In Pervasive Communications Laboratory, our research interests include Cloud computing and datacenters, Big data, Cyber-physical systems, Distributed systems, Mobile computing, High performance computing, and Social networks. The lab gathers outstanding researchers dedicated to solving the critical problems in the related research areas. Our goal is to develop advanced computing technologies that can make fundamental contributions to the real world problems and enable an intelligent future.

Pervasive Communications Laboratory

Haiying Shen
Associate Professor
hs6ms@virginia.edu
https://www.cs.virginia.edu/people/faculty/shen.html

Dept. of Computer Science
University of Virginia
Charlottesville, VA
(434) 924-8271

- Cloud computing and datacenters
- Cyber-physical systems, Internet of things
- Big data, Information retrieval, Content delivery networks
- Mobile computing
- Wireless sensor networks
- High performance computing
- Social networks

University of Virginia

SCHOOL of ENGINEERING & APPLIED SCIENCE

Virtual Resource Management in Datacenters

To provide robust infrastructure as a service (IaaS), clouds currently perform load balancing by migrating virtual machines (VMs) from heavily loaded physical machines (PMs) to lightly loaded PMs. Previous reactive load balancing algorithms migrate VMs upon the occurrence of load imbalance, while previous proactive load balancing algorithms predict PM overload to conduct VM migration. However, both methods cannot maintain long-term load balance and produce high overhead and delay due to migration VM selection and destination PM selection. To overcome these problems, we propose a proactive Markov Decision Process (MDP)-based load-balancing algorithm. We handle the challenges of allying MDP in virtual resource management in cloud datacenters, which allows a PM to proactively find an optimal action to transit to a lightly loaded state that will maintain for a longer period of time. We also apply the MDP to determine destination PMs to achieve long-term PM load balance state. Our algorithm reduces the numbers of Service Level Agreement (SLA) violations by long-term load balance maintenance, and reduces the load balancing overhead (e.g., CPU time, energy) and delay by quickly identifying VMs and destination PMs to migrate.

Towards Green Transportation: Fast Vehicle Velocity Optimization for Fuel Efficiency

To minimize the fuel consumption for driving, several methods have been proposed to calculate vehicles' optimal velocity profiles on a remote cloud. Considering the traffic dynamism, each vehicle needs to keep updating the velocity profile, which requires low latency for information uploading and profile calculation. However, these proposed methods cannot satisfy this requirement due to (1) high queuing delay for information uploading caused by a large number of vehicles, and (2) the neglect of the traffic light and high computation delay for velocity profile. For (1), considering the driving features of close vehicles on a road, e.g., similar velocity and inter-distances, we propose to group vehicles within a certain range and let the leader vehicle in each group to upload the group information to the cloud, which then derives the velocity of each vehicle in the group. For (2), we propose spatial-temporal DP (ST-DP) that additionally considers the traffic lights. We innovatively find that the DP process makes it well suited to run on Spark (a fast parallel cluster computing framework) and then present how to run STDP on Spark.

RECENT GRANTS

- NSF #1404981 Application Characterization for Adaptive Computing Platform Determination for Computational and Data-Enabled Science and Engineering (PI)
- NSF #1254006 A New Efficient and Cooperative Large-Scale Distributed Data Sharing System (PI)
- Microsoft Research, A Hierarchical DHT-Aided Chunk-Driven Overlay for Real-Time Peer-to-Peer Live Streaming (PI)

SEAS Research Information
Pamela M. Norris,
Executive Associate Dean for Research
University of Virginia
Box 400232
Charlottesville, VA 22903
EngineeringResearch@virginia.edu
434.243.7683

University of Virginia | School of Engineering & Applied Science