Red Hat Research Day EUROPE 2020

PROGRAM

January 23 | Brno | Hotel Passage

Red Hat
Research

research.redhat.com

Foreword

It gives me great pleasure to welcome you all to Red Hat Research Day Europe 2020.

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.

Red Hat Research began here in Brno, and today we are fortunate to have tracks chaired by our partners at two of Brno's great research universities. Vashek Matyáš of Masaryk University Faculty of Informatics leads a track showcasing the security work we are involved in, with some very interesting guest appearances from the Technical University in Graz, Austria. Meanwhile Tomáš Vojnar of Brno University of Technology hosts a session on formal verification, including our work on automatically verifying kernel ABI changes. Finally, our own Viktor Malík hosts a track on AI from infrastructure up to image and sound analysis.

We hope you enjoy Research Day. 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

Agenda

8:00 AM *Registration open*

TRACK: Data-Intensive Science and Software, chair: Viktor Malík

9:00 AM **Open Cloud Testbed Developing and Testbed for Research: Exploring Next-Generation Cloud Platforms**
Michael Zink

9:30 AM **Machine Learning for Adaptive Human Learning**
Radek Pelánek

10:00 AM **Acoustic Identification of Cetaceans**
Georgia Atkinson

10:30 AM **Automating Computational Placement in IoT Environments**
Peter Michalák

11:00 AM **OpenShift-based High-Performance Computing for Research in Astrophysics**
Filip Hubík, Nikolaos Moraitis, Gabriel Szász & Zdeněk Švécar

11:30 AM **Avoiding Bad Decisions and Heuristics**
Ulrich Drepper

12:00 PM *Lunch*

TRACK: Security and Privacy, chair: Vashek Matyáš

1:00 PM **Observing Developers Interacting with TLS Certificates**
Martin Ukrop

1:30 PM **Analyzing, Breaking and Improving Certified Cryptographic Hardware (TPMs and Smartcards)**
Petr Švenda

2:00 PM **Improving Disk Encryption in Linux**
Milan Brož

2:30 PM **Leaky Processors: Lessons from Spectre, Meltdown, and Foreshadow**
Jo Van Bulck & Daniel Gruss

3:00 PM **Plundervolt: Pillaging and Plundering SGX with Software-based Fault Injection Attacks**
Kit Murdock

3:30 PM *Coffee break*

TRACK: Code Analysis and Verification, chair: Tomáš Vojnar

4:00 PM **Formal Verification of a Linux Distribution**
Kamil Dudka

4:15 PM **Symbiotic: Program-Slicing Framework**
Marek Chalupa

4:30 PM **Heavy-Duty Program Analysis with DIVINE**
Vladimír Štill

4:45 PM **Run-time Verification Using Testos**
Aleš Smrčka

5:00 PM **Performance Versioning System**
Tomáš Fiedor

5:30 PM **Efficient Run-time Verification for the Linux Kernel**
Daniel Bristot de Oliveira

7:00 PM *Social Event with Reception*
VUT Dvorana (BUT Rector's office)
Antoninska 1, Brno (5-minute walk)

List of Speakers


Michael Zink


Radek Pelánek


Georgia Atkinson


Gabriel Szász


Zdeněk Švécar


Ulrich Drepper


Jo Van Bulck


Daniel Gruss


Kit Murdock


Aleš Smrčka


Tomáš Fiedor


D. Bristot de Oliveira


Peter Michalák


Filip Hubík

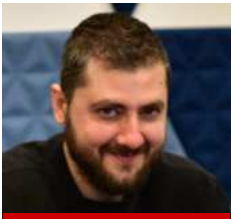
Nikolaos Moraitis


Martin Ukrop


Petr Švenda


Milan Brož


Kamil Dudka


Marek Chalupa


Vladimír Štill

Chairs:

Viktor Malík	Data-Intensive Science and Software
Vashek Matyáš	Security and Privacy
Tomáš Vojnar	Code Analysis and Verification

Talks

Data-Intensive Science and Software

9:00 AM

Open Cloud Testbed Developing and Testbed for Research: Exploring Next-Generation Cloud Platforms

