REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggesstions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any oenalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
23-02-2016	Final Report	3-Apr-2008 - 2-Jul-2017

4. TITLE AND SUBTITLE
Final Report: Secure Open Systems Institute

5a. CONTRACT NUMBER
W911NF-08-1-0105

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

6. AUTHORS
Laurie Williams

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
North Carolina State University
2701 Sullivan Drive
Admin Srvcs III, Box 7514
Raleigh, NC 27695 -7514

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
54270-NS.167

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
The Secure Open Systems Initiative (SOSI) at North Carolina State University was initiated by a federal appropriation and subsequent award through the US Army Research Office (ARO) on April 3, 2008. During the reporting period (08/01/2011 – 11/31/2014), SOSI supported 27 seed projects, with a research focus on cloud computing security, software security, wireless security, and the science of security. SOSI has also been building up a virtual computing testbed called VCL Pre-production Testbed, partially with support from and DURIP grant. SOSI also partially supported the following projects: Centmesh (a reconfigurable wireless mesh network to address

15. SUBJECT TERMS
open source, security, open systems

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dennis Kekas
UU	UU	UU	UU		19b. TELEPHONE NUMBER 919-515-5297

RPPR

as of 27-Jan-2021

Agency Code:

Proposal Number: **Agreement Number:**

Organization:
Address: , ,
Country:
DUNS Number: EIN:
Report Date: Date Received:
for Period Beginning and Ending
Title:
Begin Performance Period: **End Performance Period:**
Report Term: -
Submitted By: Email:
Phone:

Distribution Statement: -

STEM Degrees: **STEM Participants:**

Major Goals:

Accomplishments:

Training Opportunities:

Results Dissemination:

Plans Next Period:

Honors and Awards:

Protocol Activity Status:

Technology Transfer:

Final Report: Secure Open Systems Institute

The Secure Open Systems Initiative (SOSI) at North Carolina State University was initiated by a federal appropriation and subsequent award through the US Army Research Office (ARO) on April 3, 2008. During the reporting period, SOSI supported 27 seed projects, with a research focus on cloud computing security, software security, wireless security, and the science of security. SOSI has also been building up a virtual computing testbed called VCL Pre-production Testbed, partially with support from and DURIP grant. SOSI also partially supported the following projects: Centmesh (a reconfigurable wireless mesh network to address security and protocol concept investigations), KIMA (kernel security work focused work, that saw its designed implemented in an Android OS) & Eduport (work with adding security & provenance capabilities into NCSU's VCL cloud management platform to allow XaaS (anything as a service).

Table of Contents

Summary of the most important results 3

Bibliography 16

Scientific progress and accomplishments

Developing a User Profile to Predict Phishing Susceptibility and Security Technology Acceptance

Faculty: Dr. Emerson Murphy-Hill, Dr. Christopher Mayhorn

We have collected phishing susceptibility data from a subsample of NSA analysts/personnel (n=12). Subsequent data analysis is underway to assess whether the student sample data generalizes to other user populations. We have met with members of the NCSU IT group to assess the current state of anti-phishing training being employed for campus users. These discussions suggest that there may be a future opportunity to collaborate with this group and possibly engage in archival data analysis regarding past phishing instances at NCSU. We are designing a second study that will use participatory design with at-risk users to develop an anti-phishing tool. We have completed the design of the study as a 2 x 3 Mixed factorial design with the before/after (2) training manipulation as a within subjects variable and the training content (3) as a between subjects variable. We are currently evaluating the content of the training conditions.

Software Security Metrics

Faculty: Dr. Tao Xie, Dr. Laurie Williams

Soliciting and measuring security/privacy requirements. We explored leveraging the natural language requirement specification to come up with mock implementations to help solicit and measure security/privacy requirements in early stages (before the actual implementation is available). In particular, we investigated new requirements that are incremental feature requests on top of existing requirements and implementation to generate a mock implementation. The generated mock implementation could be used to solicit new security/privacy requirements (such as access control and privacy control on contents being displayed on screens) that may need developers' attention and measure the completeness of existing security/privacy requirements in the new requirements.

Leveraging the application descriptions in third-party mobile applications. We used our NLP infrastructure to detect security and privacy centric operations performed by a third-party mobile application. The work seeks to aid risk assessment of mobile applications by identifying whether permissions requested by a third-party mobile application are sensible based on the textual description of the application presented to the user. In particular, we incorporated techniques to account for semi-automated extraction of domain knowledge from API documents into our NLP framework.

Security policy metrics. We have collected three open source healthcare systems: OpenEMR, OpenMRS, and iTrust. These systems use security policies in program code. To find known security faults, we investigated bug repositories for each subject. Note that, among the three subjects, only OpenEMR and OpenMRS maintain bug repositories. We manually verified whether a bug report describes a security fault in security mechanisms (e.g., access control mechanisms). As a result, for OpenEMR and OpenMRS, we found 16 security faults (out of 366 reported bug reports) and 0 security fault (out of 21 reported bug reports), respectively.

To detect security faults, we use high-level security policy rules (e.g., a doctor role is permitted to modify patient record) extracted from either GUI, requirements, or code in a given system. By comparing such high-level security policy rules across different systems, we may find security faults in access control mechanisms in the systems. We developed a tool to detect differences of high-level security policy rules. To measure complexities of security policies written in the XACML policy specification language, we developed a tool to measure the numbers of attributes, conditional statements, and rules of policies. We measured the highest number of attributes and rules for the Continue-b project and the highest number of conditional statements for the JBoss project.

Leveraging the data in the National Vulnerability Database (NVD). We have been working on creating the NLP infrastructure for our analysis of natural language description on the national vulnerability database. In particular, the infrastructure is generic and based on an open source NLP parser. It incorporates techniques to specifically deal with the text interleaved with domain-specific keywords. Our initial results on applying our infrastructure on controlled description have been promising. In the future, we will apply and refine our infrastructure to more NVD entries. We are further working on developing mining framework to mine prevalent vulnerability configuration. We will further work to infer security metrics from the configuration perspective from the results of our natural language processing along with the configuration information available from each entry on the NVD database.

Empirical Privacy and Empirical Utility of Anonymized Data

Faculty: Dr. Ting Yu

We observe that different inference techniques for empirical privacy measure tend to offer comparable sensitive attribute prediction accuracy, which suggests the observed privacy/utility tradeoff using our framework reflects the intrinsic property of anonymized data.

Through synthetic data generation based on anonymized data and auxiliary information, our framework can be extended to improve the usability of anonymized data in two aspects: (1) it preserves the data format of the original data so that the data can be better analyzed through standard data analytic tools; (2) the integration of domain specific auxiliary information helps further improve the utility of anonymized data.

Improving the Usability of Security Requirements by Software Developers through Empirical Studies and Analysis

Faculty: Dr. Laurie Williams, Dr. Travis Breaux

A study of security expert aptitude for pattern selection and application is being conducted with security experts in network and system security to discover how security analysts solve security problems, systematically. This is based on the Situation Awareness method that frames expert thinking into three categories: state recognition, planning and projection. Our pilot study has been extended with new participants to evaluate surface vs. sub-structure in problem representation. We now have 79 complete participant responses. This study result will show how to frame security problems and solutions when evaluating pattern selection and application. In particular, we will learn whether the number of training examples or the variety of types of training examples are more likely to influence security analyst performance. By removing this confound, we can attribute any "effects" on the outcome (e.g., improved security designs) to the pattern template used as opposed to the effect of training.

We have completed a security requirements pattern specification template experimental protocol. This template is based on findings from our presentation at the RePa'12 workshop in Chicago, IL. The pattern combines the Inquiry-cycle (IC), a goal-oriented requirements refinement method, with the established software design pattern template. The technical challenge here is how to divide the problem space into cues and questions that can be used to lead novice analysts to expert conclusions. We plan to develop these patterns using common security problems, including access control, and to vet the patterns with security experts before conducting the experiment.

Attaining Least Privilege Through Automatic Partitioning of Hybrid Programs

Faculty: Dr. Will Enck, Dr. Helen Gu

This project investigates the hard problem of resilient architectures from the standpoint of enabling new potential for incorporating privilege separation into computing systems. However, privilege separation must be combined with appropriate security policy. The general hypothesis of this project is: legacy computing systems contain emergent properties that allow automatic software partitioning for privilege separation capable of supporting practical least privilege security policies.

- We have identified four natural boundaries as case studies to evaluate our hypothesis.
- We have used one of these natural boundaries to provide classified system call monitoring and demonstrated it effectively reduces false positives. This work is currently in submission to a major systems security venue.
- We have built a framework and designed an information flow policy language to control UI components in Android. We have shown that focusing on this natural boundary can effectively prevent accidental data disclosures. This work is currently in submission to a major systems security venue.

- We have identified another natural boundary in a common type of smartphone applications and leveraged the boundary to reduce third-party privileged code. This work is currently in submission to a major systems security venue.
- We have identified another natural boundary in the system installation of Android smartphones and are currently leveraging the boundary to provide hardware-based firmware integrity protection.

