Advanced Manufacturing Use Cases and Early Results in GENI Infrastructure

Alex Berryman, Prasad Calyam, J. Cecil, George B. Adams III, Douglas Comer

The Ohio State University, University of Missouri, Oklahoma State University, Purdue University, USA;
berryman@osc.edu; calyamp@missouri.edu; j.cecil@okstate.edu; gba@purdue.edu; comer@cs.purdue.edu

***Abstract*—Providing remote access and collaboration technologies to advanced manufacturing communities are exciting prospects due to the growth of the global marketplace and the pervasiveness of high-speed networks. There is a need to develop reliable protocols that extend beyond the current capabilities of typical TCP/IP connections that do not provide sufficient redundancy for controlling remote processes in manufacturing facilities. In addition, there is a need to suitably configure remote access protocol configurations that deliver satisfactory user experience amongst distributed collaborators synchronously working on manufacturing design workflows using cloud-hosted simulation software. In this paper, we present two case studies and early results that leverage the GENI Future Internet infrastructure for experimentation and development of new services that address such advanced manufacturing needs. Both case studies pivot around the idea of removing the need for users to have physical access to manufacturing resources and thus enable remote access to cloud-hosted services that use Future Internet capabilities for cost/time savings and improved convenience.**

I. Introduction

The globalization of advanced manufacturing has provided exciting prospects for using shared resources and fostering collaborations that require adoption of new enabling computing and networking technologies involving high-speed networks and cloud infrastructures. The prospects can enable novel web-architectures that foster organic, self-financed marketplaces featuring community-involved development of Apps for rapid innovation and crowd-sourced manufacturing commerce.

One adoption aim of these new technologies in advanced manufacturing communities is "remote process control", allowing manufacturers to save the money that is usually required to physically bring people to their local facilities to work on processes, and consequently using those savings to instead reinvest into their products. Also, the adoption of remote process control can create a new job market, where consultants can be hired on-demand because geographical barriers to access remote manufacturing resources can easily be overcome. However, to deliver the appropriate remote process control functions, there is a need to develop reliable protocols that extend beyond the current capabilities of typical TCP/IP connections that do not provide sufficient redundancy for controlling remote processes in manufacturing facilities.

This material is based upon work supported by the National Science Foundation under award number CMMI-1322748 and the Mozilla Foundation. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the Mozilla Foundation.

Another adoption aim is "collaborative advanced manufacturing", allowing manufacturers to utilize expertise from professionals around the world in order to efficiently produce and deliver a product in a competitive market, and for product-related workforce training/development. Using cloud-based simulation tools can enable experts at distributed sites to synchronously or asynchronously collaborate on advanced manufacturing life cycle activities. However, there is a need to suitably configure remote access protocols and simulation platforms that deliver satisfactory user experience amongst distributed collaborators over the Internet working on various manufacturing design workflows.

In this paper, we present two case studies and early results from Mozilla Ignite challenge projects [1] and [2] that leverage the NSF-supported Global Environment for Network Innovations (GENI) Future Internet infrastructure for experimentation and development of new services that could address the advanced manufacturing needs. Both case studies pivot around the idea of removing the need for users to have physical access to manufacturing resources and thus enable remote access to novel cloud-hosted services that use Future Internet capabilities for cost/time savings and improved convenience.

The first case study investigates the viability of using GENI as a testbed for the development of a reliable network protocol utilizing multiple software-defined network paths controlled via OpenFlow technology [3]. The second case study uses GENI resources to compare suitability of remote collaboration protocols that can facilitate distributed feasibility analysis of process alternatives in design of manufacturing parts. The GENI resources utilized include the ProtoGENI aggregate [4], TangoGENI testbed [5] and the OnTimeMeasure [6] instrumentation and measurement software.

The remainder of the paper is organized as follows: Section II describes the first case study in GENI involving remote process control using a reliable protocol, and Section III describes the second case study in GENI involving cloud computing for collaborative advanced manufacturing. Section IV concludes the paper.

