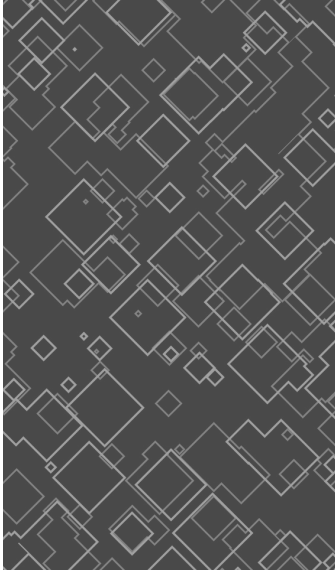

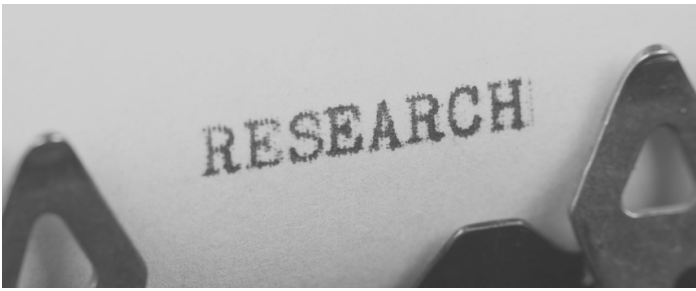

RED HAT RESEARCH DAY

Red Hat Summit

May 6 | Boston
Convention Center

Welcome

It gives me great pleasure to welcome you all to the first annual Red Hat Research Day at the Red Hat Summit 2019.

Red Hat Research is dedicated to connecting researchers with Red Hat engineers, customers, and partners, to move great research ideas into open source communities. Events like this one are a great help to that aim, and I very much appreciate all of you taking the time to come here and participate. In the morning, our data governance track examines the problems around big data, privacy, provenance and research -- as a company dedicated to the open source development model, we are deeply interested in the availability of open data sets for AI development. In the afternoon, we will look at the really interesting work that is being prompted by the increasing expense of developing smaller and faster processors. We are challenged to find smarter ways to increase performance than simply scaling up processor speed, and our speakers have some very interesting things to say on that topic.

We hope you enjoy the event. If you have any follow-up questions for Red Hat Research, you can email the group at academic@redhat.com, or contact me directly.

Thanks,

Hugh Brock
Research Director, Red Hat Office of the CTO

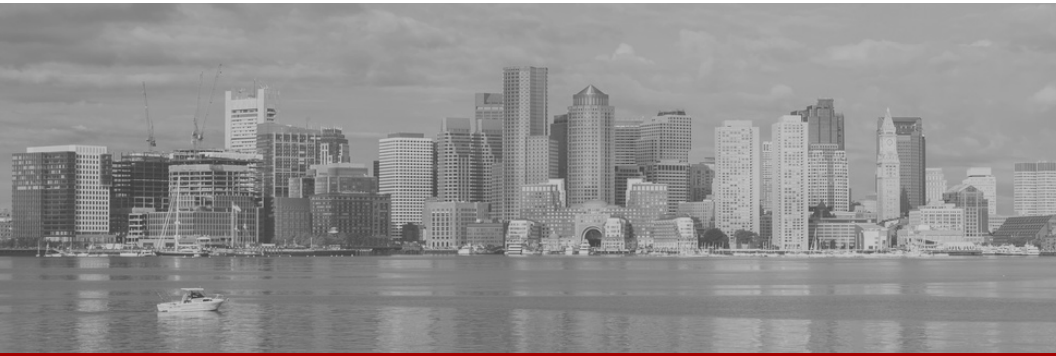

Schedule of Events

MORNING

8:30 AM	**CONTINENTAL BREAKFAST**
9:00 AM	Looking Into the Future of Collaborative Research ***Hugh Brock***
9:15 AM	Privacy and Machine Learning—Challenges and Opportunities ***Azer Bestavros***
9:30 AM	An Architecture Stack for Data-Driven Infrastructure Management ***Dimosthenis Kyriazis***
10:00 AM	Parallel Machine Learning ***Assaf Schuster***
10:30 AM	Model-Driven Discovery and the Science of Data ***Jeffrey Brock***
11:00 AM	Implementing Secure Multi-Party Computing ***Kinan Bab***
11:30 AM	Homomorphic Encryption, Why and How ***Kurt Rohloff***