Michael Zink
University of Massachusetts Amherst

Biography

Michael Zink is an Associate Professor in the Electrical and Computer Engineering Department at the University of Massachusetts Amherst. He received his Ph.D. and M.S. in Electrical Engineering from Darmstadt University of Technology in 1997 and 2003, respectively. His research interests are in cyberphysical systems, cloud computing, and Future Internet Architecture. He has been involved in the creation of research infrastructure through his involvement in the NSF GENI and NSF Cloud initiatives. In the latter case, he serves as a Co-PI on the CloudLab project. He is also PI for Open Cloud Testbed, a new NSF project, that has the goal to support the research activity of the systems community in the area of cloud computing.

Abstract

Cloud testbeds are critical for enabling research into new cloud technologies – research that requires experiments that potentially change the operation of the cloud itself. Several such testbeds have been created in the recent past (e.g., Chameleon, CloudLab, etc.) with the goal to support the CISE systems research community. It has been shown that these testbeds are very popular and heavily used by the research community. Testbed utilization often reaches 100%, especially ahead of deadlines for major systems conferences, while there are also periods of modest (<40%) testbed usage.

Data-Intensive Science and Software

In my talk, I will present our NSF "Open Cloud Testbed" (OCT) project, which has the goal to enable elastic cloud testbeds for systems research. Eventually, OCT will allow cloud testbeds to grow and shrink by allocating and deallocating additional resources from compute facilities like production clouds and HPC clusters. Within the OCT project, we will create a prototype elastic cloud testbed, which will combine proven software technologies from both the CloudLab and the Massachusetts Open Cloud (MOC) projects. It will also combine a research cloud testbed (CloudLab) with a production cloud (MOC) through OCT's tight integration with the latter and federation with CloudLab. In addition, OCT will provide programmable hardware (FPGAs) as Bump-in-the-Wire (BITW) capabilities not present in other facilities available to researchers today. The combination of a testbed and production cloud allows a) larger scale compared to isolated testbeds, b) reproducible experimentation based on realistic user behavior and applications, as well as c) a model for transitioning successful research results to practice. OCT offers a unique sustainability model, by allowing additional compute resources to be dynamically moved from institutional uses into the testbed and back again, providing a path to growth beyond the initial testbed.

9:30 AM

Machine Learning for Adaptive Human Learning

Radek Pelánek
Masaryk University

Biography

Radek Pelánek received his Ph.D. degree in Computer Science from Masaryk University for his work on formal verification. Since 2010 his research interests focus on areas of educational data mining and artificial intelligence in education. Currently, he is the leader of the Adaptive Learning group at Masaryk University and is interested in both theoretical research in user modeling and practical development of adaptive learning systems.

Talks

Data-Intensive Science and Software

Abstract

Our Adaptive Learning Group (Faculty of Informatics, MU Brno) has developed adaptive educational systems for several domains (e.g., programming, mathematics, grammar, geography). I will provide an overview of our systems and discuss machine learning techniques that we use to improve the learning of our (human) users. I will also provide examples of our research, which often deals with methodological issues in the evaluation of machine learning techniques that are relevant beyond the development of educational applications.

10:00 AM

Acoustic Identification of Cetaceans

Georgia Atkinson
Newscastle University

Biography

Georgia Atkinson is a PhD student studying bioacoustics within the Cloud Computing for Big Data EPSRC Centre for Doctoral Training scheme at Newcastle University. She has a background in Mathematics and Statistics and is currently researching techniques in Signal Processing and Deep Learning that she is applying to acoustic data gained from the North Sea to identify individual dolphins within a species.

Abstract

Modelling cetacean (dolphins, whales and porpoises) population dynamics is paramount for effective conservation and population management. Methodologies for cetacean research include passive acoustic monitoring (PAM) which allows for monitoring of cetacean occurrence and behaviour ecology through underwater recording. Due to high volumes of data collected and stored in PAM systems, there is a need for automated solutions that can detect and classify cetacean vocalisations. In this talk, we shall discuss how signal processing and deep learning techniques can be applied to this problem and see initial results for detection.