II. Case Study I: Remote Process Control Using a Reliable Protocol

A. Overview

Consider an advanced manufacturing application using a reliable communication protocol to control the process taking place within the build volume of a remote 3-D printer. With a reliable connection, a part designer could do much more than

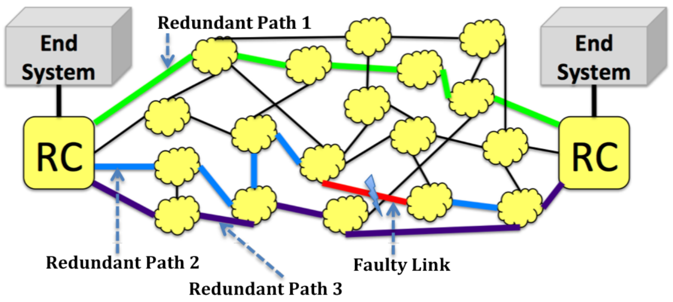

Fig. 1. Illustration of redundancy controllers (RCs) and redundant independent paths that comprise the reliable communication protocol serving two communicating end systems

simply download from his/her computer to the printer, the 3-D item shape and material(s) selection command sequence. Instead, the designer's computer could send printer commands to execute a customized fabrication path and sintering energy to use within the build volume, monitor the part as it is forming, compare those observations with the design, and apply a control algorithm to adjust future printer commands to keep the part within tolerances, achieving a feedback control loop across the Internet. The end goal could be to realize tighter control of process, materials, and tolerances, or it could be to reserve some portion of the printer command generation to a computer secured by the designer, thus concealing a portion of the designer's manufacturing knowledge from the 3-D printing location.

Extending and generalizing the above idea, one can imagine the designer establishing a connection to a third-party control algorithm service provider using a protocol that can deliver advanced process knowledge or superior computing capability or both and so on. The designer could also use above idea to remote control a robot with a camera on wheels and wifi for teleconferencing that can travel to certain parts of a manufacturing shop floor (that may include hazardous components), view what's going on and facilitate conversations between local and remote parties regarding quality control, process modification, and so on. Reliable communication allows the components of the process and the associated information in these scenarios to be located where desired, and allows dynamic control that improves user convenience and physical travel cost savings.

Our reliable communication protocol concept that can serve the above advanced manufacturing idea is illustrated in Figure 1. The communicating end systems are each connected to the Internet via a redundancy controller (RC). Using software-defined networking technologies such as OpenFlow [3], the sending RC establishes redundant paths (Redundant Paths 1 and 2 with green and blue colors), replicates packets, and sends data by both paths. The receiving RC notes the arrival of packets, forwards a copy of each packet to the end system, and matches replica packets. In matching replica packets, the receiving RC can perform on-going monitoring of path performance throughout the duration of the communication. Congestion or failure within a path (shown as a red link with lightning bolt within the blue path) is detected by the receiving RC as delayed or failed ability to match a packet with its replica. Delay or outright failure on one path does not prevent the delivery of packets to the end system, thus communication is reliable. Delayed reception of a replica packet can trigger the reliable communication protocol to build a new path (Redundant Path 3 with purple color) to restore path redundancy and maintain a desired level of reliability.

The ability for an end-user to dynamically create a redundant set of communication paths does not exist with the current TCP/IP protocol based Internet infrastructure. Hence, Future Internet infrastructures that support software-defined networking technologies such as OpenFlow can support a protocol to dynamically create and manage redundant paths between communicating devices. With replication of data packets and monitoring of path performance, Internet communication can move beyond "best effort" packet forwarding to reliable packet delivery. Such a capability is an ideal platform on which to build any number of applications in advanced manufacturing as well as other domains (e.g., health care) for which today's Internet is insufficient.

B. Testbed Setup

The goal of using the GENI infrastructure in this case study was to show that Openflow can be used to create redundant flows over multiple network paths in order to realize a reliable communication protocol.

To test the reliable communication protocol at scale, we setup a slice with two end-points and an interconnecting network with multiple exclusive paths connecting the end-points. The TangoGENI [5] project provides a suitable network architecture for our testbed by providing two distinct OpenFlow enabled paths between many reservable resources/slivers within aggregates such as ProtoGENI [4]. For the initial test setup, we reserved resources at the BBN site located in New York and at the Clemson University site in South Carolina.