Shared Perceptual Visualizations for System Security

Faculty: Dr. Christopher Healey

Our approach to constructing flexible guidelines for evaluating visualization designs suggests that basic, low-level experiments can yield results that are useful for both theoretical and practical purposes. This suggests that, in situations where it is appropriate, a controlled experimental approach can yield successful results.

- Completed experiments to measure visual interference and visual distraction. We now have perceptual guidelines on how to use luminance, hue, size, and orientation to visualize large, complex datasets, or to evaluate a visualization design that uses these features.
- Integrated our perceptual guidelines into the construction of basic charts (line, bar, scatterplot, and so on) and node-link graphs. This allows us to present information using a well recognized format, and in ways that are perceptually optimal.
- Designed and implemented a web-based visualization tool, based partially on our experimental results, for presentation at an NSA-sponsored C3E security workshop in September (see http://www.csc. ncsu.edu/faculty/healey/c3e/query.html for a live demonstration.).
- Designed and implemented a web-based visualization tool, based partially on our experimental results, for line-graph and scatterplot presentation of real-time text sentiment in tweets from Twitter (see http: //www.csc.ncsu.edu/faculty/healey/debate/client for a live demonstration).
- Designed and implemented a web-based visualization tool, based partially on our experiment results, to construct and present relationship graphs of online and social network discussions of system security attacks. A prototype of this system formed the basis of a recently funded SBIR proposal on this topic.

Quantifying Underpinnings for Network Analytics as Components of Composable Security

Faculty: Dr. Rudra Dutta

In this project, our focus is on understanding a specific sensor and ad-hoc network security mechanism in analytical terms. The system in question is one of puzzles and responses. Compromised components consist of nodes that have been compromised, and are executing a Sybil attack on consensus-based network mechanisms, such as routing or distributed sensing. The project goals are to analyze the effectiveness of the Sybil detection (which in turn would determine the ability of the consensus-based mechanism to perform correctly in the face of such an attack.

In our initial phase, we made preliminary progress by identifying the following as relevant to quantifying Sybil detection probabilities.

- Sybil computation power relative to slowest genuine node in the network
- Introduction of Sybil node versus compromise of genuine node to act as Sybil
- Existence and reachability of base station as information fusing location
- A priori alerts versus “suspected set” challenge
- Distribution of complexity of puzzles assigned to masquerade identities

- Degree of proactiveness of Sybil attacker

We have added the following to the list of issues we recognize as being of relevance (i.e. affecting detection probability)

- Whether unique a priori node identities are available, or any node can join the network
- Whether puzzles generated offline are used or puzzles must be generated by verifier
- How many nodes the Sybil attempts to masquerade

We have currently focused on bounding detection probability for the case characterized by comparatively simpler choices of assumptions:

- A Sybil model that solves all puzzles at highest possible speed of the Sybil, and reponds to each outstanding unsolved puzzle as soon as the solution is available. Such a naïve model of the Sybil is appropriate to start with since more sophisticated Sybils can only further reduce detection probability, and also such a model is not unlikely when first considering detection (as opposed to no attempt at detection).
- Verifier issues puzzles to all neighbors at the same time, but does not attempt to create any particular puzzle complexity distribution.
- Unique node identities are available to nodes (may be stolen by malefactors)
- No "suspected set" already exists for a verifier
- A base station exists and is reachable from all nodes in normal sensor reporting
- The base station is aware of the computation speed of the slowest node in the network
- The Sybil computes n times faster than the slowest genuine node
- Nodes have stores of puzzles generated offline

This case is interesting because it provides the least challenge to the detection system, and it can be expected to be most favorable to detection. In this restricted case, we are able to state that the puzzle based system will statistically detect any Sybil masquerading any more than a single node, almost surely. The detection convergence is faster with larger number of nodes masqueraded by the Sybil.

An Adoption Theory of Secure Software Development Tools

Faculty: Dr. Emerson Murphy-Hill

We have developed an initial adoption model, based on (a) a generic diffusion of innovation model and (b) our own experience on what may influence adoption. We used this model as a starting point to develop an interview protocol and to ask directed questions during our pilot interviews. We have interviewed 45 software developers at small and large firms, including several developers in management roles and several security experts familiar with development tools. Interviews lasted about 1 hour each, giving us a very large amount of data.

These interviews have been transcribed and qualitatively coded to refine our hypotheses about why developers use and do not use tools.

One challenge we face is balancing breadth and depth of data collection. On one hand, we want to have high confidence that our results are generalizable, but on the other, we want to have a complete picture of all factors affecting adoption. To mitigate this risk, we have been looking to prior work that has faced similar challenges in other domains, and observed how those challenges have been overcome.

Studying Latency and Stability of Closed-Loop Sensing-Based Security Systems

Faculty: Dr. Rudra Dutta, Dr. Meeko Oishi

We identified a suitable model of multi-modal linear system with discrete transitions that we will use initially. We also identified the specific scenario that we will address initially out of the few security systems that we proposed – namely we propose to initially focus on the area of multipath multihop wireless networking approaches (suitable for attaining selective reliability in tactical networks). We have also undertaken and made good progress on the cross-training of the two PIs on the nomenclature and terminology of the fields of network, security and dynamic systems. Subsequently, we successfully formulated the problem of representing routing policies for wireless networks in a linear system scenario, and formulated a dynamic programming representation for characterizing the near-optimality of routing policies. We then examined the feasibility of this approach to theoretically characterize simple routing policies for simple network topologies, and draw conclusions about their transient and asymptotic performances.

A significant accomplishment, in intellectual terms, in the linear system formulation of multihop wireless networks was the identification of states such that it was possible to express metrics of practical interest in terms of reachability in state space. This will allow obtaining theoretical insight into the robustness of specific network instances in essential terms, rather than merely by simulation of performance, and lead to theoretical characterization of the presence or absence of problematic behavior in the network, such as route looping, packet reverberation, etc. Our ongoing involvement in this project as well as Science of Security discussions have indicated to us that other control theoretic methods, beyond the ones we have been considering, may hold significant potential contributions for cybersecurity from a scientific approach. We are investigating these as an additional issue in our ongoing work with the hope of soon articulating such potential benefit more explicitly.

Security Metrics: Science of Security Configuration Synthesis

Faculty: Dr. Ehab Al-Shaer

Formulating of Metric-driven security architecture synthesis: We developed a formal framework for automating the security architecture design synthesis based on integrated metrics of isolation, usability and cost. ConfigSynth takes security requirements in terms of isolation, business constraints in terms of usability and deployment cost, along with the network topology, as inputs. Then it synthesizes optimal and cost-effective security configurations satisfying the constraints. ConfigSynth also provides optimal placements of different security devices in the network according to the given network topology.

Firewall policy metrics. We generated number of synthesized firewall policies for various network sizes and topologies. We then conducted various experiments to study the impact of firewall policy structure on the complexity (manageability) and vulnerability (quality of protection) offered by these polices. The manageability is measured by the potential misconfiguration when deploying or changing distributed firewall polices. The quality of protection is measured by the capability of distributed firewall polices to (1) prevent attacks, and (2) isolate or restrict the propagation of attacks in a network, and (3) mitigate attack by dynamic policy adaptation.

Towards a Scientific Basis for User Centric Security Design

Faculty: Dr. Ting Yu, Dr. Ninghui Li, Dr. Robert Proctor

We have found that presenting a high-level risk summary score can cause users to prefer lower-risk app, and to examine the permissions more. We are combining different types of user studies to understand the problem of risk communication. We conduct user studies via both Amazon Mechanical Turk and lab setting. We are also conducting quantitative experiments, where we can measure subjects' actions, instead of only conducting surveys and relying on users' subjective answers to the survey questions.

An Investigation of Scientific Principles Involved in Software Security Engineering

Faculty: Dr. Mladen Vouk, Dr. Laurie Williams

We studied capabilities of "classical" reliability models for assessment and prediction of innate security properties of open source software under "normal" usage. Assessment of Fedora series of products (version 1 through 19) shows that a) it is possible to use classical models (e.g., Yamada model) to describe both non-security and security properties of a release. It is also possible to use the model to estimate residual vulnerabilities and some either software parameters. This offers the option of using that information in developing attack models – such as those based on Rivers' hypergeometric model. We are preparing a survey geared toward industry of a comprehensive set of security practices to be used to examine the state-of-the-practice. The set of practices are based upon the union of the Building Security In Maturity Model (BSIMM), the Software Assurance Maturity Model (OpenSAMM), and the Microsoft Secure Development Lifecycle (SDL).

We believe it is necessary to improve on known and proven software reliability engineering methodologies and models, and to develop new ones, before we will be able to guarantee a certain level of security. Unfortunately, current software engineering practices appear to deal well (if appropriate, effort, time and funds are allocated) with regular, non-security, problems, but are much more ad hoc and undeterministic when it comes to security faults or vulnerabilities. They appear to avoid and eliminate vulnerabilities more by luck (aleatoric process) than through knowledge driven (epistemic) methods. The reason is often fault rarity, complexity of processes and problems, process and human lacks, as well as other limitations such as noise, policies, etc. Science normally strives to reduce epistemic (knowledge-based) uncertainty to aleatoric (or innate random, luck?) uncertainty of natural events and processes. Finding vulnerabilities by luck is not a good approach. It is therefore paradoxic that in cyber security we actually prefer that attacks are based on generally known deviations from the correct (but there perhaps due to our own process and education problems), since that results in pattern based attacks which in turn we can recognize as anomalies and defend against using classical engineering methods. If we know what the problem is, we can at least react to it. Of course, it would be much better to avoid it, or to have a built-in science-based defense against it (e.g., strong password vs. white-list and no password).

Argumentation as a Basis for Reasoning about Security

Faculty: Dr. Munindar Singh, Dr. Simon Parsons

This project involves the application of argumentation techniques for reasoning about security decisions. Specifically, we aim to produce a security-enhanced argumentation framework that (a) provides not only inferences to draw but also actions to take; (b) considers multiparty argumentation; (c) measures the mass of evidence on both attacking and supporting arguments in order to derive a defensible conclusion with confidence; and (d) develops suitable critical questions as the basis for argumentation.

We have developed a formal approach for incorporating evidence into argumentation. Based on the evidence, a decision-maker can judge how much conflict and uncertainty there are in an argument, and thus decide how to proceed—whether pursue a particular argumentation scheme or stop and make a decision.

We have also developed an account of the argumentation problems that arise in a practical firewall configuration in an organization in addition to a preliminary classification of critical questions relevant to security adequate for the argumentation problems. We have applied these questions on a simplified scenario for firewall configuration. In essence, we have created a preliminary formalization of the scenario as a way to show how argumentation can be applied therein.

Out initial evaluation approach is through modeling a military scenario. Besides the firewall scenario we are developing, we demonstrate how our approach can be applied in a practical military security problem. Doing so would show that our approach is general and domain-independent. We have also begun to develop an argumentation tool for the firewall scenario. This tool will enable network administrators to make decisions on which packet to allow or deny based on evidence, i.e., the network history and usage constraints. Our tool asks a network administrator to input available evidence, reason about the conflicts and uncertainty regarding whether a packet should be allowed or blocked, and highlight potential problems that require further investigation. We plan to use this tool to conduct a user study.

Additional research is investigating how a decision maker can obtain the amount of evidence, which is a key requisite for calculating the conflict and uncertainty in an argument. We plan to provide guidelines for the firewall scenario first and then extend to a general framework. We are also in the process of extending our argumentation framework so that it supports multiparty argumentation, emphasizing that more than one interest is often at stake in an argument about security. This will take into account the notion of an interactive dialogue between the participants.

Finally, we have increased the emphasis on empirical evaluation and have begun to develop our evaluation studies before implementing solutions.

A Science of Timing Channels in Modern Cloud Environments

Faculty: Dr. Mike Reiter

The architecture that is currently the focus of our effort is central to the goal of our research, which is to develop principled techniques for comprehensively mitigating access-driven timing channels in modern compute clouds, particularly of the "infrastructure as a service" (IaaS) variety. This type of cloud permits the cloud customer to deploy arbitrary guest virtual machines (VMs) to the cloud. The security of the cloud-resident guest VMs depends on the virtual machine monitor (VMM), e.g., Xen, to adequately isolate guest VMs from one another. While modern VMMs are designed to logically isolate guest VMs, there remains the possibility of timing "side channels" that permit one guest VM to learn information about another guest VM simply by observing features that reflect the others' effects on the hardware platform. Such attacks are sometimes referred to as "access-driven" timing attacks.

We have demonstrated for the first time an access-driven timing side-channel on modern virtualized symmetric multiprocessing (SMP) platforms with a fidelity that is sufficient to enable an attacker VM to extract a cryptographic key from a victim VM. Virtualized SMP systems are very common today, not only for IaaS clouds as described above, but also for desktops that use virtualization to sandbox application or OS compromises. Constructing such a side-channel required overcoming challenges including core migration, numerous sources of channel noise, and the difficulty of preempting the victim with sufficient frequency to extract fine-grained information from it. This work has attracted the attention of industry, and we have been consulted by commercial cloud operators and VMM vendors to discuss ways to mitigate it.

Quantifying Mobile Malware Threats

Faculty: Dr. Xuxian Jiang

In the process of characterizing mobile malware and systematizing their behavior, we have studied their installation methods, activation mechanisms as well as the nature of carried malicious payloads. Our results show that (1) 86.0% of them repackage legitimate apps to include malicious payloads; (2) 36.7% contain platform-level exploits to escalate privilege; (3) 93.0% exhibit the bot-like capability. We have widely shared our dataset in the Android Malware Genome Project within the research community. In particular, the dataset has been shared with colleagues from more than 250 research universities, industry labs and government agencies (including NSA).

Spatiotemporal Security Analytics and Human Cognition

Faculty: Dr. David Roberts

We developed a sound experimental protocol for collecting keyboard data in the typing game. We have developed and tested a game usage profile that uses action-level trace data to classify players' goals; initial evaluation indicates that classification accuracy with these features is on-par with summary speed features. We completed the pilot study data collection for the mouse movement environment. We had 179 participants complete the study, from which we obtained more than 4 million low-level input events.

We verified the 40Hz sampling rate to get high fidelity mouse movement analytics.

In our early analysis, we were able to identify different play styles using only high-level features computed from the analytics. The play styles associated with maximizing accuracy and minimizing time resulted in a clear tradeoff between the frequency of erroneous actions and the duration of play normalized for each player. These data yield classification accuracy of 82.4% using off-the-shelf machine learning techniques. We have identified a set of low-level input features. Specifically, 57 base features, across both spatial and temporal data dimensions, have been computed for each "segment" of mouse movements.

We have built two visualization tools that have driven hypothesis formation for our cognitive models. The first illustrates the path players move their mouse through. The second indicates the location of clicks on the screen given where the player should be targeting. Between these two tools, we were able to get better insight into play characteristics, which led to better analytics mining and ultimately to a new cognitive model of game play. We have developed a data collection protocol for a second game that will enable us to collect detailed analytics describing the typing patterns users display when typing under various cognitive loads.

Normative Trust: Toward a Principled Bases for Enabling Trustworthy Decision Making

Faculty: Dr. Munindar Singh

The proposed approach is centered on the notion of what we term normative relationships—or norms for short—directed from one principal to another. An example of a normative relationship is a commitment: is the first principal committed to doing something for the second principal? (The other main types of normative relationships are authorizations, prohibitions, powers, and sanctions.) Our broad research hypothesis is that trust can be modeled in terms of the relevant norms being satisfied or violated. The approach is natural and promising because it seeks to provide a conceptual grounding to the idea of trust, which in the existing literature is treated either noncomputationally or in an over-simplified manner.

We have expanded our approach to handle trust updates for group commitments and added a temporal discounting window to our approach to model the forgetting effect for an agent. We have also developed a framework, but didn't complete implementing and evaluating it, in which the discounting window sizes for different subjects can be obtained via regression from subjects' labeled trust values.

We created a small dataset extracted from the Enron emails dataset wherein subjects provided estimates of (their individual understanding of) the strength of the trust relationships between the communicating parties. Using this dataset, we developed an approach for extracting commitments from emails or chat sentences that involves a combination of text processing heuristics, recent natural language parsing tools, and machine learning. This enabled us to develop an approach for learning a model of a subject by correlating his or her estimates of the trustworthiness of different parties with how the emails suggest that their commitments to one another were created, discharged, canceled, or delegated.

Low-level Analytics Models of Cognition for Novel Security Proofs

Faculty: Dr. David Roberts, Dr. Robert St. Amant

We have developed and empirically evaluated a cognitive model of micro-strategy tradeoffs for playing a casual memory game. We have completed the first pass of a machine learning approach to predicting a player's meta-level goal (playing for speed versus for accuracy) with promising initial performance at 69%. We have finished a prototype of our adaptation algorithm that will complement the cognitive and ML models in creating our Human Subtlety Proofs. The analytics software we have developed will generalize to research on other games. The PIs have begun to explore potential links with other lablet projects, such as the roles played by norms and expectations in decision making in security contexts and how smartphone sensors can be used on mobile devices to infer things about cognition and behavior.

Modeling the risk of user behavior on mobile devices

Faculty: Dr. Ben Watson, Dr. Will Enck, Dr. Anne McLaughlin, Dr. Michael Rappa

Our primary accomplishments are the completion of the experimental design for our lab study examining the relationships between stress and insecure behavior, and the performance of a preliminary study on the signals that mobile users consider when deciding which apps are safe to install.

