Research@CSE Systems

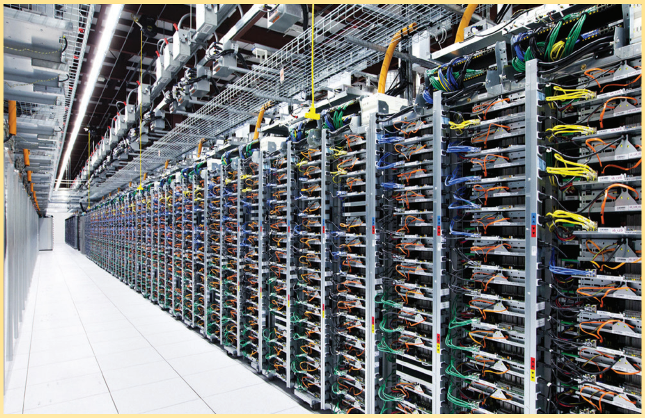

Systems Research at the University of Michigan brings together TK faculty members and their groups to study security and privacy, databases and data mining, operating systems, networking and distributed systems, languages and compilers, and mobile, embedded, and real-time systems.

DATA MINING

Faculty: Michael Cafarella, Mosharaf Chowdhury, H.V. Jagadish, Danai Koutra, Honglak Lee, Barzan Mozafari

The growth of Web services, sensor networks, and high-capacity storage devices have led to an explosion in the quantity and diversity of data suitable for data mining. Statistical techniques are critical for data mining, but are certainly not the only important part; our work also includes novel applications, software infrastructure for large-scale analytics, privacy preservation while mining data, and new interfaces that extend human capabilities to find patterns in data. Results so far include effective prediction of cardiac and epileptic events in medical patients, large-scale data extraction from the Web, large efficiency gains in Hadoop, Spark, and TensorFlow frameworks, and efficient exploration of large-scale real-world networks, including social, communication and brain networks.

OPERATING SYSTEMS AND DISTRIBUTED SYSTEMS

Faculty: Peter Chen, Mosharaf Chowdhury, Robert Dick, Jason Flinn, Sugih Jamin, Manos Kapritsos, Baris Kasikci, Harsha Madhyastha, Z. Morley Mao, Jason Mars, Satish Narayanasamy, Brian Noble, Atul Prakash, Kang Shin, Lingjia Tang

Operating systems and distributed systems research is tackling exciting new problems that span the gamut from embedded systems and sensor networks to Internet-scale distributed systems and rack-scale computers. Operating systems have become pervasive in our daily lives, controlling not just traditional computers, but also consumer electronics devices and cyber-physical systems such as automobiles and power grids. Research projects at the University of Michigan are investigating how to make computers more reliable, more secure, faster, scalable, and easier to use via new services such as deterministic replay and speculative execution.

SECURITY AND PRIVACY

Faculty: Peter Chen, Robert Dick, Roya Ensafi, Kevin Fu, Daniel Genkin, J. Alex Halderman, Peter Honeyman, Manos Kapritsos, Baris Kasikci, Harsha Madhyastha, Z. Morley Mao, Satish Narayanasamy, Brian Noble, Atul Prakash, Kang Shin, Westley Weimer

Security and privacy research at Michigan seeks to advance understanding of how systems can be attacked and defended across all layers of the computing stack. We draw on intellectual tools ranging from theoretical cryptography to applied engineering in order to explore broad aspects of this high-impact area.

CSE faculty and students are pioneering data-driven techniques for understanding security at a global scale, based on Internet-scale empirical measurements and experiments. By applying this quantitative approach to problems ranging from cryptographic vulnerabilities to censorship and blocking, Michigan researchers have discovered entirely new classes of attacks and helped develop defenses that have been implemented in widely used systems.

CSE faculty and students have been investigating privacy and security issues of various embedded devices, systems, and applications, including smartphones and wearable devices, cyber-physical systems like connected or autonomous cars, IoT and medical devices, and sensors. Investigators are also exploring system-level security across a range that spans micro-architectural threats and defenses to cloud and storage security.

Computer security broadly impacts society and public policy, and Michigan researchers play a central role in national and international conversations about cybersecurity issues ranging from election hacking to nation-state censorship, to mass surveillance.

STORAGE SYSTEMS

Faculty: Jason Flinn, Peter Honeyman

Faculty are investigating scalable storage architectures to meet the present needs of cluster and grid computing and the future requirements for enterprise-scale computing. Other projects build distributed storage systems that target mobile and pervasive clients such as laptops, smartphones, and consumer electronics.

COMPUTER NETWORKS

Faculty: Mosharaf Chowdhury, Roya Ensafi, Sugih Jamin, Harsha Madhyastha, Z. Morley Mao, Kang Shin

Topics in computer networks range from wireless networking to mobile computing, Internet networking, software-defined networks, and datacenter networks. Projects cover Internet routing, routing security, and network measurement to understand how to improve the robustness of infrastructure services such as routing, DNS systems. Projects in wireless and mobile networking cover cognitive radio and adaptive networks as well as handset security, including spectrum sensing, adaptation to radio and application environments, MAC & network-layer protocols, and mobile applications. Projects in datacenter networks include application-aware networking and network-aware application design for emerging application domains such as AI training and inference.