Using VLAN 3715 in TangoGENI provided two separate paths between the hosts. Using the National Lambda Rail (NLR) resources, a path from the BBN to Clemson was constructed by using OpenFlow switches in Chicago and Atlanta. This route was the most direct path available. Using Internet2 resources, a longer path starting from BBN, crossing to Los Angeles, then returning to Houston, and finally Atlanta was used to reach the node in Clemson.

OnTimeMeasure Node Beacons were installed at the BBN and Clemson sites, respectively to collect on-going network status measurements in terms of metrics such as packet latency and loss. Both the redundant paths topology and related performance measurements reported by OnTimeMeasure are depicted on the web-based user interface shown in Figure 2. The path that the measurement traffic took was determined by the OpenFlow controller using the source port of each packet. Measurements originating from even numbered source ports took the shorter NLR path, while odd numbered source ports took the longer Internet2 path.

C. Results Discussion

Our testbed and experiments demonstrate the feasibility of using the TangoGENI testbed as a suitable infrastructure for

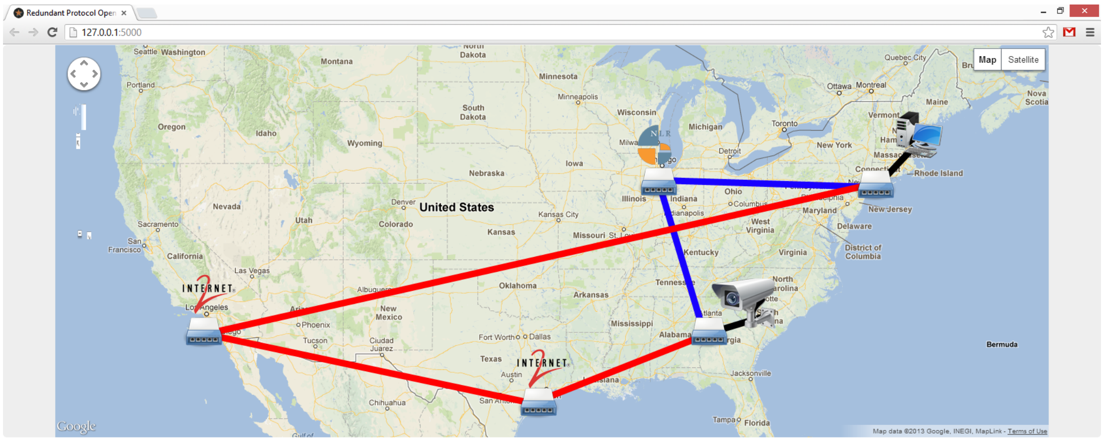

Redundant Protocol Path Information

Path	Status	Latency	Loss
NLR	up	22ms	0%
Internet2	up	17ms	0%

Fig. 2. Web-based user interface showing GENI testbed topology with redundant paths and related performance measurements

wide-scale testing of our reliable communication protocol. In order to implement our reliable communication protocol within an actual application scenario, the redundancy controllers belonging to two remote manufacturing related sites can be placed within the GENI slice, and the relevant code can be integrated with the current OpenFlow controller responsible for managing the redundant flows. Our experiments also showed the feasibility of extracting data from a web-service running on an OpenFlow controller such that the information can be presented to users through a simple web-based interface. In addition, they showed that data from the OpenFlow controller can be integrated with data from external instrumentation and measurement sources, such as those obtained from OnTimeMeasure in this case, in order to obtain performance intelligence that can be used to dynamically provision and adapt suitable application resources.

III. Case Study II: Cloud Computing for Collaborative Advanced Manufacturing

A. Overview

The potential benefits for allowing a user to collaborate with remote experts while assessing a part's design by utilizing virtual reality environments and other tools is attractive to both education and commercial manufacturing segments. Virtualization and cloud-based resources offer affordable as well as on-demand methods for students in K-12 and universities to learn about advanced manufacturing problems without the need of expensive local resources. Commercial partners can participate in virtual analysis of designs in an on-demand and scalable fashion using the elastic resource reservation capabilities of cloud platforms. The resulting designs can eventually be manufactured (or assembled), and the process can be monitored using the same cloud platforms. Thus, education communities, manufacturing and engineering enterprises will be able to collaborate globally within virtual environments and can rapidly bring advanced products to marketplaces.