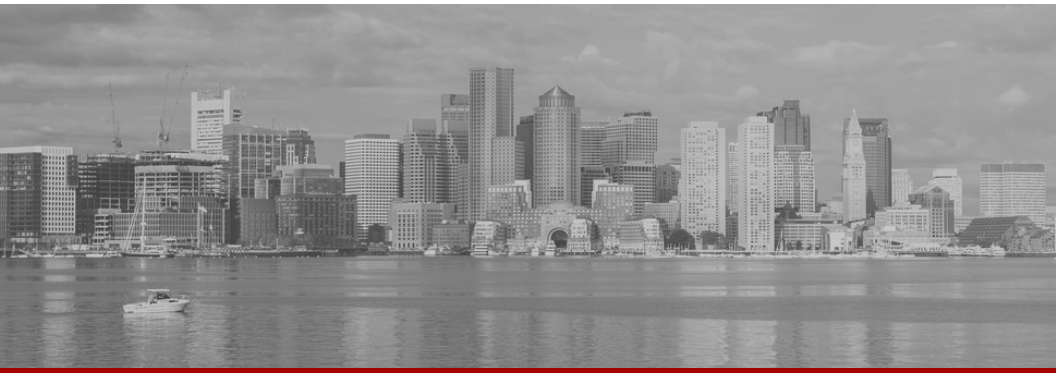

Schedule of Events

AFTERNOON

12–1 PM LUNCH AND NETWORKING

1:00 PM Keynote
Chris Wright

1:30 PM Operating Systems and Hardware Innovation
Orran Krieger, Uli Drepper

2:00 PM A Partitioning Hypervisor for Latency-Sensitive Workloads
Craig Einstein, Richard West, Bandan Das

2:30 PM UniKernel Linux (UKL)
Ali Raza, Larry Woodman

3:00 PM FPGAs Everywhere in Large-Scale Computer Systems
Ahmed Sanaullah, Martin Herbordt

3:30 PM With a Little Help from My Threads: Accelerating Single-Thread Execution with Speculating Hyperthreads
Tommy Unger, Johathan Appavoo

4:00 PM Automatic Configuration of Complex Hardware
Han Dong, Sanjay Arora

4:30 PM Removing Memory as a Noise Factor
Parul Sohal, Renato Mancuso

5–7 PM NETWORKING PARTY WITH THE OFFICE OF THE CTO

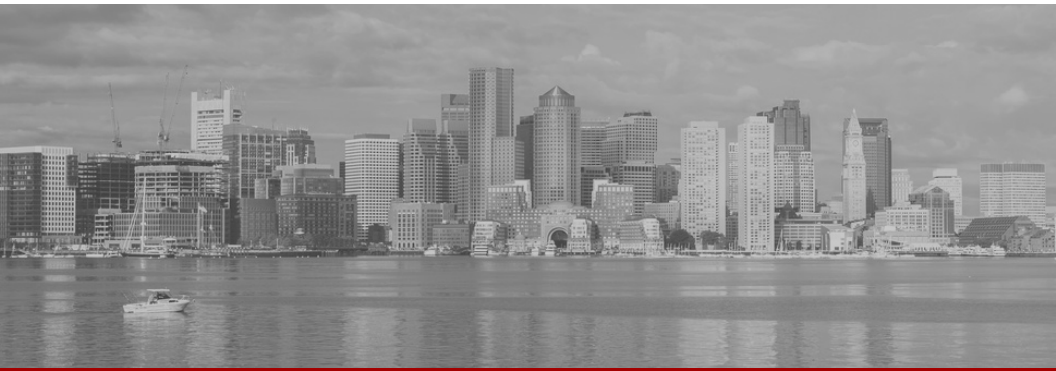

Talks

9:00 AM *Hugh Brock*

Looking Into the Future of Collaborative Research

9:15 AM *Azer Bestavros—track chair*

Privacy and Machine Learning—Challenges and Opportunities

9:30 AM *Dimosthenis Kyriazis*

An Architecture Stack for Data-Driven Infrastructure Management

In several application domains, multi-modal data are exploited towards the provision of innovative services with business value. Data management and analytics frameworks to date emphasize the computational needs and aspects of applications and services deployed in a particular infrastructure. In this talk, the architecture of a complete stack (namely BigDataStack) will be presented. This stack, based on an infrastructure management system that aims at driving decisions according to data aspects, is fully scalable, runtime adaptable and high-performant to tackle the needs of big data operations and data-intensive applications. Furthermore, the stack goes beyond purely infrastructure elements by introducing techniques for dimensioning big data applications, modeling and analyzing processes as well as provisioning data-as-a-service by exploiting a seamless analytics framework. Technical elements of the talk are linked with specific application scenarios that utilize the components of the BigDataStack architecture.