10:30 AM

Automating Computational Placement in IoT Environments

Peter Michalák
Newcastle University

Biography

Peter is a PhD student at the EPSRC Centre for Doctoral Training in Cloud Computing for Big Data at Newcastle University. He previously worked as an R&D software developer in Tieto Finland Oy, received B. Eng. in Computer Software Engineering from JAMK University of Applied Sciences, Finland, and Bc. (with distinction) in Computer Engineering from the University of Žilina, Slovakia.
Peter's research interests include distributed computing, Internet of Things, and real-time event processing. He also enjoys travelling, indoor climbing and yoga.

Abstract

The suitability of high-level declarative language use for automatic generation of run-time infrastructure with a range of non-functional requirements placed on Internet of Things environments – notably Energy and Bandwidth – will be discussed in the talk. This unique perspective allows domain experts to distribute computation over IoT devices, field gateways, and clouds with little programming knowledge. An overview of the proposed PATH2iot open source platform will be presented – walking the audience through optimisation techniques to make offloading decisions taking into consideration the current state of infrastructure, description of computation and set of non-functional requirements, as well as the automated deployment process. Significant improvements to deployment plans for two real-world use cases have been achieved: a healthcare application resulted in substantial battery savings of a wearable device; and a smart city use case designed for real-time audio signal analytics to detect train arrivals, where the optimised deployment plan satisfied strict bandwidth constraints imposed by LoRaWAN technology.

Talks

Data-Intensive Science and Software

11:00 AM

OpenShift-based High-Performance Computing for Research in Astrophysics

Filip Hubík
Red Hat

Biography

Filip is an OpenStack Quality Engineer interested in real-world practical applications of cloud computing software, High-Performance Computing, data processing and analysis. He works mainly downstream on product failure analysis, debugging, automation and CI. He has a background in cloud computing, studied Information Technology at Brno University of Technology. Filip has been with Red Hat since 2014.

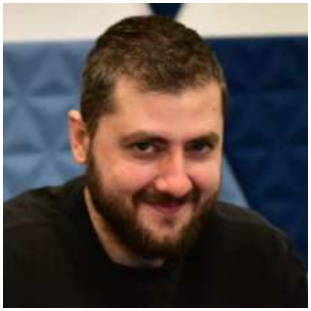

Nikolaos Moraitis
Red Hat

Biography

Nikolaos is a Software Engineer working on the Red Hat OpenShift team with a primary focus on the developer productivity test platform infrastructure. He is an active open source contributor in a variety of projects but most importantly in the Kubernetes community. Nikolaos grew up in Hellas, Greece where he did his studies as well.

Data-Intensive Science and Software

Gabriel Szász
Masaryk University

Biography

Gabriel is a graduate student at Masaryk University in Brno studying the effects of rotation on measured properties of stars – with a primary focus on the stars in the Milky Way Galaxy and Magellanic Clouds.
Gabriel has a long-term academic background in Astronomy and Astrophysics. He has been working at Red Hat for almost eight years.

Zdeněk Švécar
Red Hat

Biography

Zdenek is an IT Project Manager at Red Hat working on spearheading deployment of regional office infrastructure and services. Most recently he's been instrumental in the successful completion of all aspects of IT systems for the new Red Hat offices in Paris and Madrid. He has an academic background in Electrical Engineering and Informatics and has been with Red Hat since 2009.

Abstract

Our curiosity is constantly pushing technology towards higher goals, unveiling the secrets of distant realms far beyond our imagination.
Masaryk University, in collaboration with the Brno Red Hat office, initiated an ambitious project that has the potential to redefine what we have known about the stars in our Local Universe. Powerful open source tools, especially the Red Hat OpenShift Container Platform, are making this extensive project feasible.

Talks

Data-Intensive Science and Software

In this talk you will hear from four presenters: a Project Manager, a Graduate Student in Astrophysics, and two Red Hat Developers, all working together on a new approach involving stellar rotation as an essential parameter while focusing on reproducibility, parallelization and simplicity.

