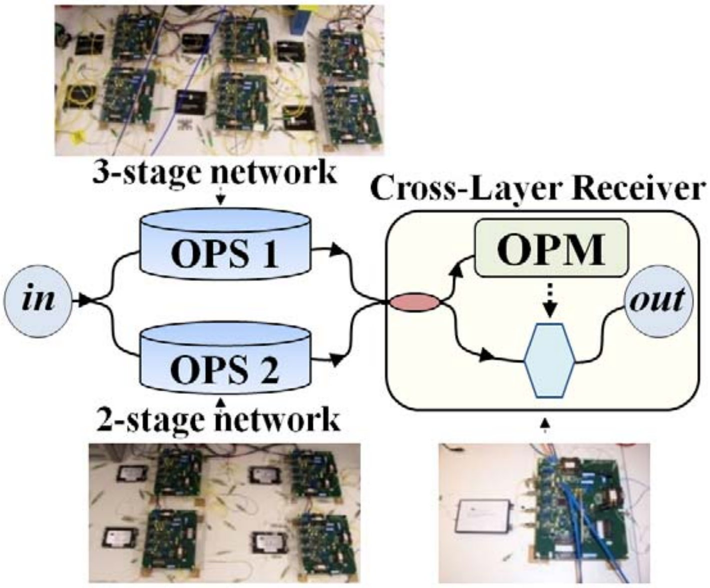

Figure 3-2: Block diagram of network test-bed with cross-layer scheme [ecoc09_2].

4 Summary and Conclusions

In this report, we first give an overview of our project findings, and subsequently discuss an experimental use-case that incorporates our vision of enabling embedded real-time measurements. The cross-layer experimental use-case that we designed implements a proactive packet protection mechanism that leverages the performance monitoring capabilities in the physical layer GENI substrate, such as the bit-error count/monitoring functionality embedded in the ORCA/BEN Infinera DTN ROADM. The proactive scheme detects a degraded packet stream with a lower BER threshold as compared to the one set by FEC to allow for near hitless protection.

We first discuss the details of the protection scheme and the feasibility of implementing this cross-layer signaling system within GENI. Building on previously designed modules, we then show simulation results in the open source simulation environment ns-2 that show that the proactive protection yields lower packet loss rates than existing FEC-based fast-rerouting schemes. Finally, we provide an experimental demonstration of the experimental use-case on our in-house network test-bed.

We endeavor for this report to provide the GPO with an experimental use-case that can be implemented in the GENI infrastructure given the availability of a measurement framework (i.e. the UMF) in combination with an appropriate software framework which allows for cross-layer information exchange (e.g. SILO). Current collaborations with the ORCA cluster provide a potential means for implementing this experimental use-case within the BEN substrate.

5 Bibliography

[1] [erm09_1] F. Fidler, C.P. Lai, and K. Bergman, "Technical Note 3, Discrete-Event Network Simulations (Project Nr. 1631, Milestone 2)," (2009, June) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements
[2] [erm09_2] F. Fidler, C.P. Lai, and K. Bergman, "Technical Note 1: GENI Requirements for Real-Time Measurements (Project Nr. 1631, Milestone 1)," (2009, February) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements
[3] [erm09_3] C.P. Lai, F. Fidler, and K. Bergman, "Technical Note 2: Specifications and Networking Protocols (Project Nr. 1631, Milestone 2)," (2009, February) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements
[4] [erm09_4] F. Fidler, C.P. Lai, and K. Bergman, "Technical Note 4: GENI Real-Time Measurements Software Architecture (Project Nr. 1631, Milestone 4)," (2009, February) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements
[5] [silo09_1] Net-Silos Team, "SILO Project – Services, Integration, controL and Optimization for the Future Internet (2009, May) [Online]. Available: http://www.net-silos.net
[6] [perfsonar09_1] perfSONAR (2009, May) [Online]. Available: http://www.perfsonar.net
[7] [erm08_1] K. Bergman, "Embedding Real-Time Substrate Measurements for Cross-Layer Communication", 3rd GENI Engineering Conference (2008, October) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements
[8] [gerstel08_1] O. Gerstel, I. Leung, G. Nicholl, H. Sohel, W. Wakim, and K. Kollenweber, "Near-Hitless Protection in IPoDWDM Networks, " Optical Fiber Communications (OFC) 2008, NWD4
[9] [ecoc09_1] F. Fidler, P. Winzer, C. P. Lai, M. K. Thottan, K. Bergman, "Cross-Layer Simulations of Fast Packet Protection Mechanisms" accepted at European Conference on Optical Communications (ECOC), Vienna, Austria (to be presented September 2009)
[10] [ecoc09_2] C. P. Lai, F. Fidler, K. Bergman, "Experimental Demonstration of QoS-Aware Cross-Layer Packet Protection Switching" accepted at European Conference on Optical Communications (ECOC), Vienna, Austria (to be presented September 2009)
[11] [geni09_1] Global Environment for Network Innovations Wiki, "Embedded Real-Time Measurements (Project Nr. 1631)" (2009, August) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements