Validating Simulations of Large-scale Computer Networks

Kevin Mills

October 27, 2015

Problem Statement

- We pioneered methods to simulate and analyze macroscopic behavior in large networks, but…
- A key hurdle remains:
  - to demonstrate convincingly that our simulations exhibit behaviors consistent with real networks.
- This *hurdle exists generally within the scientific field of big network simulation*…

Background (I)

- DoE COMBINE workshop:
  - identifies need for verification and validation (V&V) techniques to increase confidence in simulations of big networks
  - identifies need for research to characterize uncertainty in big network simulations
  - *suggests recent advances in configurable, empirical network test beds might create an opportunity to improve the situation*

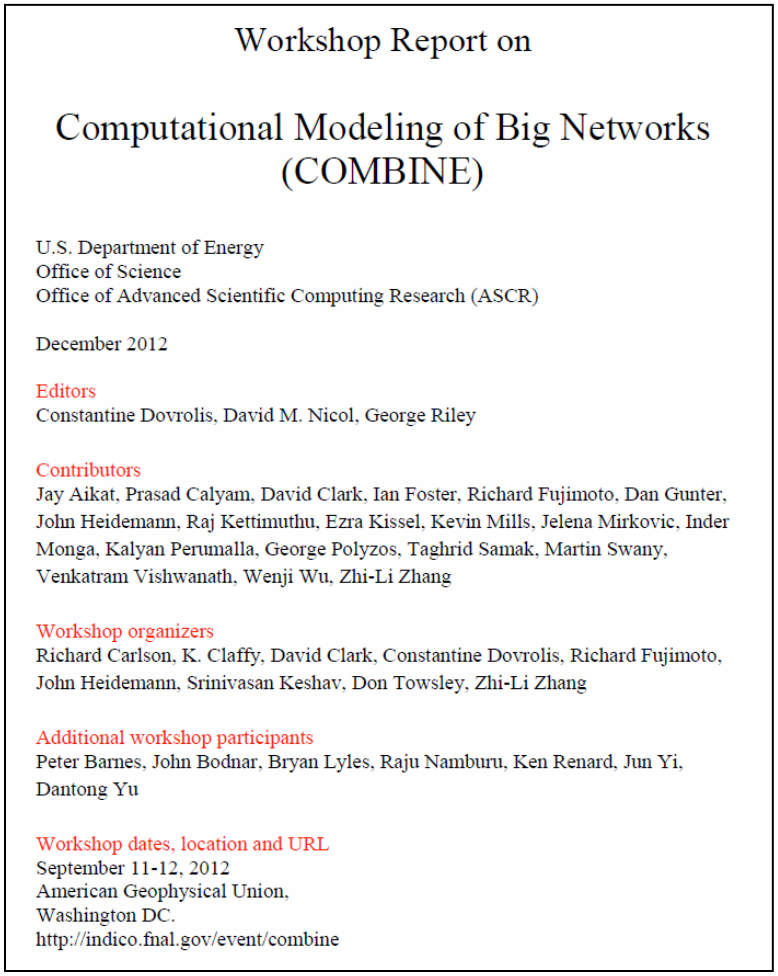

Workshop Report on

Computational Modeling of Big Networks (COMBINE)

U.S. Department of Energy
Office of Science
Office of Advanced Scientific Computing Research (ASCR)

December 2012

Editors
Constantine Dovrolis, David M. Nicol, George Riley

Contributors
Jay Aikat, Prasad Calyam, David Clark, Ian Foster, Richard Fujimoto, Dan Gunter, John Heidemann, Raj Kettimuthu, Ezra Kissel, Kevin Mills, Jelena Mirkovic, Inder Monga, Kalyan Perumalla, George Polyzos, Taghrid Samak, Martin Swany, Venkatram Vishwanath, Wenji Wu, Zhi-Li Zhang

Workshop organizers
Richard Carlson, K. Claffy, David Clark, Constantine Dovrolis, Richard Fujimoto, John Heidemann, Srinivasan Keshav, Don Towsley, Zhi-Li Zhang

Additional workshop participants
Peter Barnes, John Bodnar, Bryan Lyles, Raju Namburu, Ken Renard, Jun Yi, Dantong Yu