11:30 AM

Avoiding Bad Decisions and Heuristics

Ulrich Drepper
Red Hat

Biography

Ulrich Drepper returned to Red Hat in 2017, after a seven-year hiatus when he worked for Goldmann Sachs. He is part of the Office of the CTO and concentrates on developing new technologies for high-performance computing(and machine learning specifically), mostly in collaboration with university groups.
In his last position at Goldman Sachs he focused on the development of models and various types of stochastic algorithms to aid in operation of the technology for the entire firm. He also taught various internal classes around machine learning and other computing topics. Additionally, Ulrich was an internal consultant for all aspects related to performance, low latency, and C/C++ compilers.
His previous stint at Red Hat lasted 14 years. The last position was as member of the Office of the CTO to collect and disseminate information relevant to the Red Hat Enterprise Linux product, predominantly in the high-performance area.
His main interests are in the areas of low-level technologies like machine and processor architectures, programming language, 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.

He worked on several revisions of the POSIX standard and was invited to serve as an expert on both the C and C++ standards committees. Ulrich received his Diploma in Informatics from the University of Karlsruhe, Germany.

Abstract

In many software projects and especially in system software programmers have incomplete knowledge about the actual use cases. To provide the needed flexibility, performance, etc. these software projects often feature many, too many, configuration options. Some are exposed to the user, some internally controlled by heuristics. Unless both the developer and the user spend a huge amount of time tuning, the resulting system will not run optimally. Additionally, settings are then often applied statically while the environment constantly changes due to new software version, new hardware, or changing workloads.

One way out is to have the system learn about the best settings. If this happens for each deployment the solution could be locally optimal and adding more flexibility in the implementation can lead to better performance and not just higher complexity.

The compiler is a piece of system software that has several heuristics built in.

In this talk Uli will go over the problem statement and using the example of the compiler describe ways to achieve self-tuning system using machine learning. These are not solved problems and a goal of the talk is to get researchers interested in tackling these problems together with Red Hat.

Talks

Security and Privacy

1:00 PM

Observing Developers Interacting with TLS Certificates

Martin Ukrop
Masaryk University

Biography

A Ph.D. candidate at the Centre for Research on Cryptography and Security at Masaryk University focusing on usable security from the developers' viewpoint. He is engaged in the teacher-training initiative at the faculty, passionate about experiential learning and actively organizing educational events in the community Instruktoři Brno. "Ceaselessly fascinated by the world."

Abstract

Over the past three years, we have conducted several experiments observing IT professionals interacting with TLS certificates. The two most prominent of these experiments took place at DevConf.CZ 2017 and DevConf.CZ 2018. We investigated the usability of OpenSSL (as of 2019, the most common library for manipulating X.509 certificates), the developers' understanding of multiple certificate issues and the trust they have to such flawed certificates. Conclusions based on these observations can help us produce an environment that is more usable for IT professionals resulting in fewer vulnerabilities in developed products.

1:30 PM

Analyzing, Breaking and Improving Certified Cryptographic Hardware (TPMs and Smartcards)

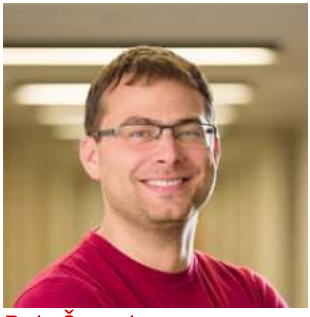

Petr Švenda
Masaryk University

Biography

Petr is an Associate Professor at the Masaryk University, Czech Republic. He engages in the research on randomness and pseudo-randomness and key distribution protocols usable for systems with multiple parties, often with devices significantly limited in performance capabilities like, e.g., cryptographic smart cards. He also focuses on utilization of secure hardware in complex scenarios and the development of secure applications on such platforms. He dreams about a more open and transparent world of cryptographic smartcards.