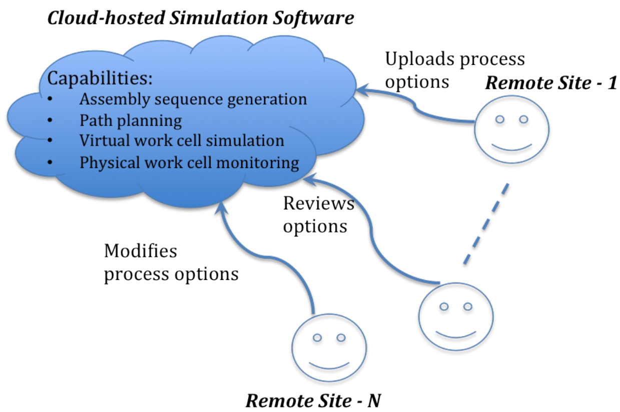

Fig. 3. Illustration of collaborative cyber-physical workflow for micro devices assembly using cloud-hosted simulation software

Our cloud computing based collaborative manufacturing concept is shown in Figure 3. Using our approach, an end user can upload a target design to a cloud-hosted virtual collaboration environment as part of a design/manufacturing activity. In our preliminary demonstration, we have considered assembly design examples from the domain of Micro Device Assembly (MDA) [7] [8]. We remark that MDA is an important emerging manufacturing domain that holds the potential of revolutionizing the manufacturing of micron sized devices in a range of domains from semiconductor manufacturing to chemlab-on-a-chip products. Using a cloud-based approach, users can upload a design which can then be accessed by other users to propose/modify assembly plans from any other location with Internet access. This enables virtual collaboration which can facilitate agile identification of process alternatives.

The collaboration involves a simulation application environment (created using C++ and Coin 3D libraries) which can be used to propose and modify candidate assembly alternatives.

Users from different locations can access (as shown in experiment described in Section III-C) the cloud based virtual assembly environment (and similar resources) and possibly collaborate on the feasibility of assembly alternatives.

In today's reality of traditional TCP/IP network stack based infrastructures, scalable and extensible collaborative manufacturing web-architectures that meet the cyber-physical environment needs of the advanced manufacturing communities are difficult to realize. Today's common infrastructure limitations include: (a) Isolation and Control: infrastructure slices (comprising of co-scheduled system and network resources) cannot be dynamically provisioned, programmed and scaled with changing user demands and application characteristics, and (b) Performance and Security: infrastructure slices cannot be optimized for performance (i.e., TCP friendliness may not be important when using UDP-based protocols such as PC-over-IP for remote desktop steering of the simulation application) and security (i.e., using application related slice certificates) when federating users and confidential design resources from multiple administration domains. We suppose that distributed elastic compute resources with layer-2 (Ethernet) networking across multi-domain networks that do not rely on traditional TCP/IP network stacks can overcome today's infrastructure limitations and can meet the bandwidth-intensive resource requirements for smooth simulation application interactions within agile MDA collaborations.

B. Testbed Setup

The goal of using the GENI infrastructure in this case study was to leverage the geographically distributed compute resources that can be dynamically configured, and use them to realize a collaborative analysis workflow whose performance can be compared with remote collaboration protocols (i.e., Opensource VNC and Microsoft RDP) over actual wide-area network conditions.

To test the remote collaboration protocols, we reserved a remote node in the Utah ProtoGENI [4] aggregate that had the Windows 7 platform provisioned on bare hardware. The node was used to host an MDA simulation application (created using C++ and Coin 3D libraries) that would be the target for remote clients interested in manufacturing planning and path planning issues virtually prior to physical assembly. Remote users with thin-clients connected to the MDA simulation application from the campus sites of Oklahoma State University, and The Ohio State University following the workflow shown in Figure 3. They were able to perform different user tasks that involve proposing and then modifying the assembly plans interactively for given part designs. OnTimeMeasure [6] initiated packet captures on the server-side were used to obtain performance measurements of instantaneous network bandwidth utilization of the thin-clients for the different user tasks shown in Figure 4. The 'simulation animation' in Figure 4 refers to the simulation occurring in the virtual environment related to the assembly activities while 'manual view' refers to the user performing navigation (zooming, turning around, etc.) in the virtual environment while exploring and studying the assembly in progress.