Workshop dates, location and URL
September 11-12, 2012
American Geophysical Union,
Washington DC.
http://indico.fnal.gov/event/combine

Background (II)

- Operational networks cannot be controlled so that simulation results can be duplicated there
- Parameters assigned in simulations might not reflect values from real networks
- Configurable network test beds could provide:
  1. reproducible platforms on which to verify simulation results
  2. foundations for making measurements that can form the basis for simulator parameterizations

Background (III)

- We had already developed the *MesoNet* simulator and had run experiments to study global patterns of congestion and user experience in large TCP/IP networks
- We had already installed a rudimentary in-house network emulation facility, based on Emulab technology developed at the University of Utah

Project Aims

1. Investigate and develop rigorous method to verify accuracy and characterize uncertainty of large-scale network simulations
2. Advance the state-of-the-art in V&V for large-scale network simulations
3. Demonstrate the findings from an earlier simulation study hold in an empirical setting

Technical Approach

1(a) Scale up & Refresh In-house Emulab

1(b) Scale down MesoNet Simulations

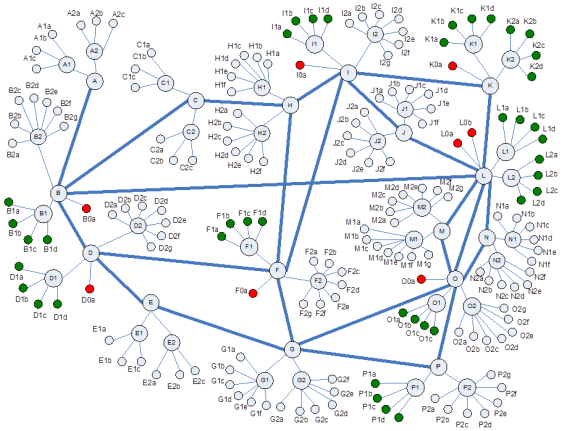

2 Replicate MesoNet Sensitivity Analysis on Emulab

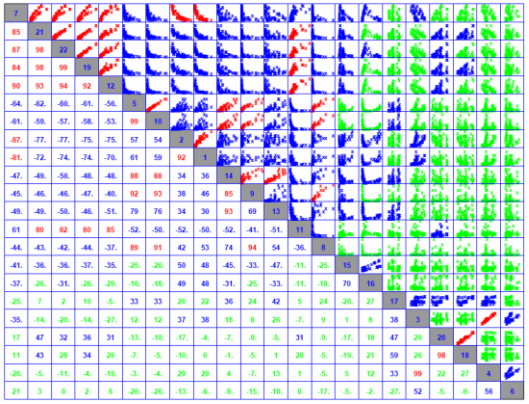

1(c) Recruit a New Employee

1(a) Progress on Emulab

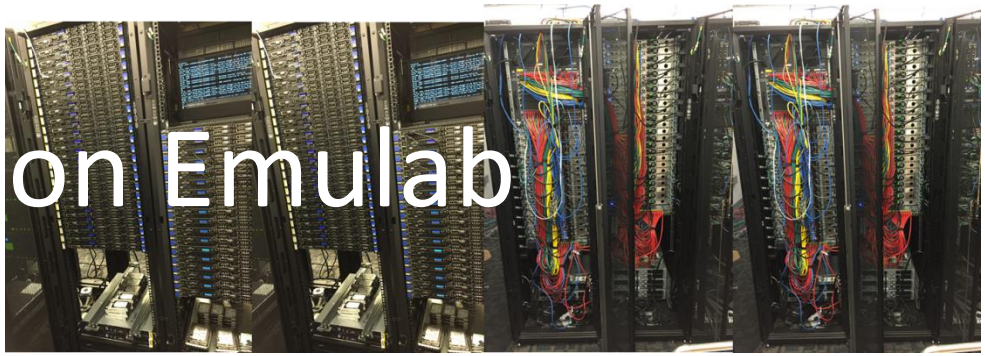

- 2x number of available nodes, with increased memory size and speed, 3x disk capacity and 2.3x more network connections
- 10x increase in backbone speed from 1 to 10 Gbps
- 2x number of switch ports
- Larger, faster control network
- Remote power cycling added
- Upgraded to most current version of Emulab control software

