Net-Centric & Cloud Software & Systems I/UCRC

Semi-Annual Highlights of Industry/Academic Cooperative Research Projects

http://netcentric.unt.edu

Issue 5, September 2016

An Application Benchmark and Simulation Approach for Cloud Computing Systems

Companies have a significant number of choices when it comes to cloud based application providers. An important selection criterion is whether the computing assets of the provider (servers, apps, networking infrastructure, bandwidth) can scale appropriately as the number of users and transactions increases over time. However, not all providers publish empirical data to support such an evaluation. Beyond this, there are myriad configurations to consider since cloud tenant requirements will be diverse.

Chief Information Officers (CIOs) are asking the question, “Should I commit my organization’s software applications to the cloud, and if so what would be the performance and cost?” The Silverlining research project provided twenty-one University of Texas at Dallas (UT Dallas) graduate students with a platform to answer this question, by benchmarking the performance and cost of applications in the Google GAE cloud.

The industry standard On Line Transaction Processing (OLTP) benchmarks and the follow-on cloud simulation forecaster were designed to guide industry CIOs, as well as system/software engineers, through the new cloud applications development and operations lifecycle. The UT Dallas Silverlining team extended OLTP benchmark to operate, for the first time, over the internet into the Google App Engine (GAE) cloud using two very different database engines (CloudSQL and Datastore/NoSQL).

UTD’s goal-oriented simulation approach for cloud-based system design captures stakeholder goals, together with such domain characteristics as workflows, and uses them to create a simulation model as a proxy for the cloud-based system architecture. Simulations are then run, in an interleaving manner, against various configurations of the model as a way of rationally exploring, evaluating and selecting among incrementally better architectural alternatives.

Using this approach, evaluators of cloud service providers can objectify multiple offerings with empirical performance data to guide selection of a provider that best meets their organizational requirements. ■

UNT’s Krishna Kavi Awarded Patent

Dr. Krishna Kavi, the Director of the NCSS IUCRC was granted a new patent in July 2016, titled, “Method and apparatus for improving computer cache performance and for protecting memory systems against some side channel attacks”. The patent was applied for in 2012. The basic innovation stems from how cache memories are addressed. In traditional designs, cache indexing can be viewed as a modulo based hash function, where a given address is “hashed” into one of the cache areas (known as cache sets). The innovation proposed by Kavi changes the hash function so that different address bits, which can be randomized, are used to create the modulo function, and the function can be changed dynamically to make it very difficult to predict which address maps to which cache set. The ability to predict addresses mapping to cache sets have led to some security attacks, leading to the disclosure of cryptographic keys.

This approach can also be used to mitigate cache conflicts among different applications or data sets, by mapping data of different applications to different cache sets. The proposed addressing can be applied to any type of cache (L-1, L-2 or Last Level Cache). ■

Frankenstein Malware for Defensive Testing

Advanced, polymorphic, zero-day threats are one of the most dangerous and increasingly common threats in the malware landscape. Anti-malware defenders are therefore devoting considerable resources to the development of more powerful static malware detection algorithms that can counter these threats. However, since tomorrow's polymorphic malware is usually unknown to defenders, they must usually resort to testing newly discovered algorithms only on known malware families that employ well-understood mutation strategies.

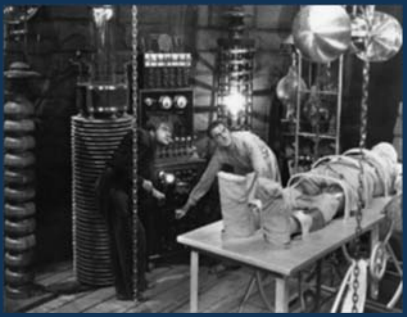

Researchers at UTD are looking to develop more challenging data sets consisting of malware that uses previously unseen mutation strategies to create more robust defenses that are prepared for tomorrow's malware attacks. UTD has developed "Frankenstein" - a new malware mutation paradigm that stitches together code fragments pilfered from other programs to recreate itself in entirely new ways. This potentially dangerous new technique yields malware mutations that are often statistically indistinguishable from benign software, and are therefore very challenging to detect using conventional matching algorithms.

The research team is extending Frankenstein with support for more realistic malware payloads, such as those that exploit zero-days. The goal is to provide defenders with new data sets that better reflect the challenges posed by future malware threats, and thereby facilitate the creation and evaluation of better, more robust defenses.

