Research Statement

Methodology

My research mainly focuses on parallel discrete-event simulation, high-performance modeling and simulation of computer networks and computer systems. My research in general is a combination of modeling and systems. A significant portion of my research involves building, experimenting, optimizing, and analyzing computer systems and computer networks. As such, a typical research cycle includes identifying problems, devising solutions, developing models, implementing prototypes, conducting experiments, and analyzing results, which subsequently lead to new problem definitions, new insights, and new solutions. Compared with other more theoretical or data-driven approaches, system-oriented research usually takes a much longer duration due to development and prototyping before meaningful results can be obtained.

Most publications related to Computer Science appear in conferences and workshops, as they generally result in a broader and more immediate impact to the research community, compared with those published in journals. I balance my publications between journals and conferences/workshops. In total, I have published 2 book chapters, 15 journal articles, and 65 conference/workshop papers. Most of my publications appeared in the premiere journals and conferences in the modeling and simulation area. I have been able to publish good quality papers, as evident by the number of citations (2,821 in total, with an h-index of 27 and an i10-index of 48, according to Google Scholar from August 2018). I have a number of papers nominated for the best paper awards (2 best paper awards and 7 best paper candidates).

Research Areas

A common theme of my research is to develop efficient modeling and simulation techniques for studying computer systems and computer networks, in particular, large-scale computer networks and high-performance computing systems. General research topics include:

- Designing high-performance and effective simulation models of large-scale systems
- Developing parallel simulation tools to handle large-scale systems
- Applying data driven methods for capturing large complex system behaviors

My specific research areas are listed as follows:

A. Parallel Discrete-Event Simulation

Parallel discrete-event simulation (PDES) is a research area crosscutting between modeling simulation and parallel computing. By exploring the potential parallelism in simulation, PDES can overcome the limitations imposed by sequential simulations in both execution time and memory space, and can thus bring substantial benefit to time-critical simulations, simulations of large-scale systems, or both. My main contribution in this area has been developing efficient parallel synchronization algorithms and high-performance discrete-event simulation tools. I designed and implemented several parallel simulation synchronization protocols, including composite synchronization for both shared-memory multiprocessors and distributed memory machines, and a lock-free scheduling algorithm for shared-memory multiprocessors. I also developed several well-known parallel discrete-event simulators, including SSF and several of its variants.

B. Large-Scale Network Simulation and Emulation Testbeds

The ability to conduct high-fidelity network experiments and allow easy exploration of the design space is crucial for studying large-scale networks and their complex behaviors. My main contribution in this area is three-fold. First, I designed and implemented network simulators (including SSFNet and RINSE) to study large-scale networks (including Internet and wireless networks) as well as network protocols and distributed applications. I applied parallel discrete-event simulation methods to enable large-scale models. Second, I developed parallel real-time simulation techniques so that the simulated large-scale network can interact with real applications and physical network entities. My students and I developed several real-time network simulators, including PRIME and PrimoGENI. The challenge is to deliver real-time performance and simultaneously maintain scalable interaction between simulation and physical network. We designed real-time scheduling algorithms and developed efficient emulation infrastructures. Third, we extended real-time simulation to incorporate machine and network virtualization techniques to facilitate scalable network experiments on network infrastructures and testbeds (such as GENI).

C. Hybrid Network Modeling for Scale

Detailed packet-oriented network simulation can be computationally expensive for large-scale networks. Using hybrid network traffic models provides the necessary trade-off between accuracy and computational efficiency. My main contribution in this area is developing hybrid models to efficiently represent large-scale network phenomena. First, my students and I designed an integrated approach combining fluid-based traffic model using stochastic differential equations with a detailed packet-oriented simulation. Second, we extended the hybrid model and developed a mixed computation method, where the fluid traffic model is implemented on GPUs for representing background traffic on large-scale networks, and the packet-oriented network simulation is run on CPUs to capture detailed protocol behaviors. Third, we developed a novel symbiotic simulation method that combines network simulation and emulation effectively without having them to exchange network traffic. The symbiotic approach has also been extended to enable distributed network emulation.