Talks

10:00 AM *Assaf Schuster*

Parallel Machine Learning

Modern deep neural networks are comprised of millions of parameters, which require massive amounts of data and time to train. Steady growth of these networks over the years has made it impractical to train them from scratch on a single GPU. Distributing the computations over several GPUs can drastically reduce this training time, however, stochastic gradient descent (SGD), which is typically used to train these networks, is an inherently sequential algorithm. As a result, training deep neural networks on multiple workers is difficult, especially when using non-dedicated cloud resources trying to maintain high efficiency, scalability, and final accuracy. In this talk we will survey some of the new ideas in this scope and discuss their potential.

10:30 AM *Jeffrey Brock*

Model-Driven Discovery and the Science of Data

In the age of machine learning, artificial intelligence, and data science, we see all around us the impact of powerful tools to extract knowledge from data. Whether in advertising, political campaigns, real-time translation tools, or in the hard sciences, the face of the knowledge frontier has a new complexion. Many of these tools, such as neural networks and deep learning, work alarmingly and uncannily well, and yet we do not fully understand why. Vital questions confront those of us in the academy: how do these tools change the way we understand knowledge acquisition? How do they change how we read texts or analyze political discourse? How do they force us to rethink the scientific method, and how do they allow us to search for new models, theories, and equations that govern the universe?

While tools from machine learning have provided elaborate taxonomies of observed data, innumerable examples show that without a coherent model for a data generating mechanism, spurious conclusions abound. Yet these tools are being used in science at an increasing rate, often with alarming apparent efficacy and impact. In this talk I will engage a pressing modern question: how as a scientific community and as a society can we come to terms with the introduction of these incredible pattern recognition machines, and what must the new mathematics of data do to assure that science remains stable, reproducible, and open?

Talks

11:00 AM *Kinan Bab*

Implementing Secure Multi-Party Computing

Secure Multiparty Computation (MPC) is a cryptographic primitive that allows several parties to jointly and privately compute desired functions over secret data. Building and deploying practical MPC applications faces several obstacles, including performance overhead, complicated deployment and setup procedures, and adoption of MPC protocols into modern software stacks. MPC applications expose trade-offs between efficiency and privacy that may be hard to reason about, formally characterize, and encode in a protocol design or implementation.

We describe technical and non-technical challenges from our experience deploying MPC applications in the real world. We showcase JIFF: an extensible general purpose MPC framework capable of running on web and mobile stacks, showing how developments in distributed systems, web development, and the SMDI paradigm can inform MPC constructs and implementation. JIFF is used to implement several MPC applications, including a successfully deployed study on economic opportunity for minority owned businesses in the Boston area, and a service for efficient privacy-preserving route recommendation.

11:30 AM *Kurt Rohloff*

Homomorphic Encryption, Why and How

The discovery of Fully Homomorphic Encryption (FHE) has been one of the major breakthroughs of computer science in the 21st century. FHE allows sensitive data to be encrypted such that arbitrary programs can be securely run over the encrypted data where the output, when decrypted, is equivalent to the result of running the original algorithm on the unencrypted data. FHE is ground-breaking in its ability to enable learning for AI and ML on encrypted data. This talk will review our advances in FHE, from theory, implementation and application perspectives, with a focus on commercially-relevant applications. We focus on regulated industry applications, such as those in the financial and healthcare domains.