C. Results Discussion

The network bandwidth utilization trend of the VNC protocol under 'Auto' setting (which uses the best encoding for the available bandwidth between the client and server) indirectly shows that fluid animation of rotating and moving parts are not smoothly displayed. More specifically, simulation application details in the target environment are not displayed that are associated with changing the view as a result of navigation in the virtual environment intrinsic to the simulation of assembly activities. The periods where the VNC protocol fails to draw updates can be seen in Figure 4, where the VNC protocol (red dashed line) shows zero bytes were transmitted. For the same case, however, the RDP protocol (blue solid line) shows a high amount of network utilization, which indirectly shows the smooth fluid animation of rotating and moving parts. The visual impact of VNC's failure to update can be seen in the side-by-side screen capture of each protocol output shown in the left and right images of Figure 5 for RDP and VNC protocols, respectively. The visual impairments and choppy artifacts with the VNC protocol occur in the 'simulation animation' user task at time 37s through 39s, and most notably in the 'manual view movement' at the end of the experiment after 80 seconds. During these time periods, the RDP protocol smoothly displayed the simulation and view movements as noted by the MDA simulation user (also evident in the high RDP protocol network utilization).

Our experiments above guide the selection of the remote collaboration protocol configuration suitable to support the MDA related simulation activities involving distributed engineers using the GENI infrastructure. They demonstrate that legacy protocols such as VNC, although widely used for enterprise remote access tasks and within collaboration frameworks such as Remote Instrumentation Collaboration Environment (RICE) [9], are not always effective to support highly interactive design tasks over wide-area networks. The collaboration capabilities in frameworks such as RICE, when utilized with high-bandwidth remote access protocol configurations can include: (a) voice communications (i.e., VoIP), (b) instant messaging (i.e., chat) for communicating efficiently in-band during remote multi-user sessions that involve controlling/viewing the session (i.e., presence), and (c) management of control privilege amongst the users (i.e., control-lock passing) such that at any given instant, only one user controls a target simulation application (as in our MDA simulation demonstration).

IV. Conclusion

Our advanced manufacturing use cases and early results in the GENI infrastructure in this paper demonstrate the exciting potential to build Future Internet services that utilize access to high-speed networks and cloud infrastructures. They can bring together manufacturing and cloud architects, developers and consumers who can help each other, and rapidly improve the functionalities of products with cost/time savings and user convenience. We hope our preliminary efforts will inform and encourage the advanced manufacturing community working on agile web-architectures to utilize cloud infrastructures such as

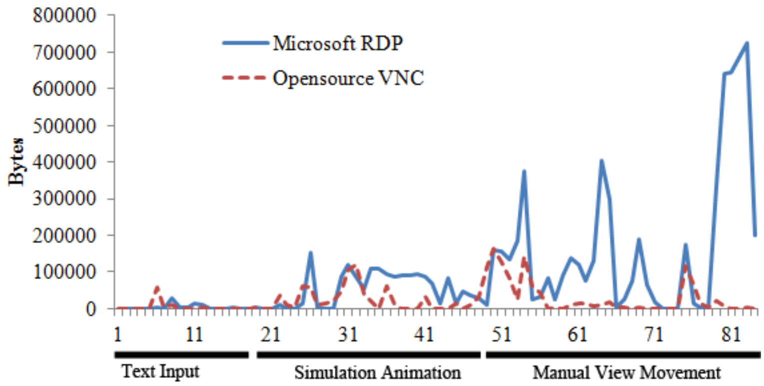

Fig. 4. Instantaneous network utilization performance comparison of RDP and VNC remote access protocols for user tasks