To date, UTD has repurposed the Frankenstein algorithm for automated software defense hardening and extended it for defense testing against implementation-aware code-reuse attacks. The research team has also done some preliminary work toward simplifying, generalizing, and extending their payload specification language. ■

Differentiated Levels of Security for IoT Devices

The Internet of Things (IoT) is quickly becoming widely accepted as a standard for smart device communication. Experts estimate that about 50 billion devices will be connected by 2020.

IoT devices are entering every aspect of society including self-driving cars (or cars with safety and driver assistance technologies), smart homes, and smart cities. The one requiring immediate attention is security. An IoT environment consists of devices communicating over a multitude of network protocols. The trouble with uniform security policies throughout the network is the lack of flexibility and control. However, device security should be rooted on the device itself so that different devices will be protected at different levels (such as multiple levels of authentication, different levels of encryption, firewalls etc.). It is also essential that the security approach should not deter users from relying on specific security features because they require a high level of expertise or are too cumbersome (such as requiring complex passwords). UNT's proposed IoT Security Hub aims to fill this gap by securing connected homes and businesses from malware and hackers.

UNT has designed a novel solution to address these issues. The IoT security hub consists of a single physical device that hosts a secure trusted environment to which the IoT devices connect to communicate with each other Machine-to-Machine (M2M) or to access the internet. The first level of defense in the trusted environment is a set of IDS (Intrusion Detection Systems) and a firewall service running at the point of interaction between the trusted environment and the Internet. Linux containers, each containing a pre-configured snapshot of the security policies for each class of devices, need to be invoked. Once the device has been identified and grouped, the Software Defined Network (SDN) controller then invokes a daemon process that runs in the background. This daemon process is tasked with invoking one of the many device-class specific security function containers that are either suspended or powered down (so in essence, there are as many containers as there are classes of devices). Hardware integrity is ensured by RADIUM (Race-free On-demand Integrity Measurement Architecture) which provides integrity measurements at regular intervals. RADIUM, along with Intel SGX ensures complete trustworthiness of the underlying hardware. The hypervisor then takes control of the platform thereby providing a trusted environment for applications. ■

A Novel IoT to Improve Home Security

People are being exposed to more and more interlopers when they're at home. People come to the door or call on the phone trying to sell various products and services, solicit donations to charity, participate in surveys, and the list goes on. A single home might receive 20 or more phone calls every week (both bots and humans) that are not identifiable on caller ID or with a voice message. The ones that come to the door might leave a card or just move to the next house hoping to confront a real person.

The good news is that the phone irritations are fairly easy to fix. For about $120 one can get an answering machine with 4 handsets that lets you save and block more than 250 numbers. This may sound like a lot of storage, but it may only be marginal for some. The same solicitor can originate from multiple numbers. If he repeatedly hits a disconnect using one number, he starts using the next available. It's not uncommon to see as many as 6 numbers used by the same solicitor (first 7 -8 digits the same, last 2-3 vary), so more memory may be important. Such a system is a viable solution to phone pests.

For potential front door annoyances, an IoT device called the Ring Video Doorbell (see it at Ring.com) is available. This device connects to a home's local network and transmits live audio and HD video to a PC or smart phone. Users can see who's at the door and talk to them in real time. Whether a user is away from home or just too lazy to get up from their favorite chair, one can simply grab their smartphone and use the free Ring app to see who's at the door and talk to them as through an intercom. It's a super convenient solution for dealing with people at the door. Almost.

A very useful feature of the Ring doorbell is being able to look back at videos of people that came to the door throughout the day, but couldn't be answered at the actual time of arrival. In many cases the user only wants to know that a package was delivered by the mailman, UPS, FedEx, etc. In other cases one may want to know that a neighbor stopped by for some reason (maybe to borrow an orbital sander or extend an invite to a weekend soirée). But herein lies the rub...

Advertising for the Ring doorbell seemed to suggest that missed rings would be saved to the user's smartphone for later recall. In fact, missed rings are sent to the cloud operated by Ring.com for which they charge $30 per year (a novel income stream). However, to be fair, one can use the Ring doorbell app to intercept rings as they happen, but you have only about 45 seconds to respond. Failing that, there will be no record of who rang unless you opt for the fee-based cloud archiving.

Being able to go back in time to see who was at the door makes the Ring a handy IoT device. The $30 annual fee for cloud storage may be an irritation for some, but not a show-stopper. Overall it's a genuinely useful device that gives users real-time situational awareness through a smartphone of who's at the door. ■

UNT Summer Research Experience for Undergraduates

The Computer Systems Research Laboratory (CSRL) at the University of North Texas sought to provide five undergraduate research students opportunities to work with graduate students this summer.

