SC21 Network Research Exhibition: Demonstration Abstract

Bottleneck and AI-Aware Traffic Engineering for Data Intensive Sciences using GradientGraph® and NetPredict Across the Pacific Research Platform and Global Network Advancement Group Multidomain Testbed

Reservoir Labs, New York

Sruthi Yellamraju, Noah Amsel, Brendan von Hove, Alison Ryan, James Ezick, Jordi Ros-Giralt, Richard Lethin

{yellamraju, amsel, vonhofe, ryan, ezick, giralt lethin}@reservoir.com

Pacific Research Platform, University of California, San Diego

John Graham, Dima Mishin, Thomas Hutton {jj graham, dmishin, hutton}@ucsd.edu

Lightwave Research Laboratory Columbia University, New York

Min Yee Teh, Zhenguo Wu, Keren Bergman {mt3126, zw2542, bergman}@columbia.edu

ESnet

Mariam Kiran, Tom Lehman, Xi Yang, John Macauley, Inder Monga (mkiran, lehman, xyang, macauley, imonga) @es.net

California Institute of Technology High Energy Physics and GNA-G

Harvey Newman, newman@hep.caltech.edu, Justas Balcas , jbalcas@caltech.edu, Raimondas Širvinskas raimis.sirvis@gmail.com

Yale Institute for Network Science, New Haven

Yuan Jiang, Aosong Feng, Leandros Tassiulas {yuan.jiang,aosong.feng leandros.tassiulas}@yale.edu

Abstract

In this Network Research Exhibition (NRE) demonstration, we will present key technologies and a new operational paradigm for next generation networks with intelligent AI-empowered control and data planes. The target applications include both the most challenging data intensive science programs such as the Large Hadron Collider (LHC) and Vera Rubin Observatory as well as many other data intensive applications using a wide area topology spanning local, regional, national and transoceanic distances.

The demonstration will showcase GradientGraph® (G2) Analytics and the NetPredict prediction tool. Developed by Reservoir Labs, G2 is a new network optimization platform that enables high-resolution traffic engineering and capacity planning for high-performance networks. G2 is based on a new mathematical theory [1,2] and algorithms that can efficiently construct the *bottleneck structure* of a network. NetPredict uses deep learning algorithms to predict high impact flows and congestion points leveraging graph neural networks [3]. In this demo, we (1) identify in real time the bottleneck network links, (2) make optimized traffic engineering decisions (e.g., flow re-routing or traffic shaping, while meeting SLA constraints), (3) create a baseline to identify flow performance issues, (4) perform capacity planning to optimize network upgrade decisions and (5) analyze network resilience, among other applications.

Our demonstration will run on the Pacific Research Platform (PRP), a research and education network (REN) connecting more than 50 institutions, national labs, and universities across the US and the world, as well as other wide area paths in the AutoGOLE/SENSE global testbed spanning the U.S., Europe, and the Asia Pacific and Latin American regions. Maintained by the University of California San Diego, PRP is a Kubernetes-based virtual network with about 200 nodes, 500 GPUs and 7000 cores.

The Software-defined network for End-to-end Networked Science at Exascale (SENSE) system will be used to dynamically provision Layer 2 and 3 virtual circuits with bandwidth guarantees and to provide mechanisms to integrate resources beyond-the-network, such as compute, storage, and Data Transfer Nodes (DTNs) into this automated provisioning environment. The persistent testbed is developed and deployed by the AutoGOLE/SENSE and Data Intensive Sciences working groups of the Global Network Advancement Group (GNA-G). Flow group identification, policy and state information, and real-time programmable traffic engineering

operations will also be enabled through the use of P4.

Goals

1. Integration of the GradientGraph™ (Reservoir) and NetPredict (ESnet) network optimization platforms into the network.
2. Generate traffic consisting of several long-lasting impactful flows; generate background traffic as a set of many smaller flows. Create congestion on one or more segments.
3. Placement of a new large data transfer flow f on the network. Use of the ‘kroutes’ G2 REST API endpoint to obtain the optimal path. G2 provides both the optimal path p* and the forecasted transmission rate r* of the flow. Use the NetPredict API to forecast the transmission rate, providing another validation point.
4. Placement of flow f according to path p*.
5. sFlow query to obtain the actual transmission rate r.
6. Comparison of r and r* for model validation.

Resources

The demo will be presented on the PRP platform, which we will connect to virtually. G2 and NetPredict will collect SNMP, NetFlow and sFlow data to return results to the optimizer, showing, in real-time, network updates.

We will link to (1) Network topology (switching and link connectivity graph, and link capacity information) and (2) flow information (which can be obtained from tools such as NetFlow, sFlow).

Involved Parties

- Reservoir Labs: Sruthi Yellamraju, Noah Amsel, Brendan von Hove, Jamie Nachbar, Alison Ryan, Dr. James Ezick, Dr. Jordi Ros-Giralt, Dr. Richard Lethin.
- Pacific Research Platform: John Graham, Dima Mishin, Thomas Hutton
- Caltech / GNA-G: Harvey Newman, Justas Balcas, Raimondas Širvinskas
- ESnet. Mariam Kiran, Xi Yang, Tom Lehman, John Macauley, Inder Monga
- Yale Institute for Network Science: Dr. Leandros Tassiulas, Yuan Jiang {leandros.tassiulas, yuang.jiang}@yale.edu
- Lightwave Research Laboratory Columbia University: Min Yee Teh, Zhenguo Wu, Dr. Keren Bergman

References

[1] J. Ros-Giralt, N. Amsel, S. Yellamraju, J. Ezick, R. Lethin, Y. Jiang, A. Feng, L. Tassiulas, Z. Wu, M. Yeh Teh, K. Bergman, "Designing Data Center Networks Using Bottleneck Structures," ACM SIGCOMM, August 2021.

[2] J. Ros-Giralt, A. Bohara, S. Yellamraju, H. Langston, R. Lethin, Y. Jiang, L. Tassiulas, J. Li, Y. Lin, Y. Tan, M. Veeraraghavan, "On the Bottleneck Structure of Congestion-Controlled Networks," ACM SIGMETRICS, Boston, June 2020.

[3] T Mallick, M Kiran, B Mohammed, P Balaprakash, Dynamic Graph Neural Network for Traffic Forecasting in Wide Area Networks, arXiv preprint arXiv:2008.12767, 2020 (IEEE Big Data 2020)