Previous 50-node In-house Emulab

Count	Server	Memory	Disk	NICs
12	860	2 GB	300 GB	6@1 Gbps
38	R610	12 GB	140 GB	6@1 Gbps

Upgraded 100-node In-house Emulab

Count	Server	Memory	Disk	NICs
56	R610	12 GB	1 TB	6@1 Gbps
24	R620	32 GB	1 TB	8@1 Gbps
20	R630	32 GB	1 TB	8@1 Gbps

http://emulab2.antd.nist.gov/

1(b) Progress on MesoNet

FIT ON EMULAB

4-Tier Topology from US ISP

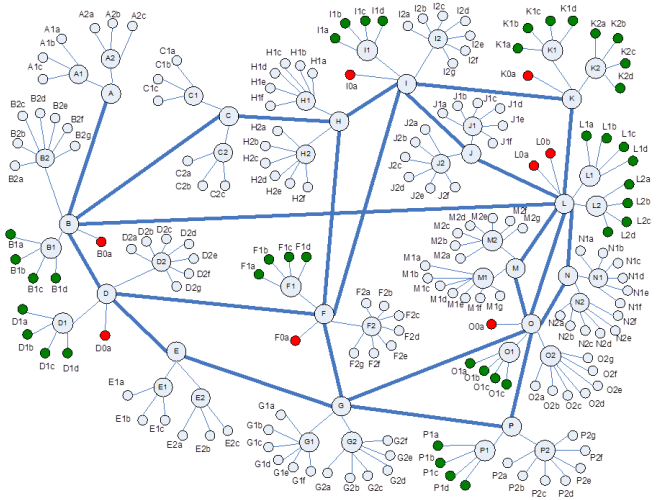

Tier	Node Type	Count
1	Backbone	16
2	PoP	32
3	Access	170
4	Source	51,588
4	Receiver	206,352
Total	All	258,158

40x SLOWER SPEED

38.4 Gbps Backbone

Parameter	Value	Speed Relationships		Speed Scaling with X3
*s*1	X3	**Router Class**	**Speed**	**X3 = 1600 p/ts**
*s*2	4	Backbone	*s*1 x *BBspeedup*	3200 p/ts
*s*3	10	PoP	*s*1/ *s*2	400 p/ts
BBspeedup	2	**N**-Class	*s*1/ *s*2/ *s*3	40 p/ts
Bfast	2	**F**-Class	*s*1/ *s*2/ *s*3 x *Bfast*	80 p/ts
Bdirect	10	**D**-Class	*s*1/ *s*2/ *s*3 x *Bdirect*	400 p/ts
Hbase	8	Source/Receiver	*Hbase*	8 p/ts
Hfast	80	Fast Src/Rcvr	*Hfast*	80 p/ts

960 Mbps Backbone

Parameter	Value	Speed Relationships		Speed Scaling with X3
*s*1	X3	**Router Class**	**Speed**	**X3 = 40 p/ts**
*s*2	4	Backbone	*s*1 x *BBspeedup*	80 p/ts
*s*3	10	PoP	*s*1/ *s*2	10 p/ts
BBspeedup	2	**N**-Class	*s*1/ *s*2/ *s*3	1 p/ts
Bfast	2	**F**-Class	*s*1/ *s*2/ *s*3 x *Bfast*	2 p/ts
Bdirect	10	**D**-Class	*s*1/ *s*2/ *s*3 x *Bdirect*	10 p/ts
Hbase	8/X3	Source/Receiver	*Hbase*	0.2 p/ts
Hfast	80/X3	Fast Src/Rcvr	*Hfast*	2 p/ts

Planned Next Steps

- 1(c) Complete recruiting of new hire (interviews conducted, candidate identified and has applied)
- Merge scaled-down MesoNet onto scaled-up Emulab to determine envelope of feasibility
  - How many physical Emulab nodes required for 258,158 virtual nodes?
  - Can Emulab support 40x scaled down MesoNet speeds?
- 2 Replicate MesoNet sensitivity analysis for 40x scaled down MesoNet speeds on Emulab and evaluate results
  - Same qualitative results?
  - Reasonable basis for V&V method?
  - Possible to characterize uncertainty of measured outputs?


Questions?