EFFICIENT WAREHOUSE AND RACK SCALE COMPUTERS

Faculty: Mosharaf Chowdhury, Jason Mars, Lingjia Tang

Today's warehouse scale computers (WSC) face two major challenges: inefficiency in resource utilization and a mismatch between expected and experienced performance. One of the major sources of inefficiency in WSC lies in the drastic gap between the state-of-the-art software systems in WSCs and modern hardware platform capabilities. We find this SW/HW efficiency gap exists even in the most sophisticated data centers. This gap manifests itself in several aspects and in every level within the WSC system, and our work focuses on designing the software stack, including cluster-level scheduling and server-level runtime systems, to eradicate these inefficiencies.

As networks become faster, an emerging platform for computing is rack scale computers (RaSC). We are building RaSCs that can provide a single machine abstraction out of an entire rack using RDMA-enabled networking and resource disaggregation. The challenges in terms of efficiency and performance still arise from the mismatch between SW and HW layers. We are addressing a similar set of challenges as in WSCs in this exciting, new context of RaSCs for the very first time. efficiency and performance still arise from the mismatch between SW and HW layers. We are addressing a similar set of challenges as in WSCs in this exciting, new context of RaSCs for the very first time.

EMBEDDED AND REAL-TIME SYSTEMS

Faculty: , Robert Dick, Kevin Fu, Sugih Jamin, Jean-Baptiste Jeannin, Brian Noble, Kang Shin

Research in this area includes the following projects:

- Cyber-Physical Systems, integrating a network of embedded real-time computing systems with physical systems like time embedded application ground and space vehicles.
- Wired and wireless networking, including robust Internet protocols, wireless LANs/WANs, cognitive radio networks, and sensor networks.
- Computation and network security, including lightweight systems (e.g., handhelds and sensor networks) and enterprise systems & networks.
- Low-power embedded real-time OS, middleware, and storage systems.
- Virtualization-based resource management for large-scale Internet services and servers.
- Body-Area Networks

Bob and Betty Beyster Building
2260 Hayward Street, Ann Arbor, MI 48109-2121
cse.umich.edu

DATABASE SYSTEMS

Faculty: Michael Cafarella, H.V. Jagadish, Barzan Mozafari

The database group's research is focused on building the data management infrastructure for the twenty-first century, with particular emphasis on issues surrounding Big Data, including stream processing, approximate query answering, text mining, data integration, information extraction, and data sharing. We have a strong emphasis on database usability. Our approach is to understand at a fundamental level what it is about the data model and representation that make it hard to use and query.

In addition, we have a very strong data science effort, with particular emphasis on the effective integration and efficient querying of materials and biological data.

LANGUAGES, COMPILERS, AND RUNTIME SYSTEMS

Faculty: Peter Chen, Jason Flinn, Jean-Baptiste Jeannin, Manos Kapritsos, Baris Kasikci, Scott Mahlke, Jason Mars, Satish Narayanasamy, Lingjia Tang, Westley Weimer

Static and run-time compiler systems are used to get more performance, robustness, and energy efficiency through intelligent program analysis, transformation, and adaptation. Deep analysis of application behavior identifies opportunities to customize the software to the underlying hardware as well as providing an automatic approach to customize hardware functionality to a set of applications. Dynamic profile or run-time data expose opportunities to optimize the common case or even input-specific optimizations.

CSE faculty are developing new compiler methods to optimize applications from mobile to cloud server. Techniques such as software approximate computing unlock performance in applications where 100% output accuracy is neither necessary or even possible. Lightweight dynamic compilation allows an application to be optimized while it is running without disruptions. CSE faculty are also developing effective approaches to automatically fix bugs in software, reducing the cost of software maintenance.

Other projects explore problems at the interface of architecture, operating systems and program analysis, to help programmers write and maintain reliable parallel software, including support for tolerating concurrency bugs, memory consistency models and deterministic replay. Researchers are using advanced medical imaging to explore the neural correlates of programming tasks and expertise.

MOBILE AND CLOUD COMPUTING

Faculty: Mosharaf Chowdhury, , Robert Dick, Jason Flinn, Sugih Jamin, Harsha Madhyastha, Scott Mahlke, Z. Morley Mao, Barzan Mozafari, Brian Noble, Atul Prakash, Kang Shin

Smartphones are increasingly popular as an alternative compute platform. We are investigating several key challenges to address the growing user demand while leveraging the new technology trends in cloud computing and multicore smartphone compute platforms. Ongoing research activities include:

- Accurate online power modeling of mobile applications.
- Dynamic offloading of executions.
- Cellular network policy investigation and security analysis.
- Smartphone workload modeling and benchmarking.
- Large-scale cellular network performance measurement and diagnosis.
- Mobile application profiling and cross layer analysis.
- Multi-resource allocation in cloud computing.
- Geo-distributed analytics across multiple clouds.
- Predictable performance in the cloud.