Design for lab study. Using mobile phones, we will ask participants to install several mobile apps using an store we build and populate. Some of these apps will be malicious. As they perform this task, we will sometimes impose stressors: noise, time pressure, or multitasking. Security signals will either be subtle or obvious. We will measure whether or not participants correctly identify the malicious apps, and the speed with which they perform each task.

Preliminary study on security signal importance. 176 participants on Amazon's Mechanical Turk ranked eight possible signals in app stores of app security and reliability. Results were that users ranked app permissions first by a wide margin, with review content a clear second, and number of stars, downloads and reviews clustered closely in the third, fourth and fifth ranks. Correspondingly, we will emphasize these signals in our lab study.

Empirically Evaluating and Quantifying the Effects of Inspections and Testing on Security Vulnerabilities

Faculty: Dr. Jeffery Carver

We have developed a research plan to empirically evaluate the effectiveness of peer code review to prevent security defects and have built a set of empirically validated keywords, which can be used to find security defects. We have added steps to empirically generate the keywords set based on text mining and have added steps to validate the completeness of our keywords set. We have populated our database with peer code review data mined from 34 Open Source repositories. We are training two REU students on analyzing code review request to find security defects.

Understanding the Fundamental Limits in Passive Inference of Wireless Channel Characteristics

Faculty: Dr. Huaiyu Dai, Dr. Peng Ning

Link signature, which refers to the unique and reciprocal wireless channel between a pair of transceivers, is used for signal authentication and shared secret construction for various wireless security applications. Particularly, half a wavelength is the key distance metric used inexisting link signature based security mechanisms for security assurance. However, based on our study of existing literature (e.g., the "sum-Kronecker" model), it is found that the correlation between the Angle of Arrival (AoA) and Angle of Departure (AoD) of a wireless link is an important factor that affects the channel correlation, and the spatial correlation of wireless channels increases as the AoA-to-AoD correlation increases.

We have also observed in the simulations that even for spatial separation of 20 wavelengths, significant channel correlation (about 0.9) still exists when the AoA-to-AoD correlation is sufficiently strong (which physically implies highly correlated scattering mechanisms at the transmitter and receiver). For a 900MHz wireless link with wavelength of 1/3 meter, this indicates that the legitimate channel and the adversary channel are still highly correlated even when the adversary is about 7 meters away from the legitimate device. Therefore the safety distance metric needs to be reevaluated in such cases.

An Investigation of Scientific Principles Involved in Attack-Tolerant Software

Faculty: Dr. Mladen Vouk

Deliberate attacks and strategies on information technology resources invariably have a psychosocial component. When considering psychosocial effects involved in attacks (both human-driven, and automated), we deal with both quantitative and qualitative models since many of the human drivers and reactions (of attackers and the victims) are hard to quantify beyond an ordinal scale. It is not just network and system security performance information that is helpful (e.g., how many vulnerabilities there are in the system, what are field failure rates, etc.), but also what the education and training level of both attackers and defenders is, what their intentions and motivations are, and whether any of the human participants is in an altered state of mind – e.g., under a lot of stress or for some other reason is not his or her usual self. In considering psychosocial component, it is important to note that there is an independent contribution of psychological (for instance personality, cognition) and sociological (for instance socioeconomic background, peer group values and norm) factors to the development of comprehensive theoretical models of human online behaviors (including attackers, victims, end users and IT managers). When dealing with risk that involves humans we also need to account for possible deception and bad intent. For credible science-basedmodels, we need sophisticated analyses of attack profiles and exploits that peel back confounding information of the intentional or unintentional transfer functions. This is true for both human-directed and automated (robo) attacks. Latter will invariably have a human-developed built-in strategy that probably has limited self-adaptation capabilities. To assess intent of an attacker, one needs to have information about the value of the asset being protected, information about vulnerabilities of the asset, and information about the way attack is conducted.

Human element (users, administrators, attackers, …) and end-to-end risk assessment and mitigation remains the key unknown, and probably the largest source of uncertainty. As part of implementation of automatic attack tolerance, we plan to investigate new theories and mechanisms of coupling risk metrics with human and social elements (such as already mentioned individual, group and organizational utility functions). In a large system, administrator has to intersect, evaluate, understand and respond to a large number of security alerts. That position also needs to communicate and coordinate with administrators at higher layers such as middleware, applications and ultimately end-users. As part of the attack-tolerance dashboard sub-project, the goal is to investigate a) techniques that would determine and summarize security risks based on analysis of the whole end-to-end stack – from networks, to applications, to users, and b) present this information in a human-centric format, and with attack tolerance advice, to increase situation awareness through a holistic view of the risk with the ability to zoom in on specific risk regions if necessary.

Centennial Mesh Network & CentMesh Drone Challenge

Faculty: Dr. Mihail Sichitiu, Dr. Rudra Dutta

SOSI funds were used to develop and support a general sensor/actuator interface seamlessly integrated with the Centennial Mesh Network Testbed (CentMesh), create and showcase a reference implementation of the same in a class of airborne autonomous CentMesh nodes, and conduct an outreach activity in the form of a student programming challenge. This activity was funded by other sources of funding to CentMesh during the same period. For context, we review CentMesh briefly. Designing and building this teaching and research facility was one of the specific areas of effort in the original SOSI vision, and has been funded using SOSI funds as well as a separate DURIP grant from

ARO. The purpose of the testbed is to provide a deeply programmable, flexible, realistic, outdoor facility for experimentation in support of research, development, transition and teaching in secure pervasive computing and networking applications. The testbed consists of 14 stationary WiFi nodes mounted on poletops and 8 mobile nodes (on pushcarts). All nodes are fully programmable from the MAC layer all the way up the network stack. This programmability affords the researchers the ability to investigate security approaches from a number of vectors such as physical layer jamming, MAC layer attacks, routing approaches, etc. Since 2012, CentMesh has been operational as originally conceived, and operating as a comparatively stable facility; although we continue to make ongoing evolutionary improvements as befits a research facility that seeks to keep up with researcher needs. We designed and realized the CentMesh Mobile Node, designed to house sensing and actuating capabilities, offer portability and mobility, while allowing connectivity to the CentMesh fixed infrastructure. The nodes can be mounted on a "vehicle". The vehicle can be anything capable of moving a node around, e.g., a student backpack, a bus, a car or a hexcopter. When the vehicle is capable of complying with MAVLink, a standardized interface, the CentMesh Mobile Node can also control the motion of the vehicle. We realized the CentMesh Mobile Node as an embedded Linux board - we chose the Beaglebone Black, due to its relatively low cost, reliability, excellent software support and flexibility: with a host of exposed interfaces, A/D and D/A connections, the platform can be extended almost indefinitely. The main communication links with the ground stations and inter-drone communications are achieved using standard high power USB WiFi cards and TCP/IP protocols. Using these as the base, we designed and realized the CentMesh Drones. Each CentMesh drone is comprised of several modules interconnected through standard interfaces:

- The airframe module includes the frame, propellers, brushless motors, electronic speed controllers (ESCs), battery, as well as the 5V regulator that provides power to the autopilot and the sensor node. The airframe receives standard RC PWM control signals to the ESCs from the autopilot module.
- The autopilot module has the main role to stabilize the platform and to fly it at locations received from the sensor node. The autopilot requires an inertial measurement unit (IMU) typically hosting three axis accelerometers and gyroscopes, and a barometric sensor for stabilization. For navigation at the desired points it requires a GPS unit and a compass. A battery monitoring unit allows for failsafe landings on low battery. An abort signal (kill switch), sent in a separate frequency band (72MHz) allows for a forced landing or for a safe return to the launch point in case of major software failure on the sensor node. The commands from the sensor node are sent using the MAVLink format through an USB link.
- The CentMesh Mobile Node described above.

Virtual environment: We created a virtual environment corresponding to CentMesh drones, for ease of programming and easy student access. In such a case, the airframe is replaced with a physics simulation, the autopilot by a software version of the same autopilot, and the Mobile Node platform by a similar Linux platform with the same software.

The first ever programming challenge at NCSU on unmanned aerial computing platforms culminated in a day-long final challenge in which the finalist teams took turns running their programs on CentMesh drones to see if their code would fly. All the code submitted by the finalist teams performed admirably. There was only one quasi-crash (forced landing) which turned out, upon examination of the logs, to have been the fault of a firmware bug and not the code submitted for the challenge.

The CentMesh drones are fabricated using standard hexacopter frame and the well-known Arducopter controllers; but instead of communicating to the Arducopter using a Ground Control Station software (or simply flying it manually like a radio-controlled toy plane), we put a BeagleBone Black, a small Linux computer, on each of them. The challenge consists of pre-programming these computers to fly the drones, so they can perform in true autonomous fashion. For example, a simple challenge could be: "Write a program so that upon execution, the drone will take off vertically, hold altitude at 50 feet, and

wait for a message sent via WiFi containing a set of coordinates; upon receiving the message, will move to those coordinates, hover for 30 seconds, then land."

Trusted Smart Room

Faculty: Dr. Mladen Vouk, John Streck

The project investigated the cloud enabling of a fixed asset such as a multi-media enabled meeting room partnered with the necessary security capabilities paired to the needs of the particular user group. The cloud enablement was to add to the cloud OS, in this case NCSU's VCL cloud system, the architecture components to allow a federated pool of resources (in this case the TSR) as an XaaS (Anything as a Service) component to the traditional compute and storage resources of a cloud infrastructure. To do this, the cloud architecture/implementation needed to have a resource reservation system that enforced security profiles per user group upon the resource requested for the period of operation requested. The time and access of the TSR would control access (security) to the physical room, to assets in the TSR, to connected cloud compute and storage devices and network connectivity from the TSR to another cloud asset (simplest case of a TSR to another cloud enabled TSR via an on-demand VPN). The cloud components for the TSR were architected so any high valued asset could be securely added to any cloud environment (as Anything as a Service[XaaS]). Additionally, the design was expanded to allow mining of the provenance data not only for security auditing but also for billing of services. The billing could be either monetary or a simple credit system for resource allotment. The initial system was targeted to be implemented on NCSU's VCL cloud first but could be implemented on other cloud platforms.

Access Control Policies

Faculty: Dr. Mladen Vouk, Dr. Peng Ning

Controlling access based on access control policies (ACPs) is a necessary security mechanism to protect critical resources in operating systems or applications, such as health care and military applications. As operational and security requirements of a system evolve, policy authors often change access control policies. Our research goal is to assist policy authors in effectively managing access control policies by providing insights on historical trends and evolution patterns in access control policies. To achieve this goal, we performed an empirical study to analyze how ACPs evolve in three systems: Security Enhanced Linux (SELinux), Virtual Computing Laboratory (VCL), and a network intrusion detection system (Snort). We first empirically observe growth trends of ACPs and corresponding systems in terms of the number of policy lines and lines of code (LOC), respectively. We propose an approach to extract evolution patterns based on analyzing ACP historical change data. An evolution pattern indicates an abstraction of change in the permissions/privileges assigned to a group or a user. We further develop a prediction model based on our evolution pattern ranking to predict how ACPs evolve. We found eight frequently occurring evolution patterns across the three systems. Our results indicate that our model can predict evolution patterns in ACPs with a precision of 50-80%, a recall of 70-90% and an F-measure of 65-75%.

References

Ting Yu, PhD. Data Anonymization: The Challenge from Theory to Practice , (03 2014)

Nirav Ajmeri, Munindar P. Singh, Simon Parsons, Chung-Wei Hang. Argumentation, Evidence, and Schemes, ( )

Shalini Chauhan, Michael Devetsikiotis. AGENT BASED FRAMEWORK FOR AVATAR INTERACTIONS IN AN ADAPTIVE VIRTUAL WORLD GAME ENVIRONMENT, ( )

TOTAL: 3

Books

Patrick Dreher, Mladen Vouk. Utilizing Open Source Cloud Computing Environments to Provide Cost Effective Support for University Education and Research, Hershey, Pennsylvania: IGI Global, (04 2012)

TOTAL: 1

Non-Peer-Reviewed Conference Proceeding publications (other than abstracts)

M. Riaz, J. Slankas, J. King, L. Williams. Hidden in Plain Sight: Automatically Identifying Security Requirements from Natural Language Artifacts, IEEE International Requirements Engineering Conference (RE) 2014. 25-AUG-14, . : ,

Anup Kalia, Nirav Ajmeri, Kevin Chan, Jin-Hee Cho, Sibel Adali, Munindar Singh. A Model of Trust, Moods, and Emotions in Multiagent Systems and its Empirical Evaluation, 17th AAMAS Workshop on Trust in Agent Societies. , . : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer. Metrics for Automated Network Security Design, 20th Annual Network & Distributed System Security Symposium (NDSS). 25-FEB-13, . : ,

Chengzhi Li, Huaiyu Dai, Liang Xiao, Peng Ning. Analysis and Optimization on Jamming-resistantCollaborative Broadcast in Large-Scale Networks, 2010 Asilomar Conference on Signals, Systems, and Computers. 07-NOV-10, . : ,

TOTAL: 4

Peer-Reviewed Conference Proceeding publications (other than abstracts)

Rucha Tembe, Olga Zielinska, Yuqi Liu, Kyung Wha Hong, Emerson Murphy-Hill, Chris Mayhorn, Xi Ge. Phishing in international waters: Exploring Cross-national Differences in Phishing Conceptualizations Between Chinese, Indian and American Samples, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Anthony Thyron Rivers, Mladen A. Vouk, Laurie A. Williams. On Coverage-Based Attack Profiles, 2014 IEEE 8th International Conference on Software Security and Reliability-Companion (SERE-C). 30-JUN-14, San Francisco, CA, USA. : ,

Maria Riaz, Jason King, John Slankas, Laurie Williams. Hidden in plain sight: Automatically identifying security requirements from natural language artifacts, 2014 IEEE 22nd International Requirements Engineering Conference (RE). 25-AUG-14, Karlskrona, Sweden. : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer, Rajesh Kavasseri. Impact Analysis of Topology Poisoning Attacks on Economic Operation of the Smart Power Grid, 2014 IEEE 34th International Conference on Distributed Computing Systems (ICDCS). 30-JUN-14, Madrid, Spain. : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer, Rakesh B. Bobba. Moving Target Defense for Hardening the Security of the Power System State Estimation, First ACM Workshop on Moving Target Defense. 07-NOV-14, Scottsdale, Arizona, USA. : ,

Vasant Tendulkar, Adwait Nadkarni, William Enck. NativeWrap: Ad Hoc Smartphone Application Creation for End Users, The ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2014). 23-JUL-14, Oxford, United Kingdom. : ,

Tsung-Hsuan Ho, Daniel Dean, Xiaohui Gu, William Enck. PREC: practical root exploit containment for android devices, The 4th ACM conference on Data and application security and privacy. 03-MAR-14, San Antonio, Texas, USA. : ,

Victor Heorhiadi, SeyedKaveh Fayaz, Michael K. Reiter, Vyas Sekar. SNIPS: A Software-Defined Approach for Scaling Intrusion Prevention Systems via Offloading, 10th International Conference on Information Systems Security, ICISS 2014. 16-DEC-14, . : ,

Xiaofan He, Huaiyu Dai, Yufan Huang, Dong Wang, Wenbo Shen, Peng Ning. The security of link signature: A view from channel models, 2014 IEEE Conference on Communications and Network Security (CNS). 29-OCT-14, San Francisco, CA, USA. : ,

Amit K. Chopra, Munindar P. Singh. The thing itself speaks: Accountability as a foundation for requirements in sociotechnical systems, 2014 IEEE 7th International Workshop on Requirements Engineering and Law (RELAW). 26-AUG-14, Karlskrona, Sweden. : ,

Amiangshu Bosu, Jeffrey C. Carver, Munawar Hafiz, Patrick Hilley, Derek Janni. Identifying the characteristics of vulnerable code changes: an empirical study, the 22nd ACM SIGSOFT International Symposium. 16-NOV-14, Hong Kong, China. : ,

Trisha Biswas, Kendra Lesser, Rudra Dutta, Meeko Oishi. Using linear system reliability to obtain theoretical understanding of wireless routing, GLOBECOM 2014 - 2014 IEEE Global Communications Conference. 08-DEC-14, Austin, TX, USA. : ,

Amiangshu Bosu, Jeffrey C. Carver, Munawar Hafiz, Patrick Hilley, Derek Janni . When Are OSS Developers More Likely to Introduce Vulnerable Code Changes? A Case Study, 10th International IFIP WG 2.13 Conference on Open Source Systems. 06-MAY-14, . Berlin Heidelberg: Springer, Springer

Xibin Gao, Munindar P. Singh. Extracting normative relationships from business contracts, 2014 international conference on Autonomous agents and multi-agent systems . 06-MAY-14, . : ,

R. Slavin, J.-M. Lehker, J. Niu, T. Breaux. Managing Security Requirement Patterns Using Feature Diagram Hierarchies, 22nd IEEE International Requirements Engineering Conference. 25-AUG-14, . : ,

Paulo Shakarian, Gerardo I. Simari, Geoffrey Moores, Simon Parsons, Marcelo A. Falappa. An argumentation-based framework to address the attribution problem in cyber-warfare, 3rd ASE International Conference on Cyber Security. 27-MAY-14, . : ,

Adwait Nadkarni, Vasant Tendulkar, William Enck. NativeWrap, the 2014 ACM conference. 23-JUL-14, Oxford, United Kingdom. : ,