Talks

1:00 PM *Chris Wright*

Keynote

1:30 PM *Orran Krieger, Uli Drepper—track chairs*

Operating Systems and Hardware innovation

2:00 PM *Craig Einstein, Richard West, Bandan Das*

A Partitioning Hypervisor for Latency-Sensitive Workloads

Quest-V is a separation kernel that partitions services of different criticality levels across separate virtual machines, or sandboxes. Each sandbox encapsulates a subset of machine physical resources that it manages without requiring intervention from a hypervisor. In Quest-V, a hypervisor is only needed to bootstrap the system, recover from certain faults, and establish communication channels between sandboxes. Partitioning VMs onto separate machine resources offers an opportunity for per-sandbox power management. Depending on the latency and power demands of each sandbox, it can be suspended to RAM or disk and optionally migrated across hosts to balance system resources and reduce power consumption. Shared machines can be placed into low power states when all sandboxes migrate away from them. Quest-V allows VMs to suspend and resume individual hardware resources without interfering with the operation of other VMs on the same physical platform. This allows for the creation of systems that are both power and latency aware.

Talks

2:30 PM *Ali Raza, Larry Woodman*

UniKernel Linux (UKL)

Unikernels are small, lightweight, single address space operating systems, with the kernel included as a library with the application. Because unikernels run a single application, there is no sharing or competition for resources among different applications, improving performance and security. Unikernels have thus far seen limited production deployment. This project aims to turn the Linux kernel into a unikernel with these characteristics: 1) easily compiled for any application, 2) uses battle-tested, production Linux and glibc code, 3) allows the entire upstream Linux developer community to maintain and develop the code, and 4) provides applications normally running vanilla Linux to benefit from unikernel performance and security advantages. UniKernel Linux (UKL) provides the opportunity to pursue many interesting research ideas, e.g. studying advantages of bypassing syscalls and directly invoking internal kernel functionality, studying impacts of link time optimizations across application/kernel boundaries, studying performance benefits from profile driven optimizations, and observing performance gains from simplified user level synchronization mechanisms.

3:00 PM *Ahmed Sanaullah, Martin Herbordt*

FPGAs Everywhere in Large-Scale Computer Systems

As modern data center workloads become increasingly complex, constrained and critical, mainstream "CPU-centric" computing can no longer keep pace. Future data centers are moving towards a more fluid model, with computation and communication no longer localized to commodity CPUs and routers. Next generation "data-centric" data centers will "compute everywhere," whether data is stationary (in memory) or on the move (in network). Reconfigurable hardware, in the form of Field Programmable Gate Arrays (FPGAs), are transforming ordinary clouds into massive supercomputers. We will highlight many ways to deploy FPGAs in a data center node, such as traditional back-end accelerators, tightly coupled off-load processors, Smart NICs, Bump-in-the-Wire, and even in the router itself. We will also discuss our efforts to make these devices globally accessible, through deeper integration into software stacks, transparent generation of custom hardware stacks, and device management using reconfigurable hardware operating systems.

Talks

3:30 PM *Tommy Unger, Jonathan Appavoo*

With a Little Help from My Threads: Accelerating Single-Thread Execution with Speculating Hyperthreads

To solve problems, humans tend to synthesize known facts with some amount of new effort. When interesting problems are solved mainly with the former, the solution might be recognized as elegant. To date, we've struggled to get computing machines involved in much elegant problem solving. Considering constraints like budget caps and the polar ice caps, this lack of elegance becomes more than an aesthetic issue.

In this talk, I'll present our research into using read-only "snapshots" as a tool for reducing the "new effort" required to solve computational problems in a cloud "Functions as a Service" (FaaS) case study. Dropping our prototype in as a replacement backend for the Apache OpenWhisk FaaS platform allows a compute node to cache multiplicatively more functions, and to reduce latencies on cache misses. I'll conclude with plans for summer work making RISC-V FPGA softcores amenable to this trick.

4:00 PM *Han Dong, Sanjay Arora*

Automatic Configuration of Complex Hardware

A modern network interface card (NIC) such as the Intel X520 10 GbE is complex, with hardware registers that control every aspect of the NIC's operation, from device initialization to dynamic runtime configuration. The Intel X520 datasheet documents over 5600 registers; only ~1890 are initialized by a modern Linux kernel. It is thus unclear what the performance impact of tuning registers on a per application basis will be. We pursue three goals towards this understanding: 1) Identify, via a set of microbenchmarks, application characteristics that will illuminate mappings between hardware register values and their corresponding microbenchmark performance impact. 2) Use these mappings to frame NIC configuration as a set of learning problems, such that an automated system can recommend hardware settings corresponding to each network application. 3) Introduce either new dynamic or application instrumented policy into the device driver in order to better attune dynamic hardware configuration to application runtime behavior.