For Troy King, a Senior Computer Science major at UNT, working on Dataflow PIM taught him about how to research in general and more about architecture, which had always been a difficult topic for him in the classroom.

Zach Poycattle, a Senior CS major at UNT, focused his research on looking at vulnerabilities in software and hardware to see what types of attacks were possible because of them. After gaining that knowledge, he put everything he learned into an ontology using a program called Protégé. In Protégé, he had to use logic rules make the ontology more efficient and to make sure everything was in the right place.

This summer was the first time Mukundan Kuthalam, a Junior Computer Engineering major at UT Austin, was involved in research. He joined the CSRL lab to get more experience in what grad school would be like and to learn more about the field of software. His research on identifying and analyzing cloud application vulnerabilities really opened his eyes to how vast and complex the field of software engineering is.

From left to right: Troy King, Zacharia Poycattle, Mukundan Kuthalam, and Clement Cole.
Not pictured: Margarita Sanchez

The main task in Clement Cole's research was to build an ALU (Arithmetic Logic Unit) for testing data flow architecture in support of another research project. Clement is a Senior CE major at UNT.

Margarita Sanchez (not pictured), a Senior CS major at UNT, spent this summer creating a program that works with the traces generated by the Hopper Application. The application takes an executable, reverse compiles it, and provides her with a trace file that she then uses to generate data flow graphs (DFG). The program takes in this trace and generates the DFG and allows her to compare it with other graphs to see if there are any similarities that she can use to improve performance.

The REU students describe their experience working with Dr. Kavi and the CSRL graduate students as rewarding, challenging, and inspiring. ■

Collaborative Industry-University Workshop of the ASU NCSS Site and Tech de Monterrey Telecommunications Center

The ASU NCSS site and the SenSIP Center co-organized a collaborative workshop with the Tecnológico de Monterrey (ITESM) Center for Electronics and Telecommunications on sensors and machine learning in May 2016. The International Sensors, Signal Processing and Communication Workshop brought together government, industry and academic leaders from the United States and Mexico to share new signal/information processing research and identify collaborations that can meet the challenges of an increasingly digital world. The workshop included presentations from ITESM, ASU, NXP, Intel, MEMS Industry and Sensors Group, IBM Research, and Lawrence Livermore.

The Plenary speaker, Stephen Whalley, CSO of MIGS, talked on "Forecasts for sensor demand are as high as 100 trillion by 2030." Stretchable plastic and thin film substrates and ultimately roll-to-roll printing on paper will serve various application, cost, performance and form factor needs. Talks by IBM research and Lawrence Livermore stressed the Importance of machine learning in emerging internet of things applications. A talk on deep learning by IBM research attracted a lot of interest.

According to NCSS Site director, Andreas Spanias, ASU faculty are helping the ITESM Center for Electronics and Communications to establish an industry consortium in telecommunications and sensors. Drs. Spanias and Vargas have since jointly proposed projects to the Consejo Nacional de Ciencia y Tecnología, known as Conacyt and to Samsung. The projects will provide scholarships for postgraduate studies and managing programs to encourage industry and private sector involvement in science and technology research and development. ■

The collaborative ASU-ITESM workshop, May 2016

Use of Machine Learning for Condition Monitoring

The ability to perform reliable hardware prognostics for mission and safety critical systems is essential for preventing death or serious injury to people and loss or severe damage to equipment and property. Effective implementation of prognostics depends largely on the choice of appropriate sensing devices, fusion of information from multiple sensors and trending of data over time for more precise prediction of eventual component failures.

A research team at ASU is investigating how to optimally develop condition and health monitoring applications using a common set of machine learning primitives (a *toolbox*) that can be used to develop a broad range of prognostics applications. They are also developing a hardware feature and sensor selection guidelines to ease the cost of entry for condition monitoring applications.

Using devices and test equipment provided by their project sponsor NXP (formerly Freescale), ASU researchers have collected data for DC motor run-to-failure testing and various water pump conditions. The team also conducted experiments using a multi-function sensor incorporating accelerometer, magnetometer, gyroscope, and pressure measurement with a focus on minimally supervised or unsupervised sensor data analytics.

The team also implemented an embedded machine learning algorithm based on the Gaussian Mixture Model in an NXP embedded sensor board. Data collection and statistical analysis were performed as embedded applications on the board. Density functions were estimated from normal conditions to provide criteria for detecting abnormal behaviors.

