G E N I

Global Environment for Network Innovations

Milestone 2
Specifications and Networking Protocols

Document ID: GENI-MS2-ERM-March09-v1.0

February 27, 2009

Prepared by:
C.P. Lai, F. Fidler, and K. Bergman
Dept. of Electrical Engineering, Columbia University New York,
500 W. 120th Street,
New York, 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 2: Specifications and networking protocols" of Project Nr. 1631, "Embedding real-time substrate measurements for cross-layer communications."

The purpose of this milestone is to develop a set of specifications for enabling real-time measurements within the substrate and specifications for networking protocols based on the GENI requirements for real-time user-accessed cross-layer measurements. In Section 2, we identify a set of specifications for the implementation of a unified, integrated measurement framework with the ultimate purpose of limiting the hardware and software overhead and complexity associated with accessing measurement data. Section 3 proposes recommendations for the implementation of networking protocols dealing with timing requirements, interfacing issues, and the separation between control and data signals.

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_Dec08	4Q08 Status Report
GENI_MS1_ERM_March09	Technical Note ERM: Milestone 1
GENI_MS3_ERM_March09	Technical Note ERM: Milestone 3

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	27 February 09	C.P. Lai	Initial draft

2 Specifications for the Unified Measurement Framework

2.1 Motivation

Within the scope of this milestone, we endeavor to present a set of specifications for networking protocols based on the GENI requirements for real-time user-accessed cross-layer measurements [geni08_1]. In order to create a seamless and transparent interface between the substrate, the control framework, and the GENI experimenter, our project's Milestone 1 [erm_MS1] recommends the implementation of a unified measurement framework (UMF). The UMF represents a universal, integrated measurement platform that is comprised of a number of well-defined interfaces. The platform will mitigate the management of multiple PMONs by achieving a level of abstraction desirable when handling multiple vendor-designed devices on the substrate level. Here, we define a set of specifications for the implementation of the UMF. Ultimately, the goal is to limit the hardware and software overhead when upgrading the substrate with additional PMONs and to reduce the complexity associated with accessing measurement data such that cross-layer based performance enhancements may be realizable.

2.2 Specifications

In order for the GENI researchers to have access to real-time measurements in an efficient and seamless manner, we have determined that an integrated measurement framework to acquire the measurement data must have the following specifications:

Timing constraints: A maximum allowed timing must be associated with each (control and measurement) information exchange between the substrate, the GENI control framework, and the experimenter. Depending on which performance measurement parameter is requested, between which entities information has to be exchanged, and the purpose of the communication, the signal exchange must be regulated within a certain reasonable time range. Well-defined timing requirements are especially important if decisions based on the measured parameters are made within cross-layer based networking protocols. Since the exchange of control information occurs either before the experiment is started (e.g. between the GENI control framework and the UMF for slicing measurement resources) or relatively seldom during the course of an experiment (e.g. changing the measurement update rate between the experimenter and the UMF), the timing for these control signals is less important. Consequently, we identify the following timing parameters as the most critical (cf. Figure 2-1):

- Measurement update rate: For example, in a network which uses SONET framing, measurements (e.g. bit error ratio (BER) deduced from forward error correction (FEC) devices) should be updated at least every 125 μs, which corresponds to the duration of a SONET frame [ch95_1]. For other framing methods (e.g. OTN) or packet networks (e.g. IP over WDM), other timing constraints may be required. They may also depend on other parameters, such as the channel line rate in the case of OTN framing where the frame duration decreases with increasing data rate due to its constant frame size (cf. ITU-T recommendation G.709). Therefore, some control information must be exchanged between the substrate and the UMF to determine these timing limits.
- Maximum allowed "time of flight": The maximum allowed "time of flight" (ToF) of measurement information at the substrate delivered to the experimenter via the UMF will be determined by the GENI experiment itself (and thus has to be communicated from the researcher to the UMF). For example, if fast protection mechanisms should be tested in a cross-layer optimized way, i.e. relying on real-time measurements, the ToF is limited by the requirements placed by the used network protocols. For example, fast rerouting (FRR) is defined to have a 50-ms recovery time [raj07_1].

Storage capability: The UMF should exhibit a certain level of storage capability. After measurement data is gathered from the substrate level PMONs, it may be desirable to store some of the measurement data or data about the device, link, or network for a finite time before or while being forwarded to the GENI access point. In this way, the measurement platform may be able to minimize the number of measurement collection scans that must be generated within a given time period. The UMF may also be required to store access information from the control framework or control parameters from the experimenter. Section 3 will provide an overview how this storage capability may be realized.