Talks

4:30 PM *Parul Sohal, Renato Mancuso*

Removing Memory as a Noise Factor

Memory bandwidth is increasingly the bottleneck in modern systems and a resource that, until today, we could not schedule. This means that, depending on what else is running on a server, performance may be highly unpredictable, impacting the 99% tail latency which is increasingly important in modern distributed systems. Moreover, the increasing importance of high-performance computing applications, such as machine learning, and Real-Time Systems demands more deterministic performance, even in shared environments. Alternatively, many environments resist running more than one workload on a server, reducing system utilization. Recent processors have started introducing the first mechanism to monitor and control memory bandwidth. Can we use these mechanisms to enable machines to be fully used while ensuring that primary workloads have deterministic performance? We present early results from using Intel's Resource Director Technology and some insight into this new hardware support. We also look at an algorithm to use these tools to provide deterministic performance on different workloads.

Speakers

JONATHAN APPAVOO

Prior to joining Boston University as an Associate Professor in 2009, Jonathan Appavoo was a Research Staff Member at IBM Research in New York. He graduated with his PhD in 2006 from the University of Toronto, under the supervision of Professor Michael Stumm.

SANJAY ARORA

Sanjay Arora is a data scientist in the AI Center of Excellence at Red Hat. He first started his data science journey as a physicist at the CMS Experiment at CERN, followed by stints in finance and technology consulting. His main interests are machine learning applied to systems (operating systems, compilers, cloud deployments etc.), theoretical machine learning, and applied mathematics.

KINAN BAB

Kinan is a PhD student in the Department of Computer Science at Boston University. His research lies in implementing and reasoning about software and systems. In particular, Kinan works on designing cryptographic protocols for distributed systems at scale, to ensure security and privacy, and on developing formal program analysis techniques and novel programming models and frameworks to reason about the correctness and security of whole systems at large. A Hariri Institute Graduate Student Fellow, and a Software Engineering Fellow at the Software and Application Innovation Lab (SAIL) in BU, Kinan is also a visiting researcher with the NETMIT research group at CSAIL MIT, working with Prof. Dina Katabi on the Emerald project. He received a B.S. in Computer Science in 2015 from the American University of Beirut.

Speakers

AZER BESTAVROS

Azer Bestavros is Warren Distinguished Professor of Computer Science and the Founding Director of the Hariri Institute for Computing at Boston University, which was set up in 2011 as an incubator for high-risk, high-reward cross-disciplinary collaborations. Prior to his role at the Hariri Institute, he chaired the Computer Science Department from 2000 to 2007, having joined it in 1991 after completing his PhD at Harvard University. His research in networking, distributed computing, cybersecurity, and high-assurance systems led to seminal contributions that include pioneering studies of web push caching through content distribution networks, self-similar Internet traffic characterization, game-theoretic cloud resource management, and safety certification of networked systems and software. In 2017, he was named a William Fairfield Warren Distinguished Professor, the highest distinction bestowed upon senior faculty members at Boston University, for “representing our community with distinction, enriching the academic experience for our students, and raising our stature as a major research university.”

HUGH BROCK

Hugh Brock is the Research Director for Red Hat, coordinating Red Hat’s research collaboration with universities, governments, and industry worldwide. He relocated to Boston in 2018 to found Red Hat Research after a three-year stint in Brno, Czechia, where he led a large team of developers working on Red Hat’s OpenStack management tool, OSP-director. A Red Hat manager since 2008 and a Red Hatter since 2002, Hugh brings intimate knowledge of the complex relationship between upstream projects and shippable products to the task of finding and cultivating research that Red Hat will be able to bring into the open source world.

Speakers

JEFF BROCK

Jeffrey Brock is Professor of Mathematics, and Dean of Science at Yale University. His research focuses on low dimensional geometry and topology, particularly hyperbolic geometry. His work on William Thurston's program to understand hyperbolic 3-manifolds led to their geometric classification in joint work with R. Canary and Y. Minsky. More recently, he has worked to study renormalized volume and its gradient flow on deformation spaces. He is also interested in geometric and topological methods in analysis of large, complex data sets. He was an undergraduate at Yale, and obtained his PhD at U.C. Berkeley, after which he held positions at Stanford and the University of Chicago before moving to the Brown University Math Department, which he chaired from 2013 to 2017. In 2016 he served as founding Director of Brown's Data Science Initiative. He was a Guggenheim Fellow in 2008.