Abstract

Cryptographic hardware like smartcards or Trusted Platform Modules (TPMs) is a crucial component of many security systems, serving as an authentication token, digital signature device, secure storage for encryption keys, or providing a platform's root of trust. Despite the existence of extensive security certification schemes like Common Criteria or NIST FIPS140-2, the security vulnerabilities are still found in such devices, partially due to the overall closeness of the secure hardware ecosystem.

The talk will present a suite of the open tools for black-box security analysis of cryptographic hardware developed by CRoCS laboratory at Masaryk University, and vulnerabilities found like ROCA (CVE-2017-15361) or Minerva (CVE-2019-15809) which lead to large practical impact with estimated 1-2 billion devices affected worldwide.

Security and Privacy

As the analysis is performed without the need for knowledge of the hardware design or firmware source code, it can be applied not only by specialized certification laboratories but also by the end-users of these devices – all with the goal of better and continuous security analysis and more transparent certification.

2:00 PM

Improving Disk Encryption in Linux

Milan Brož
Masaryk University, Red Hat

Biography

Milan Brož is a developer and maintainer of Linux disk encryption open source code and a researcher in the area of storage security. He received his Ph.D. from the Masaryk University in 2019. With over 20 years of industrial experience as a software engineer, he is now focusing mainly on storage and security systems design. He is also working as a program manager for the Red Hat Research Program and its university and research projects.

Abstract

Disk Encryption has become a widely used security feature used in all types of data processing systems today – from mobile devices up to large cloud systems. This talk is an overview of several recent improvements in Linux disk encryption like data integrity protection based on authenticated encryption or using memory-hard key derivation for protection of disk unlocking passphrases. It demonstrates that an applied research collaboration between a university research laboratory and upstream developers can produce not only academic publications but also new features in open source projects.

2:30 PM

Leaky Processors: Lessons from Spectre, Meltdown, and Foreshadow

Jo Van Bulck
The Katholieke Universiteit Leuven

Biography

Jo Van Bulck is a PhD student at imec-DistriNet, KU Leuven (BE). His research explores security limitations along with the hardware-software interface, with particular attention to privileged side-channel attacks in trusted execution environments. Over the past years, Jo has uncovered several innovative microarchitectural side-channel attack vectors in commodity Intel x86 processors, and more recently was among the first to discover transient-execution CPU vulnerabilities. In the aftermath of Spectre and Meltdown, his research on the high-profile Foreshadow attack led to a complete collapse of the intel SGX ecosystem and ultimately even dismantled widespread virtual machine and operating system isolation.

Daniel Gruss
Graz University of Technology

Biography

Daniel Gruss is an Assistant Professor at Graz University of Technology. He finished his PhD with distinction in less than 3 years. He has been involved in teaching operating system undergraduate courses since 2010. Daniel's research focuses on software-based side-channel attacks that exploit timing differences in hardware and operating systems. He implemented the first remote fault attack running in a website, known as Rowhammer.js. He frequently speaks at top international venues, such as Black Hat, Usenix Security, IEEE S&P, ACM CCS, Chaos Communication Congress, and others. His research team was one of the teams that found the Meltdown and Spectre bugs published in early 2018.

Talks

Security and Privacy

Abstract

Over the past decades, security has been largely regarded as a software developer's responsibility, while hardware vendors have focused on making processors faster. However, now that security researchers are exploring ever-deeper parts of the system stack, it has been shown that these impressive performance improvements also come with a cost. With the announcement of the Spectre, Meltdown, and Foreshadow CPU vulnerabilities in 2018, an entirely new and dangerous class of transient execution attacks has arisen. Crucially, these attacks cross the hardware-software boundary, thereby challenging decades of performance gains and ultimately eradicating the viewpoint of security as an exclusive software responsibility. This talk will review transient-execution processor vulnerabilities from the ground up. We will discuss Spectre, Meltdown, Foreshadow, and more recent variants. Finally, the talk will discuss the long-term insights behind this recent wave of attacks, as well as mitigation strategies across the application, compiler, operating system, and CPU microcode levels.