Shweta Subramani, Mladen Vouk, Laurie Williams. An analysis of Fedora security profile, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Qian Liu, Juhee Bae, Benjamin Watson, Anne McLaughhlin, William Enck. Modeling and sensing risky user behavior on mobile devices, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Agnes Davis, Ashwin Shashidharan, Qian Liu, William Enck, Anne McLaughlin, Benjamin Watson. Insecure behaviors on mobile devices under stress, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Roopak Venkatakrishnan, Mladen A. Vouk. Diversity-based detection of security anomalies, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Trisha Biswas, Kendra Lesser, Rudra Dutta, Meeko Oishi. Examining reliability of wireless multihop network routing with linear systems, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Ashwini Rao, Hanan Hibshi, Travis Breaux, Jean-Michel Lehker, Jianwei Niu. Less is more?, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Rucha Tembe, Olga Zielinska, Yuqi Liu, Kyung Wha Hong, Emerson Murphy-Hill, Chris Mayhorn, Xi Ge. Phishing in international waters, the 2014 Symposium and Bootcamp. 08-APR-14, Raleigh, North Carolina. : ,

Amiangshu Bosu. Characteristics of the vulnerable code changes identified through peer code review, Companion the 36th International Conference. 31-MAY-14, Hyderabad, India. : ,

Patrick Dreher, Mladen Vouk. Integration of high-performance computing into a VCL cloud, the 8th Workshop. 17-NOV-13, Denver, Colorado. : ,

Tsung-Hsuan Ho, Daniel Dean, Xiaohui Gu, William Enck. PREC: practical root exploit containment for android devices, the 4th ACM conference. 03-MAR-14, San Antonio, Texas, USA. : ,

Shundan Xiao, Jim Witschey, Emerson Murphy-Hill. Social influences on secure development tool adoption, the 17th ACM conference. 15-FEB-14, Baltimore, Maryland, USA. : ,

Yinqian Zhang, Michael K. Reiter. Düppel, the 2013 ACM SIGSAC conference. 04-NOV-13, Berlin, Germany. : ,

Adwait Nadkarni, William Enck. Preventing accidental data disclosure in modern operating systems, the 2013 ACM SIGSAC conference. 04-NOV-13, Berlin, Germany. : ,

Jim Witschey. Secure development tool adoption in open-source, the 2013 companion publication for conference. 26-OCT-13, Indianapolis, Indiana, USA. : ,

Graham Cormode, Entong Shen, Xi Gong, Ting Yu, Cecilia M. Procopiuc, Divesh Srivastava. UMicS, the 22nd ACM international conference. 27-OCT-13, San Francisco, California, USA. : ,

Yinqian Zhang, Ari Juels, Michael K. Reiter, Thomas Ristenpart. Cross-VM side channels and their use to extract private keys, the 2012 ACM conference. 16-OCT-12, Raleigh, North Carolina, USA. : ,

Rucha Tembe, Kyung Wha Hong, Emerson Murphy-Hill, Christopher B. Mayhorn, Christopher M. Kelley. American and Indian Conceptualizations of Phishing, 2013 3rd International Workshop on Socio-Technical Aspects in Security and Trust (STAST). 29-JUN-13, New Orleans, LA, USA. : ,

Wenbo Shen, Peng Ning, Xiaofan He, Huaiyu Dai. Ally Friendly Jamming: How to Jam Your Enemy and Maintain Your Own Wireless Connectivity at the Same Time, 2013 IEEE Symposium on Security and Privacy (SP) Conference dates subject to change. 19-MAY-13, Berkeley, CA, USA. : ,

Amiangshu Bosu, Jeffrey C. Carver. Peer Code Review to Prevent Security Vulnerabilities: An Empirical Evaluation, 2013 IEEE 7th International Conference on Software Security and Reliability-Companion (SERE-C). 18-JUN-13, Gaithersburg, MD, USA. : ,

Anup K. Kalia, Hamid R. Motahari-Nezhad, Claudio Bartolini, Munindar P. Singh. Monitoring Commitments in People-Driven Service Engagements, 2013 IEEE International Conference on Services Computing (SCC). 28-JUN-13, Santa Clara, CA, USA. : ,

Da Young Lee, Mladen Vouk, Laurie Williams. Using software reliability models for security assessment — Verification of assumptions, 2013 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW). 04-NOV-13, Pasadena, CA, USA. : ,

Shweta Subramani, Mladen Vouk, Laurie Williams. Non-operational testing of software for security issues, 2013 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW). 04-NOV-13, Pasadena, CA, USA. : ,

Xiaofan He, Huaiyu Dai, Wenbo Shen, Peng Ning. Is link signature dependable for wireless security?, IEEE INFOCOM 2013 - IEEE Conference on Computer Communications. 14-APR-13, Turin, Italy. : ,

Jason King. Measuring the forensic-ability of audit logs for nonrepudiation, 2013 35th International Conference on Software Engineering (ICSE). 18-MAY-13, San Francisco, CA, USA. : ,

Graham Cormode, Cecilia M. Procopiuc, Entong Shen, Divesh Srivastava, Ting Yu. Empirical privacy and empirical utility of anonymized data, 2013 IEEE 29th International Conference on Data Engineering Workshops (ICDEW). 08-APR-13, Brisbane, Australia. : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer. A Formal Framework for Network Security Design Synthesis, 2013 IEEE 33rd International Conference on Distributed Computing Systems (ICDCS). 08-JUL-13, Philadelphia, PA, USA. : ,

Peng Li, Debin Gao, Michael K. Reiter. Mitigating access-driven timing channels in clouds using StopWatch, 2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). 24-JUN-13, Budapest, Hungary. : ,

Mohammed Noraden Alsaleh, Saeed Al-Haj, Ehab Al-Shaer. Objective metrics for firewall security: A holistic view, 2013 IEEE Conference on Communications and Network Security (CNS). 14-OCT-13, National Harbor, MD, USA. : ,

Brent Harrison, David L. Roberts. Analytics-driven dynamic game adaption for player retention in Scrabble, 2013 IEEE Conference on Computational Intelligence and Games (CIG). 11-AUG-13, Niagara Falls, ON, Canada. : ,

Jim Witschey, Emerson Murphy-Hill, Shundan Xiao. Conducting interview studies: Challenges, lessons learned, and open questions, 2013 1st International Workshop on Conducting Empirical Studies in Industry (CESI). 20-MAY-13, San Francisco, CA, USA. : ,

Kendra Lesser, Meeko Oishi, R. Scott Erwin. Stochastic reachability for control of spacecraft relative motion, 2013 IEEE 52nd Annual Conference on Decision and Control (CDC). 10-DEC-13, Firenze. : ,

Shundan Xiao, Jim Witschey, Emerson Murphy-Hill. Social Influences on Secure Development Tool Adoption: Why Security Tools Spread, 2013 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW). 23-FEB-13, . : ,

Jim Witschey. Security Tool Adoption in Open-Source, 2013 Conference on Systems, Programming, Languages and Applications: Software for Humanity. 26-OCT-13, . : ,

Wei YangYang, Mukul R. Prasad, Tao Xie. A Grey-Box Approach for Automated GUI-Model Generation of Mobile Applications, 16th International Conference, FASE 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013. 16-MAR-13, . Berlin: Springer Berlin Heidelberg, Springer Berlin Heidelberg

Rucha Tembe, Kyung Wha Hong, Emerson Murphy-Hill, Christopher B. Mayhorn, Christopher M. Kelley. American and Indian conceptualizations of phishing, Third International Workshop on the Socio-Technical Aspects of Security and Trust. 29-JUN-13, . : ,

R. Slavin, H. Shen, J. Niu. Characterizations and Boundaries of Security Requirements Patterns, 2nd IEEE Workshop on Requirements Engineering Patterns (RePa'12). 24-SEP-12, . : ,

Munindar P. Singh. Can't We All Just Get Along? Agreement Technologies and the Science of Security, 2nd International Conference on Agreement Technologies. 01-AUG-13, . : ,

Maria Riaz, Laurie Williams. Security Requirements Patterns: Understanding the Science Behind the Art of Pattern Writing, Requirements Patterns (RePa) at Requirements Engineering 2012. 24-SEP-12, . : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer. A Formal Approach for Network Security Management Based on Qualitative Risk Analysis, IFIP/IEEE International Symposium on Integrated Network Management (IM). 27-MAY-13, . : ,

Mohammad Ashiqur Rahman, Ehab Al-Shaer. A Formal Framework for Network Security Design Synthesis, 33rd International Conference on Distributed Computing Systems (ICDCS). 08-JUL-13, . : ,

Ari Juels, Michael K. Reiter, Thomas Ristenpart, Yinqian Zhang. Cross-VM side channels and their use to extract private keys, 19th ACM Conference on Computer and Communications Security. 16-OCT-12, .