PARALLEL COMPUTING

Faculty: Satish Narayanasamy, Quentin Stout

Prof. Stout's research is primarily in the area of parallel computing/supercomputing applied to scientific simulations. These projects usually involve large interdisciplinary teams. Since the computers used are sometimes among the largest in the world, there is an emphasis on developing parallel data structures and algorithms that scale to thousands of processors. They also develop software frameworks which allow them to compose large complex simulations, such as a climate model, from pieces such as atmospheric, ocean, and land models. Prof. Narayanasamy works on parallel computer architecture, and its interaction with parallel languages, compilers and operating system. The goal is to make parallel programming as easy as sequential programming. Some of the specific projects include research on deterministic parallel computing, memory consistency models and heterogeneous programming models and runtime support for cloud computing.

APPROXIMATE QUERY PROCESSING

Faculty: HV Jagadish, Barzan Mozafari

With the abundance of massive datasets, traditional data processing tools have now become the bottleneck in many data science scenarios. Even with a few terabytes of data, the fastest database systems today take hours or days to answer the simplest queries. This response time is simply unacceptable to many users and applications. The exploratory data analytics, however, is often an interactive and iterative process: a data scientist forms an initial hypothesis (e.g., by visualizing the data), consults the data, adjusts his/her hypothesis accordingly, and repeats this process until a satisfactory answer is discovered. Thus, slow and costly interactions with data can severely inhibit a data scientist's productivity, engagement, and even creativity. Driven by this growing demand for interactive analytics, Approximate Query Processing (AQP) aims to return immediate response times by providing answers that are 99.9% accurate in most cases.

Bob and Betty Beyster Building
2260 Hayward Street, Ann Arbor, MI 48109-2121
cse.umich.edu

Research@CSE Systems

FACULTY IN SYSTEMS

Michael Cafarella
Associate Professor
michjc
4709 Beyster

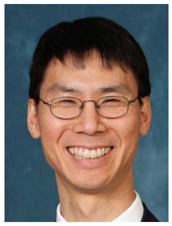

Peter Chen
Arthur F. Thurnau Professor
pmchen
4640 Beyster

Mosharaf Chowdhury
Assistant Professor
mosharaf
4820 Beyster

Roya Ensafi
Research Asst. Professor
ensafi
4745 Beyster

Jason Flinn
Professor
jflinn
4641 Beyster

Daniel Genkin
Assistant Professor
genkin
4753 Beyster

J. Alex Halderman
Professor
jhalderm
4717 Beyster

Peter Honeyman
Lecturer
honey
4777 Beyster

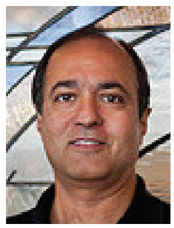

H.V. Jagadish
Bernard A. Galler Professor of EECS
jag
4601 Beyster

Sugih Jamin
Associate Professor
jamin
4737 Beyster

Manos Kapritsos
Assistant Professor
manosk
4824 Beyster

Baris Kasikci
Assistant Professor
barisk
4816 Beyster

Danai Koutra
Assistant Professor
dkoutra
3633 Beyster

Walter Lasecki
Assistant Professor
wlasecki
2636 Beyster

Harsha Madhyastha
Associate Professor
harshavm
4611 Beyster

Z. Morley Mao
Professor
zmao
4629 Beyster

Barzan Mozafari
Assistant Professor
mozafari
4769 Beyster

Brian Noble
Professor and CSE Chair
bnoble
3713 Beyster

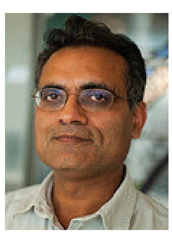

Atul Prakash
Professor
aprakash
4741 Beyster

Kang G. Shin
Kevin and Nancy O'Connor Professor of Computer Science
kgshin
4605 Beyster

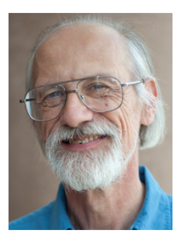

Quentin Stout
Professor
qstout
3605 Beyster

Westley Weimer
Professor
weimerw
4636 Beyster

Research@CSE Systems

AFFILIATED FACULTY

Mark Ackerman - Computer Science and Engineering, Interactive Systems

Eytan Adar - School of Information

Robert Dick - Electrical and Computer Engineeing

Alfred Hero - Electrical and Computer Engineering

Jean-Baptiste Jeannin - Aerospace Engineering

Honglak Lee - Computer Science and Engineering, Artificial Intelligence

Scott Mahlke - Computer Science and Engineering, Hardware Systems

Jason Mars - Computer Science and Engineering, Hardware Systems

Qiaozhu Mei - School of Information

Rada Mihalcea - Computer Science and Engineering, Artificial Intelligence

Satish Narayanasamy - Computer Science and Engineering, Hardware Systems

Shai Revzen – Electrical and Computer Engineering

Lingjia Tang - Computer Science and Engineering, Hardware Systems