3:00 PM

Plundervolt: Pillaging and Plundering SGX with Software-based Fault Injection Attacks

Kit Murdock
The University of Birmingham

Biography

Kit is currently pursuing a PhD in cyber security at The University of Birmingham. Her research interests include embedded hardware and software based fault injections. Kit has been building and researching a tool to enable testing and evaluation of hardware fault injection using software emulation. Kit currently runs the University's Ethical Hacking Club, AFNOM which encourages students to learn offensive security in a friendly, informal environment.

Abstract

Many modern processors expose privileged software interfaces to dynamically modify the frequency and voltage. These interfaces were introduced to cope with the ever-growing power consumption of modern computers. In this talk we show how these privileged interfaces can be exploited to undermine the system's security. We present the Plundervolt attack – demonstrating how we can corrupt the integrity of Intel SGX computations.

4:00 PM

Formal Verification of a Linux Distribution

Kamil Dudka
Red Hat

Biography

Kamil Dudka joined Red Hat as an intern in 2008 when he was finishing his Master's degree in Intelligent Systems at Brno University of Technology. In 2009 he started to work on a formal verification tool named Predator, which has won several gold medals in the International Competition on Software Verification (SV-COMP). Since 2011 Kamil develops open source tools for fully automatic static analysis of RPM packages.

Abstract

Red Hat uses static analyzers to automatically find bugs in the source code of Red Hat Enterprise Linux, consisting of approx. 3000 RPM packages and 300 million lines of code. There are open source tools that can statically analyze this amount of software in a fully automatic way. Would it be possible to use formal verification tools to find bugs in (or even prove the correctness of) the important pieces of code in our Linux distribution? Red Hat is now experimenting with formal verifiers Symbiotic and Divine, which are developed by research groups of Masaryk University in Brno.
Are these tools ready for industrial software? How much are we able to integrate them into our release pipeline?

Talks

Code Analysis and Verification

4:15 PM

Symbiotic: Program-Slicing Framework

Marek Chalupa
Masaryk University

Biography

Marek Chalupa got his Bachelor's and Master's degree in computer science from the Faculty of Informatics of Masaryk University in Brno. During his undergraduate studies, he also worked for Red Hat on the desktop team. Currently, he is a PhD candidate at the Faculty of informatics of Masaryk University and his research focuses on program analysis and verification.

Abstract

Symbiotic is a framework that takes a C or LLVM program and generates a new program suited to analysis and verification. To reach its aims, Symbiotic uses instrumentation combined with static analyses, program slicing and compiler optimizations. It can also seamlessly run several program analysis tools on the generated program. In this talk, I will briefly describe how Symbiotic works and how it can be used.

4:30 PM

Heavy-Duty Program Analysis with DIVINE

Vladimír Štill
Masaryk University

Biography

Vladimír Štill is a PhD candidate at the Faculty of Informatics, Masaryk University in Brno. His research topic is the correctness analysis of parallel programs written in C and C++. His work includes finding new analysis techniques and their implementation in the DIVINE tool. His most notable contributions include support for discovery of bugs related to relaxed memory behaviour on modern processors, and work on detection of parts of parallel programs which do not terminate. Apart from his PhD research, he is also involved in teaching several courses on the faculty, mostly in programming and formal languages. He also maintains an evaluation engine for programming courses.

Abstract

DIVINE is a tool for analysis of programs written in C and C++ that primarily targets hard-to-find bugs, including concurrency-related bugs, assertion violations, and memory errors. It is developed in the Parallel and Distributed Systems Laboratory at the Faculty of Informatics, Masaryk University. One of the main strengths of DIVINE is that it is precise and can discover bugs which are hard or impossible to find by other means. Unsurprisingly, this precision comes at some cost in the human and computing effort put to the analysis. In this presentation, we outline some of the strengths and weaknesses of DIVINE. We describe how we can use DIVINE for the analysis of programs which have source code and tests available, and what can we do if that is not the case. We also show how DIVINE can be used with programs that run in the POSIX environment, i.e., use system calls and interact with other programs using the filesystem or network. Finally, we present how we can apply DIVINE to the discovery of sporadically occurring concurrency bugs.