Xusheng Xiao, Wei Yang, Rahul Pandita, William Enck, Tao Xie. WHYPER: Towards Automating Risk Assessment of Mobile Applications, 22nd USENIX Security Symposium. 14-AUG-13, . : ,

Adwait Nadkarni and William Enck. Preventing Accidental Data Disclosure in Modern Operating Systems, 20th ACM Conference on Computer and Communications Security (CCS). 04-NOV-13, . : ,

Patrick Morrison, Laurie Williams. An Analysis of HIPAA Breach Data, USENIX HealthSec 2012. 06-AUG-13, . : ,

P. Morrison, C. Holmgreen, A. Massey, L. Williams. Proposing Regulatory-Driven Automated Test Suites, Agile 2013: Agile Software Development. 05-AUG-13, . : ,

D. Gao, M. K. Reiter, P. Li. Mitigating access-driven timing channels in clouds using StopWatch, 43rd IEEE/IFIP International Conference on Dependable Systems and Networks. 24-JUN-13, . : ,

Kendra Lesser, Meeko Oishi, R.S. Erwin. Stochastic reachability for control of spacecraft relative motion, IEEE Conference on Decision and Control. 10-DEC-13, . : ,

Jason King, Laurie Williams. Secure Logging and Auditing in Electronic Health Record Systems: What Can We Learn from the Payment Card Industry, USENIX HealthSec 2012. 08-AUG-12, . : ,

Ben Smith, Laurie Williams, Jason King. Modifying Without a Trace: General Audit Guidelines are Inadequate for Electronic Health Record Audit Mechanisms, International Health Informatics Symposium (IHI 2012). 28-JAN-12, . : ,

Christopher M. Kelley, Kyung Wha Hong, Christopher B. Mayhorn, Emerson Murphy-Hill. Something smells phishy: Exploring definitions, consequences, and reactions to phishing, Human Factors and Ergonomics Society 56th Annual Meeting. 22-OCT-12, . : ,

Anup K. Kalia, Zhe Zhang, Munindar P. Singh. Trustworthy Decision Making via Commitments, 16th International Workshop on Trust in Agent Societies (Trust). 06-MAY-13, . : ,

Anup K. Kalia, Hamid R. Motahari Nezhad, Claudio Bartolini, Munindar P. Singh. Monitoring Commitments in People-Driven Service Engagements, 10th IEEE International Conference on Services Computing (SCC). 27-JUN-13, . : ,

K. W. Hong, C. M. Kelley, C. B. Mayhorn, E. Murphy-Hill. Keeping up with the Joneses: Assessing phishing susceptibility in an email task, Human Factors and Ergonomics Society 57th Annual Meeting. 30-SEP-13, . : ,

X. He, H. Da, W. Shen, P. Ning. Is Link Signature Dependable for Wireless Security?, 2013 IEEE Conference on Computer Communications (INFOCOM). 14-APR-13, . : ,

David Roberts, Brent Harrison. When Players Quit (Playing Scrabble), Eighth Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2012). 08-OCT-12, . : ,

Brent Harrison, David L. Roberts. Analytics-Driven Dynamic Game Adaption for Player Retention (in Scrabble), IEEE 2013 Conference on Computational Intelligence in Games (IEEE CIG). 11-AUG-13, . : ,

Graham Cormode, Cecilia M. Procopiuc, Entong Shen, Divesh Srivastava, Ting Yu. UMicS: From Anonymized Data to Usable MicroData, ACM Conference of Information and Knowledge Management. 27-OCT-13, . : ,

G. Cormode, C. Procopiuc, E. Shen, D. Srivastava, T. Yu. Empirical Privacy and Empirical Utility of Anonymized Data, Workshop on Privacy-Preserving Data Publication and Analysis (PrivDB). 08-APR-13, . : ,

Travis Breaux, Hanan Hibshi, Ashwini Rao, Jean-Michel Lehker. Towards a Framework for Pattern Experimentation: Understanding empirical validity in requirements engineering patterns, 2nd IEEE workshop on Requirements Engineering Patterns. 24-SEP-12, . : ,

Jeffrey Carver, Amiangshuv Bosu. Peer Code Review to Prevent Security Vulnerabilities: An Empirical Evaluation, 2013 IEEE Seventh International Conference on Software Security and Reliability (SERE). 18-JUN-13, . : ,

Titus Barik, Brent Harrison, David L. Roberts, Xuxian Jiang. Spatial Game Signatures for Bot Detection in Social Games, Eighth Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2012). 08-OCT-12, . : ,

Jim Witschey, Emerson Murphy-Hill, Shundan Xiao. Conducting Interview Studies: Challenges, Lessons Learned, and Open Questions, 2013 Workshop on Conducting Empirical Studies in Industry. 20-MAY-13, . : ,

Ninghui Li, Christopher Gates, Jing Chen, Robert Proctor. CodeShield: towards personalized application whitelisting, the 28th Annual Computer Security Applications Conference. 03-DEC-12, Orlando, Florida.

Arpan Chakraborty, Brent Harrison, David L. Roberts, Robert St. Amant, Titus Barik. Modeling the Concentration Game with ACT-R, 12th International Conference on Cognitive Modeling (ICCM). 14-JUL-13, . : ,

Jagan Srinivasan, Wei Wei, Xiaosong Ma, Ting Yu. EMFS: Email-based Personal Cloud Storage, 6th IEEE International Conference on Networking, Architecture and Storage (NAS), 2011 . 28-JUL-11, . : ,

Wei Wei, Ting Yu, Rui Xue. iBigTable: practical data integrity for bigtable in public cloud, CODASPY 2013 : Third ACM Conference on Data and Application Security and Privacy. 18-FEB-13, . : ,

Eric Helms, Laurie Williams. Evaluating Access Control of Open Source Health Record Systems, 3rd Workshop on Software Engineering in Healthcare at the International Conference (SEHC) on Software Engineering (ICSE) 2011. 22-MAY-11, . : ,

JunBum Lim, P. H. Pathak, U. Patel, G. Deuskar, M. Pandian————————————, A. Danivasa, M. Sichitiu————————————, R. Dutta. CentMesh: Modular and Extensible Wireless Mesh Network Testbed, TRIDENTCOM. 17-APR-11, . : ,

Trisha Biswas, Rudra Dutta. Spatially Diffuse Pathsets for Robust Routing in Ad Hoc Networks, Globalcon 2011. 05-DEC-11, . : ,

Parth Pathak, Rudra Dutta. Using Centrality-based Power Control for Hot-spot Mitigation in Wireless Networks, IEEE Globecom 2010. 06-DEC-10, . : ,

Chung-Wei Hang, Munindar Singh. From Quality to Utility: Adaptive ServiceSelection Framework, 8th International Conference on Service Oriented Computing (ICSOC). 07-DEC-10, . : ,

Chung-Wei Hang, Anup Kalia , Munindar Singh. Behind the Curtain: Service Selection via Trust in Composite Services, 2012 IEEE 19th International Conference on Web Services. 24-JUN-12, . : ,

Mladen A. Vouk. A Note on Uncertainty in Real-Time Analytics, the IFIP WoCo 10 on Uncertainty Quantification in Scientific Computing, IFIP WG 2.5 on Numerical Software. 01-AUG-11, . : ,

Meiyappan Nagappan, Aaron Peeler, Mladen Vouk. Modeling Cloud Failure Data: A Case Study of the VirtualComputing Lab, The ICSE 2011 Software Engineering For Cloud Computing Workshop. 22-MAY-11, . :

Meiyappan Nagappan, Brendan Murphy, Mladen Vouk. Which Code Construct Metrics are Symptoms of PostRelease Failures, The 2nd International Workshop on Emerging Trends in Software Metrics (WETSoM 2011). 24-MAY-11, . : ,

Mladen A. Vouk, Pierre A. Mouallem. On High-Assurance Scientific Workflows, 2011 IEEE 13th International Symposium on High-Assurance Systems Engineering. 10-NOV-11, . : ,

Jianchun Jiang, Liping Ding, Ennan Zhai, Ting Yu. VRank: A Context-Aware Approach toVulnerability Scoring and Ranking in SOA, Sixth IEEE International Conference on Software Security and Reliability (SERE). 20-JUN-12, . : ,

Chung-Wei Hang, Anup K. Kalia, Munindar P. Singh. Behind the Curtain: Service Selection via Trust inComposite Services, 10th International Conference on Web Services (ICWS). 24-JUN-12, . : ,

Michael Grace, Yajin Zhou, Zhi Wang, Xuxian Jiang. Systematic Detection of Capability Leaks in Stock Android Smartp, 19th Network and Distributed System Security Symposium (NDSS 2012). 05-FEB-12, .

