A Trans-Pacific Programmable Network Testbed for Future Real-time Science Applications
Position Statement

Joaquin Chung*, Eun Sung Jung,^ and Rajkumar Kettimuthu*
Argonne National Laboratory, USA*
Hongik University, South Korea^

Rajkumar Kettimuthu will attend GEFI 2019

Motivation

A growing number of scientific fields require the ability to analyze data in near real-time, so that results from one experiment can guide selection of the next—or even influence the course of a single experiment. The experiments are often tightly scheduled, with timing driven by factors ranging from the physical processes involved in an experiment to the travel schedules of on-site researchers. With improvements in the sensor and detector technologies at experimental facilities (e.g., synchrotron light sources and neutron sources), data produced at these facilities significantly exceed their own local processing capabilities. Thus, the data needs to be moved to remote compute facilities both within and outside a country (or continent) as the users of these facilities often span diverse geographic locations. The computing and network resources must be available at a specific time, for a specific period. On-demand network bandwidth, though provided by backbone research and education networks such as ESnet and Internet2, is not easy to get end-to-end in an automated fashion. Even though compute resources can be obtained on-demand (at least in some institutions), those resources are not typically connected to the wide-area network (WAN). The typical model is that the data coming from the WAN goes into the parallel file system via the dedicated data transfer nodes (DTNs) and compute nodes access the data from the parallel file system. This model does not work well for near real-time analysis of the data streams coming from an experiment or simulation. We need international (and intercontinental) testbeds to evaluate solutions to enable these emerging science workflows.

Description of Research

In the past we have proposed and demonstrated an architecture for reserving and provisioning dedicated paths for scientific network flows. We called this architecture software-defined network science flows (SDN-SF). We also presented the design and implementation of an elastic data transfer infrastructure (DTI) that leverages elastic resources for large data transfers. Although SDN-SF can provision dedicated network paths end-to-end, it assumes software-defined networking (SDN) availability in all networks involved in the end-to-end path. Similarly, the elastic DTI requires a programmable network to connect elastic resources with WAN perimeter. We call this programmable network a dynamic Science DMZ. Researchers in South Korea have proposed solutions similar to SDN-SF and elastic DTI under the KREONET-S project.

Our ongoing research focuses on enabling near real-time data analysis for the elastic DTI and the dynamic Science DMZ. We are investigating techniques for enabling high accuracy network telemetry and data plane programmability in campus networks. We will consider both end-to-end availability of SDN and partial availability of programmable networks (e.g., only available at the end nodes as virtual swtches). For future work we will like to understand how science workflows that span geographically diverse locations affect real-time analysis of scientific data. We propose to optimize the resource usage and users' flow allocation in the network. We are also interested in developing resource allocation techniques for scenarios in which there is partial availability of the high accuracy network telemetry.

History and Plans for International Collaboration

Jung (currently at Hongik University in South Korea) and Kettimuthu (currently at Argonne National Laboratory, USA) have a long history of collaboration. Jung and Kettimuthu worked together on developing and optimizing wide area file transfer tools for 3 years (2013-2016) at Argonne. Jung moved to Korea in 2016 but Jung and Kettimuthu have continued their collaboration in the areas of high speed file transfers and SDN.

We plan to extend our collaboration and include institutions such as KISTI, KREONET in South Korea and ESnet, Univeristy of Chicago (especially the Chameleon cloud) in the United States to create a trans-Pacific programmable network testbed with high accuracy network telemetry. This testbed will help us understand how science workflows that span geographically diverse locations affect real-time analysis of scientific data. The first phase of our research will start with a knowledge transfer by comparing implementation experiences of SDN-SF, elastic DTI, and KREONET-S. In the second phase, we will evaluate a dynamic Science DMZ for real-time analysis between South Korea and the US with high accuracy network telemetry and programmable data planes. In the third phase we will study optimization of resource allocations. Experiment-based research is essential for this position statement because we are interested in studying these systems under real load and across geographically disperse locations.