Talks

Code Analysis and Verification

4:45 PM

Run-time Verification Using Testos

Aleš Smrčka
Brno University of Technology

Biography

Aleš Smrčka is an Assistant Professor at Brno University of Technology. He received his PhD on the topic of formal verification of hardware designs in which he continued on research of formal verification of pipelined processors. He pushed his experience in formal methods to the area of software testing, where making quality software is his main point of interest. He contributes to different topics of software testing and dynamic analysis ranging from automated test design, testing of concurrent programs and run-time verification, to automated generation of complex test inputs. He teaches programming and software testing and mentors a team of software testing enthusiasts which develop different testing tools emphasising practical solutions to real problems from industry area.

Abstract

Proof of formal correctness can be provided by formal verification. Unfortunatelly, exhaustive verification does not scale well when applied on some programs incorporating features like complex data structures, synchronisation primitives, floating-point computations, or the combinations of above. These are the cases where runtime verification takes its place. Runtime verification does not statically reason about the program, it instead traces a program under test and monitors the program correctness. Such an approach can be easily combined with systematic or random testing. Correct behaviour is typically expressed in some kind of finite automata or temporal logic well known in other formal methods. In this talk, runtime verification will be briefly introduced and Testos, a tool set developed in BUT, will be used as a practical example of how to monitor ptLTL formulae during testing of C programs' concurrency bugs.

5:00 PM

Performance Versioning System

Tomáš Fiedor
Brno University of Technology

Biography

Tomas is a researcher of the VeriFIT group (BUT FIT). His main focus is on static and dynamic performance analysis. In particular, he develops novel techniques for efficient performance measurement; effective interpretation and modeling; and adapting some recent techniques, such as fuzzing, for performance analysis. He is the main author of the Perun tool.

Abstract

Long-term management of the performance or the resource consumption of any programming project is, indeed, a tiresome job. As a consequence, many projects are needlesly slow and deplete a computer's precious resources. Yet, so far, there has been no effective solution! This talk introduces Perun: a tool suite, which aims at helping developers track the performance history of their projects. Perun serves as a wrapper over existing repositories and offers a suite of performance collectors and analysers. The talk will provide a brief overview of Perun's main features and introduce selected novel techniques for, e.g. automatic detection of performance changes, generating inputs leading to bad performance, or effective modeling of the program performance.

Talks

Code Analysis and Verification

5:30 PM

Efficient Run-time Verification for the Linux Kernel

Daniel Bristot de Oliveira
Scuola Superiore Sant'Anna
Red Hat

Biography

Daniel is a Principal Software Engineer at Red Hat, working in the real-time team. He researches the run-time behavior of Linux, both from the logical perspective, with run-time verification and timing perspective, with real-time analysis. He is a member of the Real-time System lab at Scuola Superiore Sant'Anna, Pisa – Italy.

Abstract

Formal verification of the Linux kernel has been receiving increasing attention in recent years, with the development of many models, from memory subsystems to the synchronization primitives of the real-time kernel. The effort in developing formal verification methods is justified considering the large code-base, the complexity in synchronization required in a monolithic kernel and the support for multiple architectures, along with the usage of Linux on critical systems, from high-frequency trading to self-driven cars. Despite recent developments in the area, none of the proposed approaches are suitable and flexible enough to be applied in an efficient way to a running kernel. Aiming to fill such a gap, this research proposes a formal verification approach for the Linux kernel, based on automata models. It presents a method to auto generate verification code from an automaton, which can be integrated into a module and dynamically added into the kernel for efficient on-the-fly verification of the system, using in-kernel tracing features. Finally, a set of experiments demonstrates verification of three models, along with performance analysis of the impact of the verification, in terms of latency and throughput of the system, showing the efficiency of the approach.

Connect with Research

RED HAT RESEARCH QUARTERLY

Want to know more about Red Hat Research?
Check out our magazine, the *Red Hat Research Quarterly*.

Download the magazine, subscribe to the digital version or get the printed version delivered to your mailbox:
research.redhat.com/quarterly.

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 the Red Hat Research Program, email us at academic@redhat.com.

Red Hat
Research