D. Modeling and Simulation of High-Performance Computing (HPC)

To accommodate the increasing computational demand of scientific applications, HPC systems have been constantly changing with new architectures and programming models. Accurate, fast, and scalable performance models and simulation tools are essential for evaluating design decisions and predicting application performance. My main contribution in this area is in two aspects: one in simulation, and the other in modeling. In the first aspect, we designed and developed high-performance simulation tools, using parallel simulation implemented in just-in-time compiled languages. We developed HPC architecture models with detailed interconnection networks, application models with detailed communication behaviors, and job scheduling and resource provisioning models based on the architecture and application models. In the second aspect, we developed models to study energy consumption of the HPC systems. More specifically, we designed job scheduling and resource allocation methods, and also applied game theoretical approach such as contract theory to enable energy demand response of HPC systems.

E. Other Areas

I have a broad interest in algorithms, modeling, and data analysis, and applying them for design optimization and performance evaluation of both traditional and emerging networks and systems areas, including datacenter networks, software-defined networks, storage and I/O systems, cloud computing and virtualized systems.

Funded Research

My research has been supported by grants from various funding agencies, including NSF, DOE, DOD, and DHS. Overall, I have received a total of $5.8M of grants, as a PI, Co-PI, or senior investigator. Among these grants, a total of $3.6M has been received as PI. The following shows the list of grants that I have received:

1. Center for Advancing Education and Studies on Critical Infrastructures Resilience. Department of Homeland Security (2017-ST-062-000002). $1.2M. 2017.8–2022.7. PI: Jason Liu; Co-PIs: S.S. Iyengar, Scott Graham, Leonardo Bobadilla, Bogdan Carbunar, Mark Finlayson, Liting Hu, Monique Ross, Ning Xie.
2. GENI Engineering Conference 25. National Science Foundation (CNS-1724805). $50K. 2017.2–2018.2. PI: Julio Ibarra; Co-PIs: Heidi Morgan, Jason Liu.
3. CSR: Medium: Collaborative Research: NVM-Enabled Host-Side Caches. National Science Foundation (CNS-1563883). $648K. 2016.4–2019.5. PI: Raju Rangaswami; Co-PIs: Giri Narasimhan, and Jason Liu (33%).
4. Security Analyses and Applications of Complex Networks: From Theory to Practice. Florida Center for Cybersecurity at University of South Florida. $25K. 2016.3–2017.12. PI: Jason Liu.
5. Creating and Composing SDN Security Modules. Florida Center for Cybersecurity at University of South Florida. $25K. 2016.3–2017.12. PI: Wei Zeng (replacing Xin Sun); SI: Jason Liu (50%).
6. Scalable Discrete Event Simulation for Performance Prediction. Los Alamos National Laboratory. $252K. 2014.10–2017.9. PI: Jason Liu.
7. Vulnerability and Survivability of Cyberspace: Basic Science to Applications. Florida Center for Cybersecurity at University of South Florida. $25K. 2015.3–2016.8. PI: Jason Liu (replacing Ming Zhao).

8. EAGER: SwitchOn – Exploring and Strengthening US-Brazil Collaborations in Future Internet Research. National Science Foundation (CNS-1443285). $200K. 2014.8–2018.2. PI: Jason Liu; Co-PIs: Julio Ibarra, Heidi Alvarez.
9. Enabling Time-sensitive Applications on Virtualized Computing Systems. Department of Defense (#W911NF-13-1-0157). $643K. 2013.5–2017.4. PI: Jason Liu (replacing Dr. Ming Zhao).
10. PrimoGENI Constellation for Distributed At-Scale Hybrid Network Experimentation. National Science Foundation (through Raytheon/GENI Project Office, CNS-1346688). $269K. 2013.10–2016.9. PI: Jason Liu.
11. SoftPM: Streamlining High-End Computing with Software Persistent Memory. National Science Foundation (CCF-0937964). $712K. 2009.9–2013.8. PI: Raju Rangaswami; Co-PIs: Jason Liu (33%), Ming Zhao.
12. PrimoGENI–Developing GENI Aggregates for Real-time Large-scale Network Simulation. National Science Foundation (through Raytheon/GENI Project Office, CNS- 0714770). $502K. 2009.10–2013.7. PI: Jason Liu; Co-PIs: Julio Ibarra, Heidi Alvarez.
13. CREST: Center for Innovative Information Systems Engineering, Subproject 5: Complex System Modeling, Analysis and Realization (CS-MAR). National Science Foundation (HRD-0833093). $825K. 2008.8–2016.12. PI: Xudong He; Co-PIs: Shu-Ching Chen, Peter Clarke, Jason Liu, S. Masoud Sadjadi.
14. CAREER: Immersive Large-Scale Network Simulations. National Science Foundation (CNS-0546712, CNS-0836408). $436K. 2006.3–2012.2. PI: Jason Liu.