Separation of control signals and measurement information: The UMF should be able to accurately recognize and separate control signals from measured data signals, which usually will be sent at different update rates. This can be realized by either having separate physical ports for the control and data signals, or having the control information encoded differently from the data signals (format, etc.).

Interfaces: The UMF will specifically require interfacing between three separate and distinct GENI components (cf. Figure 2-1) - the substrate, the control framework, and the researcher.

- *Substrate-from/to-UMF:* The UMF must be able to interface with the substrate and the substrate components, i.e. its PMONs, in a *bidirectional* and relatively *fast* manner. The substrate-to-UMF signals will allow for the acquisition of measurement data, while the UMF-to-substrate control signals will be used to configure the PMONs, for example communicate the required sampling rate, measurement threshold levels, data exchange format, and so on. When the substrate's physical measurement parameters change, the measurement signals should be brought up directly to the UMF depending on the researcher-defined update rate. Ideally, these interfaces with the substrate should be data format transparent: the substrate data format should not affect the way in which the UMF interfaces with the physical layer PMONs. It can be assumed, that the exchange data format for the measurements will be determined by the interfacing capabilities of the PMONs embedded within the substrate. Thus, a number of different drivers should be developed for every available PMON.
- *Control framework-to-UMF:* The UMF must be able to interface with the control framework to reserve and provide an experimental slice for the GENI researcher in a unidirectional manner. The slice reservation speed is not as time critical as compared to the substrate measurement collection speed. Along with this notion, the UMF should also be able to process multiple measurement data as required by the control framework. The UMF should implement an algorithm that incorporates all the measurement data it is required to gather/store in way that optimizes network performance and gives priority to more important slices. That is, if multiple measurement collections are needed by multiple researchers, the UMF should optimize the process to optimize overall performance.
- *Researcher-from/to-UMF:* The UMF must be able to interface with the GENI researcher and its network architecture framework in order to export the measurements and to deduce control signals. This interface must be *bidirectional* and *medium fast*, with its timing constraints depending on the measurement parameter, its purpose, and the abstraction level as defined by the experimenter. Furthermore, the researcher should be able to request precisely which data measurements should be monitored and stored. For example, if the researcher is primarily concerned with the physical layer's BER performance, the UMF should prioritize this measurement.

2.3 Implementation

This section discusses how we propose the integrated unified measurement framework to be implemented within the GENI infrastructure while incorporating the required specifications delineated above. The hardware options are as follows:

- The UMF may be realized directly on the substrate equipment. However, this option is not adequately flexible and may give rise to equipment or vendor specific restrictions going forward as more substrate devices and/or PMONs are added to the GENI node. Further, the ability to realize storage capabilities may be hindered in this implementation option.
- The UMF may also be realized in an external hardware unit incorporating for example (next-generation) NetFPGAs [netf08_1] for time-demanding measurements, Ethernet-, GPIB-, or serial ports for non-time-demanding measurements, and a server for storage.

One important consideration is that the measurement framework hardware should be *extendable* over time. For example, the first stage may be a computer with a NetFPGA, then subsequently a next-generation 10G-NetFPGA within a bigger server to store more data, etc. The software and drivers within the UMF must be compatible and reconfigurable when hardware is improved or expanded.

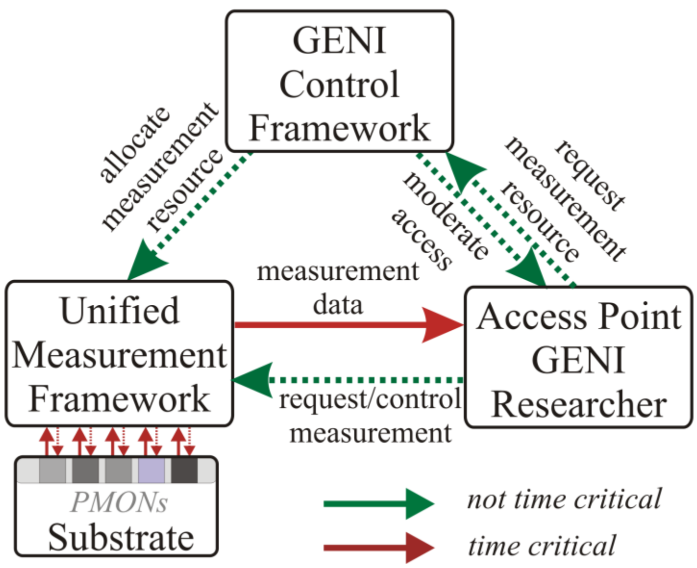

Figure 2-1: Schematic of interaction between performance monitors (PMONs) of substrate, unified measurement framework (UMF), control plane, and GENI researcher.