BANDAN DAS

Bandan Das is a software engineer in the Virtualization group at Red Hat. He works on KVM in the Linux kernel and device emulation in Qemu and considers himself a generalist. He has previously worked in development roles, both in industry and academia and is passionate about systems architecture and performance.

HAN DONG

Han Dong is a 5th year PhD candidate in Computer Science at Boston University. His research interests lie in the areas of distributed systems, high performance computing, and cloud computing. Currently, he is exploring the interactions between hardware and software systems and its implications due to the proliferation of hardware accelerators.

Speakers

ULI DREPPER

Ulrich Drepper joined Red Hat again in 2017, after a seven year hiatus when he worked for Goldman Sachs. He is part of the office of the CTO and concentrates on technologies for machine learning and high-performance computing. His main interests are in areas of low-level technologies like machine and processor architectures, programming languages, compilers, high-performance and low-latency computing. In addition, he is interested in using statistics and machine learning for performance analysis of programs and security of application and OS environments. Ulrich received his Diploma in Informatics from the University of Karlsruhe, Germany.

CRAIG EINSTEIN

Craig Einstein is a Computer Science PhD student at Boston University under the advisement of Professor Richard West. Prior to starting his PhD, he received his BA in Geophysics and Planetary Sciences in 2016 from Boston University. He researches computer systems with an emphasis on real-time systems and autonomous control. Upon the completion of his doctorate, he would like to work in the space industry developing systems to make space travel more efficient and accessible.

MARTIN HERBORDT

Martin Herbordt is Professor of Electrical and Computer Engineering at Boston University where he directs the Computer Architecture and Automated Design Lab. His research spans Architecture and High Performance Computing. He and his group have been working for many years in accelerating HPC applications with FPGAs and GPUs, and in building systems integrating FPGAs. More recently their focus has been on middleware and system aspects of large-scale FPGA clusters and clouds.

Speakers

ORRAN KRIEGER

Orran Krieger is the lead on the Massachusetts Open Cloud, Founding Director for the Cloud Computing Initiative (CCI) at BU, Resident Fellow of the Hariri Institute for Computing and Computational Science & Engineering, and a Professor at the Department of Electrical and Computer Engineering Boston University. Before coming to BU, he spent five years at VMware starting and working on vCloud. Prior to that he was a researcher and manager at IBM T. J. Watson, leading the Advanced Operating System Research Department. Orran did his PhD and MASc in Electrical Engineering at the University of Toronto.

DIMOSTHENIS KYRIAZIS

Dimosthenis Kyriazis is an Assistant Professor at University of Piraeus (Department of Digital Systems). He received his diploma from the school of Electrical and Computer Engineering of the National Technical University of Athens (NTUA) in 2001 and his MSc degree in Techno-economics in 2004. Since 2007, he holds a PhD in the area of Service Oriented Architectures with a focus on quality aspects and workflow management. His expertise lies with service-based, distributed and heterogeneous systems. Before joining University of Piraeus, he was a Senior Research Engineer at the Institute of Communication and Computer Systems (ICCS) of NTUA, having participated and coordinated several European-funded projects (e.g. BigDataStack, 5GTANGO, MATILDA, CrowdHEALTH, ORBIT, COSMOS, VISION Cloud, IRMOS, etc.) focusing his research on issues related to distributed computing architectures, data management and analytics, performance modelling, deployment and management of virtualized infrastructures and platforms.

Speakers

RENATO MANCUSO

Renato Mancuso is an assistant professor at the department of Computer Science at Boston University. He received his PhD from the University of Illinois at Urbana-Champaign (UIUC) in 2017. His research focuses on real-time and embedded systems. He is especially interested in partially-reconfigurable platforms and OS-level multi-core resource management technologies for high-performance, safety-critical systems. He is also interested in applications and methodologies to design, deploy and analyze Cyber-Physical Systems (CPS), and in real-time cloud computing. His research also touches on aspects of security for embedded systems and technologies for UAVs.

ALI RAZA

Ali Raza is a PhD candidate in the Department of Computer Science at Boston University. He works with Prof. Orran Krieger and Prof. Renato Mancuso. His area of interest is Operating Systems, specifically turning the Linux kernel into a unikernel, optimizing it based on the application, and studying the performance gains.