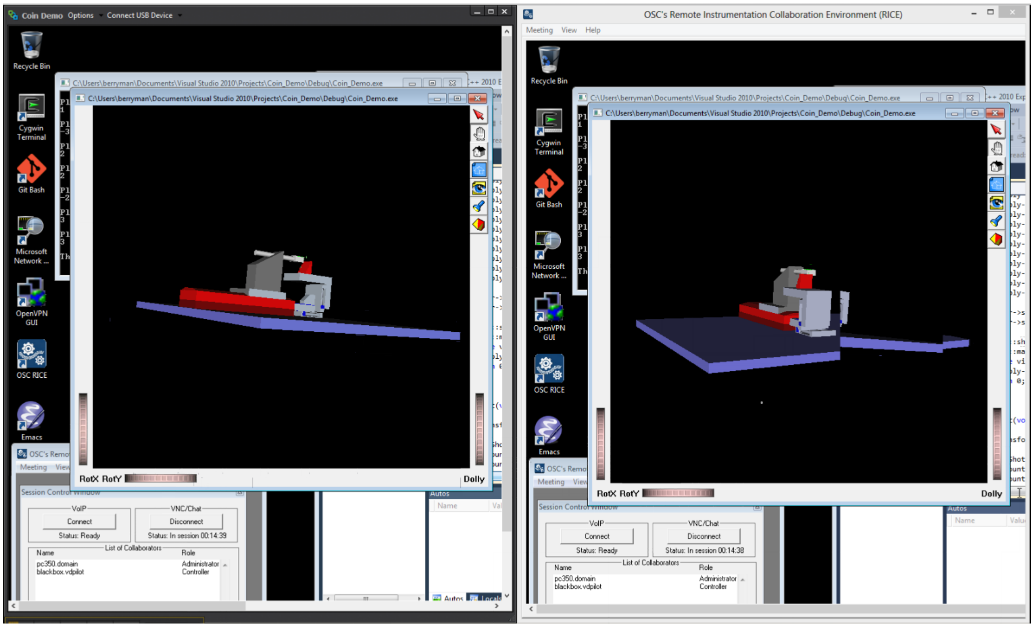

Fig. 5. Micro Assembly Simulation Output: Visually-smooth RDP rendering (left image) and visually-choppy VNC rendering (right image)

GENI, and latest protocol advances in federated networking, distributed computing and security.

REFERENCES

[1] G. B. Adams III, D. Comer, P. Calyam, A. Berryman, B. Sharma, R. Karandikar, D. Nguyen, J. Clark, P. Kennedy, J. Geske, "Remote Process Control Using a Reliable, Real-Time Protocol", *Mozilla Ignite Project* - https://mozillaignite.org/apps/418, 2012.

[2] J. Cecil, P. Calyam, A. Berryman, V. Rahneshin, "Cloud Computing for Collaborative Advanced Manufacturing", *Mozilla Ignite Project* - https://mozillaignite.org/apps/415, 2012.

[3] N. Mckeown, T. Anderson, H. Balakrishnan, et. al., "OpenFlow: Enabling Innovation in Campus Networks", *ACM SIGCOMM Computer Communication Review*, Vol. 38, No. 2, Pages 69-74, 2008.

[4] ProtoGENI Aggregate in GENI - http://www.protogeni.net, 2012.

[5] TangoGENI Operations - http://groups.geni.net/geni/wiki/TangoGENI, 2012.

[6] P. Calyam, M. Sridharan, Y. Xu, K. Zhu, A. Berryman, R. Patali, "Enabling Performance Intelligence for Application Adaptation in the Future Internet", *Journal of Communications and Networks (JCN)*, 2011.

[7] J. Cecil, N. Gobinath, "Development of an Integrated Framework for Micro Assembly", *Journal of Intelligent Manufacturing*, Vol.18, No.3, 2007.

[8] J. Cecil, "A Cyber Physical Framework for Micro Assembly", *Technical Report, Center for Information Centric Engineering, Oklahoma State University*, 2012.

[9] P. Calyam, A. Kalash, R. Gopalan, S. Gopalan, A. Krishnamurthy, "RICE: A Reliable and Efficient Remote Instrumentation Collaboration Environment", *Journal of Advances in Multimedia's special issue on Multimedia Immersive Technologies and Networking*, 2008.