3 Networking Protocols

This section proposes recommendations for the networking protocols that would be required when implementing and using the UMF. By realizing an additional platform that interfaces between the substrate, control framework, and experimenter, we would need to develop a novel set of networking protocols to deal with timing requirements, potential interfacing issues, as well as the separation between the control and data signals. The signal format would also need to be established to ensure high performance of the UMF within the future GENI network infrastructure.

The specifications for these networking protocols and their implementations are as follows:

- The networking protocols would have to fulfill stringent timing requirements. The latency associated with the transmission of measurement data from the substrate to the UMF, and ultimately to inform the GENI researchers, will represent the limit for the experiment (cf. also Section 2.2).
- The interfaces between the three main infrastructure components (the substrate, the control framework, and the experimenter) should be seamless. Following Figure 2-1, the interaction between certain components should be bidirectional, allowing for transparent interfacing capabilities across the different network layers. This will readily support cross-layer communications in which the substrate devices dynamically interact with the UMF to enable routing and network reconfigurability.
- The supported signal formats should provide a clear separation between the control and data signals. This will further decrease the latency associated with each of the interfaces and allow the researcher to extract the measurement data from the substrate (via the UMF) in an abstracted fashion. It should be possible to seamlessly encode measurement data into cross-layer based protocols.

Regarding the control framework networking protocols, two alternatives for the involvement of the control framework are as follows:

- All the components that comprise the network infrastructure require direct regulation from the control framework (i.e. the control framework should always be aware of the PMON measurement data).
    - For example, the experimenter reserves one fiber ring and uses it. If its performance deteriorates, the UMF requests another fiber ring from the control framework. In this way, the control framework is involved in the experiment and is affected by the PMON data.

  We do not recommend this approach, as it would impose more stringent requirements on the timing constraints for control information exchange between GENI researcher and control framework as well as control framework and UMF.
- All the components that comprise the network infrastructure should interact dynamically without intrusion from the control framework (i.e. the control framework should not be aware of the PMON measurement data).
    - For example, the experimenter reserves two fiber rings in advance and uses one. If its performance deteriorates, the experiment slice automatically switches to the second reserved ring. In this way, the control framework is not involved in the experiment.

  While this method is certainly more resource demanding, it might adequately represent the only way of effectively implementing and testing cross-layer based networking protocols, where timing constraints are critical.

In a final UMF implementation, the experimenters might want to be able to set their preference on how involved they wish the control framework to be in their given experiment. If they are testing the cross-layer capabilities of the network, they may wish to be informed of the real-time measurements and PMONs' data in a fast and transparent way. If they are running a more high-level slice, they may then choose to have the real-time measurements hidden from view and timing constraints might be less critical. Switching between these alternatives can represent an essential functionality for the control framework protocol.

4 Summary

Within this milestone's work, we have shown the advantages of a unified measurement framework and how it would enable cross-layer communications research based on real-time measurements deduced from performance monitors within the substrate. The UMF is required if one wants to

- control and acquire measurement data in a unified way,
- abstract measurement capabilities and equipment,
- provide a single point of access for the GENI researcher and the control framework,
- enable basic processing of the measurement data,
- provide some storage capacity,
- allow for easy reconfiguration, and
- reduce software, hardware, and design overhead when extending the measurement capabilities of the GENI network.

We have developed a set of specifications that clearly shows the importance of taking into account the timing requirements, the storage capabilities, and the separation of data and control signals when designing the UMF and its interfaces. We proposed possible implementations and realizations of the measurement framework and gave recommendations for the design of (cross-layer based) networking protocols that make use of the measured data.

5 Bibliography

[1] [geni08_1] Global Environment for Network Innovations – Wikipedia (2008, December) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements

[2] [erm_MS1] F. Fidler, C. P. Lai, K. Bergman, "Milestone 1: GENI Requirements for Real-Time Measurements," 01 March 2009

[3] [netf08_1] J.W. Lockwood, "NetFPGA - A quick update for the GEC3 conference" (2008, October) [Online].Available: http://groups.geni.net/geni/attachment/wiki/presentations/NetFPGA_Status_Update_at_GEC3.ppt?format=raw

[4] [ch95_1] M. Ch. Chow, "Understanding SONET/SDH", Andan Publisher, NJ, ed. 1, 1995

[5] [raj07_1] A. Raj and O.C. Ibe, "A survey of IP and multiprotocol label switching fast reroute schemes", Computer Networks, Vol. 51, No. 8, 2007, pp. 1882-1907