KURT ROHLOFF

Kurt Rohloff is the CTO of Duality Technologies. He is a recipient of a DARPA Young Faculty Award and has been the PI on multiple DARPA and IARPA projects in areas of data science and privacy. He is an associate professor of computer science at NJIT. Prior to Duality, he was a senior scientist in the Distributed Systems research group at Raytheon BBN Technologies. He is the co-founder of the open-source PALISADE lattice crypto software library that supports homomorphic encryption. He received his Bachelor's degree in Electrical Engineering from Georgia Tech and his Master's and PhD. in Electrical Engineering from the University of Michigan.

Speakers

AHMED SANAULLAH

Ahmed Sanaullah is a fourth year Computer Engineering PhD student at Boston University. His research in the BU-Red Hat Collaboratory is aimed at reconfigurable-hardware Operating Systems for FPGA enhanced High Performance and Cloud computing. Other research projects include performance-programmability using High Level Synthesis, and real-time machine learning architectures.

ASSAF SCHUSTER

Prof. Assaf Schuster of the Computer Science Department is the head of the new AI center at the Technion. He is a Fellow of the ACM and the IEEE, with more than 200 published papers in highly selected venues. His interests and publications are in the wide scope of distributed and scalable data mining, big and streaming data technologies including management, analytics & prediction, cyber security and system/IoT vulnerabilities, privacy preserving, cloud resource management and more. He consulted leading hi-tech companies and participated in the bumpy journey of several startups, two of which he co-founded.

PARUL SOHAL

Parul Sohal is a second year PhD student in Computer Science from Boston University, working with Prof. Orran Krieger and Prof. Renato Mancuso. Her interests lie in operating systems and performance determinism in Real-Time, and Cloud Systems.

Speakers

TOMMY UNGER

Tommy Unger is a PhD candidate at Boston University and a summer 2019 Red Hat intern. He builds operating systems that learn about runtime characteristics of workloads for the purpose of speeding up execution. He's excited to learn a bit about RISC-V cores and to look for opportunities to use one "helper core" to accelerate another. When he's not in the lab, you might see him on a bike or in a canoe.

RICHARD WEST

Rich West joined the BU Department of Computer Science in 2000 after completing his PhD at Georgia Tech. Rich is a tinkerer of systems, notably, but not limited to, those in embedded and real-time computing. He likes to take a principled approach to system design, having dabbled in the development of standalone kernels and resource management policies where safety and predictability are paramount. He has studied real-time scheduling and resource management, cache-aware performance of multicore processors, and machine virtualization, amongst other topics. He is currently leading the development of the Quest real-time operating system for multicore processors. Its sister system, Quest-V is a secure and predictable separation kernel that forms a distributed system on a chip, providing efficient, predictable and safe execution of sandboxed guest systems including Linux.

Speakers

LARRY WOODMAN

Larry Woodman has over 36 years of operating system kernel development experience. He started contributing to the AT&T system V UNIX kernel in 1981 at Bell Labs. Over the years he has worked on several versions of UNIX as well as several proprietary operating systems for a variety of hardware and software vendors. For the past 20 years, Larry has been working on the Linux kernel, mostly for Red Hat. The majority of his work has been in the memory management, process management and inter-process communication spaces but also includes work on file systems,hardware support code and device drivers. Larry holds Bachelor's degrees in Aeronautical Engineering and Computer Science and Master's degrees in Computer Science and Computer Engineering.

CHRIS WRIGHT

Chris Wright is vice president and chief technology officer (CTO) at Red Hat. Wright leads the Office of the CTO, which is responsible for incubating emerging technologies and developing forward-looking perspectives on innovations. During his more than 20 years as a software engineer, Wright has worked in the telecommunications industry on high availability and distributed systems, and in the Linux industry on security, virtualization, and networking. He has been a Linux developer for more than 15 years, most of that time spent working deep in the Linux kernel. He is passionate about open source software serving as the foundation for next generation IT systems.

Notes

Connect

RED HAT RESEARCH QUARTERLY

Want to know more about Red Hat Research? Check out our new magazine Red Hat Research Quarterly, downloadable from research.redhat.com/quarterly.

If you would like to receive a print copy, email your mailing address to academic@redhat.com.

TALK TO US

Have an idea for a research project, or want to see what we are working on worldwide? research.redhat.com has status on all our current research projects, plus ways to submit new ideas.

If you have any questions at all about Red Hat Research, email us at academic@redhat.com.