In the next phase of this project, an embedded machine learning sensor system will be developed for monitoring hardware conditions and providing libraries that can be reused in Cloud-based networks. ■

Dr. Guang Zhou was very active in our NSF NCSS I/UCRC at the University of Texas at Dallas. He worked on the Internet of Things project, especially on the Autonomous Agriculture Robotic system project, called the AgriBot project. He worked with other students, including two high school students, and actually developed and deployed a prototype of the AgriBot system and tested it on a farm.

As part of his Doctoral research, Guang also worked on the development, deployment, and evaluation of a very high quality Automated UAV based Power Line Inspection system. The goal of this system is to make power line inspections safer and more efficient as compared with manual inspection procedure. The system collects images from the camera in the drone and performs real-time detection and recognition using advanced image processing algorithms to detect abnormal conditions, such as wires that may brake soon. The system provides a simple, low-cost and easy-to-use alternative to existing power line inspection methods. Based on this research, he received the 1st Place Award in the DJI Developer Challenge competition. Guang graduated in Summer 2016 and is working at Baidu USA (San Francisco) as a Research Scientist in the Baidu Autonomous Driving Group.

Another UTD student, Dr. Jinwei Yuan, worked on the project on robust appearance model construction for a long-term tracking system. In addition to this project, Jinwei did very innovative research on image processing as part of his Doctoral research, especially for real-time object tracking. Jinwei received the Top 10% Award at the *2014 IEEE International Conference on Image Processing*. Jinwei graduated in Spring 2016 and, based on his innovative research, he is currently working as a Software Engineer in the Google Image Processing Research Team.

Recent Publications on NCSS Related Research

C.Y. Lee and K. Kavi, "Evaluation of Security Service Level Agreements", 10th International Conference on Software Engineering Advances (ICSEA-2015), Nov 2015.

C. Shelor, K. Kavi, S. Adavally, "Dataflow based near-data processing using coarse grained reconfigurable logic", 3rd Workshop on Near Data Processing (WoNDP-3) in conjunction with the 48th IEEE/ACM International Symposium on Microarchitecture (MICRO-48), Dec 2015.

Book Monograph, S. Miller, X. Zhang, A. Spanias, Multipath Effects in GPS Receivers, Synthesis Lectures on Communications, Morgan & Claypool Publishers, ISBN 978-1627059312, 70 pages, Ed. William Tranter, Vol. 8, No. 1 , Pages 1-70, Dec 2015.

V. Berisha, A. Wisler, A. Hero, A. Spanias, "Empirically Estimable Classification Bounds Based on a Nonparametric Divergence Measure," *IEEE Transactions on Signal Processing,* Vol. 64, No. 3, pp.580-591, Feb 2016.

H. Braun, S. T. Buddha, V. Krishnan, C. Tepedelenlioglu, A. Spanias, M. Banavar, and D. Srinivansan, "Topology reconfiguration for optimization of photovoltaic array output," *Elsevier Sustainable Energy, Grids and Networks* (*SEGAN),* pp. 58-69, Vol. 6, Jun 2016.

P. Kamongi, K. Kavi, M. Gomathisankaran, "Predicting Unknown Vulnerabilities using Software Metrics and Maturity Models", 11th International Conference on Software Engineering Advances (ICSEA-2016), Aug 2016.

G.R. Lundquist, V. Mohan, and K.W. Hamlen, "Searching for Software Diversity: Attaining Artificial Diversity through Program Synthesis," Proceedings of the New Security Paradigms Workshop (NSPW), Sept 2016.

Upcoming Events

The semi-annual meeting of the NCSS Industrial Advisory Board will be held Wednesday and Thursday, October 19-20, 2016 at the University of North Texas.

For more information about the NCSS I/UCRC or how you can join our center, please contact:

Professor Krishna Kavi
Director of NCSS I/UCRC
Krishna.Kavi@unt.edu

or

Melanie Dewey
940-565-4764
Melanie.Dewey@unt.edu

University Membership

University of North Texas (UNT)
Krishna Kavi, Site Director

University of Texas at Dallas (UTD)
Farokh Bastani, Site Director

Arizona State University (ASU)
Andreas Spanias, Site Director

Industry Membership

AMD
Armor
Ashum Corp.
Briggs Freeman Sotheby's International Realty
CompuMatrice
Endometric
Interactive Flow Studies
Intel
NTT Data
NXP (formerly Freescale)
Poundra
Raytheon
Sprint
Texas Instruments
US Army RDECOM CERDEC