Yajin Zhou, Zhi Wang, Wu Zhou, Xuxian Jiang. Hey, You, Get Off of My Market:Detecting Malicious Apps in Official and Alternative Android Markets, 19th Network and Distributed System Security Symposium (NDSS 2012). 05-FEB-12, . : ,

Zhi Wang, Chiachih Wu, Michael Grace, Xuxian Jiang. Isolating Commodity Hosted Hypervisors with HyperLock, ACM SIGOPS/EuroSys Conference on Computer Systems (EuroSys 2012). 12-APR-12, . : ,

Emre Yetginer, George Rouskas. Power Efficient Traffic Grooming in Optical WDM Networks, Proceedings of the 28th IEEE conference on Global telecommunications. 30-NOV-09, . : ,

Zhenhuan Gong, Prakash Ramaswamy, Xiaohui Gu, Xiaosong Ma. SigLM: Signature-Driven Load Management for Cloud Computing Infrastructures, IEEE International Conference on Quality of Service (IWQoS). 13-JUL-09, . : ,

Juan Du, Xiaohui Gu, Douglas Reeves. Highly Available Component Sharing in Large-Scale Multi-Tenant Cloud Systems, Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing. 20-JUN-10, . : ,

Juan Du, Wei Wei, Xiaohui Gu, Ting Yu. Towards Secure Dataflow Processing in Open Distributed Systems, Proceedings of the 2009 ACM workshop on Scalable trusted computing. 13-NOV-09, . : ,

Wei Wei, Juan Du, Ting Yu, Xiaohui Gu. SecureMR: A Service Integrity Framework for MapReduce, Proceedings of 2009 Annual Computer Security Applications Conference (ACSAC). 07-DEC-09, . : ,

Juan Du, Wei Wei, Xiaohui Gu, Ting Yu. RunTest: Assuring Integrity of Dataflow Processing in Cloud Computing Infrastructures, Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security. 13-APR-10, . : ,

Christopher Zimmer, Balasubramanya Bhat, Frank Mueller, Sibin Mohan. Time-based intrusion detection in cyber-physical systems, Proceedings of the 1st ACM/IEEE International Conference on Cyber-Physical Systems. 13-APR-10, . : ,

Parth Pathak, Divya Gupta, Rudra Dutta. Loner Links Aware Routing and Scheduling in WMNs, Proceedings of IEEE Advanced Network and Telecommunications Systems (ANTS). 14-DEC-09, . : ,

Parth Pathak , Rudra Dutta. Impact of Power Control on Capacity of Large Scale Wireless Mesh Networks, Proceedings of the 3rd international conference on Advanced networks and telecommunication systems. 14-DEC-09, . : ,

Dheeraj Kandula, Parth Pathak, Rudra Dutta. MF-TCP : Design and Evaluation of TCP for Message Ferry Delay Tolerant Networks, 2009 Australasian Telecommunication Networks and Applications Conference. 12-NOV-09, . : ,

Parth Pathak, Rudra Dutta. Impact of Power Control on Relay Load Balancing in Wireless Sensor Networks, Proceedings of 2010 IEEE Wireless Communications and Networking Conference. 18-APR-10, . : ,

Andrew Meneely, Laurie Williams. Secure Open Source Collaboration: An Empirical Study of Linus' Law, Proceedings of the 16th ACM conference on Computer and communications security. 09-NOV-09, .

Andrew Meneely, Mackenzie Corcoran, Laurie Williams. Improving Developer Activity Metrics with Issue Tracking Annotations, Proceedings of the 2010 ICSE Workshop on Emerging Trends in Software Metrics. 02-MAY-10, . : ,

Nuo Li, Tao Xie, Nikolai Tillmann, Jonathan Halleux, Wolfram Schulte. Reggae: Automated Test Generation for Programs using Complex Regular Expressions, Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering. 16-NOV-09, . : ,

Suresh Thummalapenta, Tao Xie. Alattin: Mining Alternative Patterns for Detecting Neglected Conditions, Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering. 16-NOV-09, . : ,

Lu Zhang, Hao Zhong, Tao Xie, Hong Mei. Inferring Resource Specifications from Natural Language API Documentation, Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering. 09-NOV-09, . : ,

Hao Zhong, Suresh Thummalapenta, Tao Xie, Lu Zhang, Qing Wang. Mining API Mapping for Language Migration, Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering. 02-MAY-10, . : ,

Michael Gegick, Pete Rotella, Tao Xie. Identifying Security Bug Reports via Text Mining: An Industrial Case Study, 2010 7th IEEE Working Conference on Mining Software Repositories. 02-MAY-12, . : ,

Attila Yavuz, Peng Ning. BAF: An Efficient Publicly Verifiable Secure Audit Logging Scheme for Distributed Systems, Proceedings of the 2009 Annual Computer Security Applications Conference. 07-DEC-09, . : ,

Ahmed Azab, Peng Ning, Can Emre, Xiaolan Zhang. HIMA: A Hypervisor-Based Integrity Measurement Agent, Proceedings of the 2009 Annual Computer Security Applications Conference. 07-DEC-10, . : ,

Jinku Li, Zhi Wang, Xuxian Jiang, Michael Grace, Sina Bahram. Defeating return-oriented rootkits with "Return-Less" kernels, Proceedings of the 5th European conference on Computer systems. 13-APR-10, .

Zhi Wang, Xuxian Jiang. HyperSafe: A Lightweight Approach to Provide Lifetime Hypervisor Control-Flow Integrity, Proceedings of the 2010 IEEE Symposium on Security and Privacy. 16-MAY-10, . : ,

Ahmed M. Azab, Peng Ning, Zhi Wang, Xuxian Jiang, Xiaolan Zhang, Nathan C. Skalsky. HyperSentry: Enabling Stealthy In-context Measurement of Hypervisor Integrity, the 17th ACM conference on Computer and Communications Security (CCS '10). 04-OCT-10, Chicago, Illinois, USA. : ,

Ahmed M. Azab, Peng Ning, Xiaolan Zhang. SICE: A Hardware-Level Strongly Isolated ComputingEnvironment for x86 Multi-core Platforms, the 18th ACM Conference on Computer and Communications Security (CCS '11). 17-OCT-11, . : ,

Wei Wei, Ting Yu. he Design and Enforcement of a Rule-basedConstraint Policy Language for Service Composition, 2010 IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT). 20-AUG-10, . : ,

Jagan Srinivasan, Wei Wei, Xiaosong Ma, Ting Yu. EMFS: Email-based Personal Cloud Storage, 6th IEEE International Conference onNetworking, Architecture, and Storage (NAS '11). 28-JUL-11, . : ,

Juan Du, Xiaohui Gu, Ting Yu. On Verifying Stateful Dataflow Processing Services inLarge-Scale Cloud Systems, 17th ACM Conference on Computer and Communications Security (CCS)(Poster). 05-OCT-10

Chung-Wei Hang, Munindar P. Singh. From Quality to Utility: Adaptive Service Selection Framework, the 8th International Conference on Service-Oriented Computing (ICSOC). 07-DEC-10, . : ,

Meiyappan Nagappan, Aaron Peeler, Mladen Vouk. Modeling Cloud Failure Data: A Case Study of the VirtualComputing Lab, The ICSE 2011 Software Engineering For Cloud Computing Workshop. 21-MAY-11, . : ,

Prasanth Anbalagan, Mladen Vouk. Towards a Bayesian Approach in Modeling the Disclosure of Unique SecurityFaults in Open Source Projects, 21st International Symposium on Software Reliability Engineering 2010 (ISSRE 2010). 01-NOV-11, . : ,

Minh Tran, Mark Etheridge, Tyler Bletsch, Xuxian Jiang, Vincent Freeh, Peng Ning. On the Expressiveness of Return-into-libc Attacks, 14th International Symposium on Recent Advances in Intrusion Detection (RAID 2011. 20-SEP-11, . : ,

Michael Grace, Zhi Wang, Deepa Srinivasan, Jinku L, Xuxian Jiang, Zhenkai Liang, Siarhei Liakh. Transparent Protection of Commodity OS Kernels Using Hardware Virtualization, the 6th International Conference on Security and Privacy in Communication Networks (SecureComm 2010). 07-SEP-10, . : ,

Siarhei Liakh, Michael Grace, Xuxian Jiang. Analyzing and Improving Linux Kernel Memory Protection:A Model Checking Approach, 26th Annual Computer Security Applications Conference (ACSAC 2010). 06-DEC-10, . : ,

Liang Xiao, Huaiyu Dai, Peng Ning. Jamming-Resistant Collaborative Broadcast In WirelessNetworks, Part I: Single-hop Networks, 2010 IEEE GLOBECOM. 06-DEC-11, . : ,

TOTAL: 132

Conference Proceeding

. Jamming-Resistant Collaborative Broadcast In WirelessNetworks, Part II: Multihop Networks, . , . : ,

TOTAL: 1

MS Thesis

Shweta Subramani. A Study of Fedora Security Profile, (05 2014)

Roopak Venkatakrishnan. Redundancy-Based Detection of Security Anomalies in Web-Server Environments, (05 2014)

NIKHIL VIVEK TALPALLIKAR. High-Performance Cloud Computing: VCL Case Study, (04 2012)

KAUSHAL DIDWANIA. Performance Evaluation of Greenplum in VCL Cloud Computing Environment, (03 2012)

TOTAL: 4