Handout 2
Articles for Oral Reports in CompSci 725
v1.01 of 30 July 2013

Clark Thomborson
Giovanni Russello

July 30, 2013

A large fraction (15%) of your marks in CompSci 725 will be awarded for your performance of an oral presentation during a lecture period, during weeks 6 to 12 of the semester. Your presentation will be based on your careful reading and analysis of a professional publication that appears on the list at the end of this handout.

Students will deliver their oral presentations according to a randomly-selected "Order" that is determined at random, by the instructor, at the end of the first week of lectures.

Students with low Order numbers will present before students with higher numbers. Lower-numbered students also have a better chance of being able to present on their first-choice article.

During the first week of lectures, you should be deciding which of the (approximately) two dozen articles on the reference list of this handout you'd like to use as the basis for your oral presentation. You should also make some backup choices, in case you draw a high number.

Deadline: 5pm Monday 29 July. All students should send an email to Clark (`mailto:cthombor@cs.auckland.ac.nz`), indicating their first, second, and third choice of article for their oral presentation. Technical articles can (usually) be uniquely identified by the surname of the first author and the year of publication. For example, your preference list might be "Akhawe 2013, Gong 2012, Jackson 2007".

Clark will digest your emails on Tuesday, and before lecture on Thursday 1 August, he will webpost the first draft of an "Articles to be Presented" handout. This handout will fix the date of your oral presentation; if you have a problem with this date please let him know as soon as possible.

Your attendance at lecture on Monday 5 August will be very important, because on that day we will be finalising the "Articles to be Presented" using the algorithm specified below.

Further information on the oral presentations and written reports will be supplied later in the term. You will also be given some tuition – we understand that, for many of you, this will be your first experience at constructing, and delivering, an oral presentation on a technical subject. We also understand that English is not the native language for many of our students, and we will not be marking you on the fine points of English grammar or spelling. However we do insist that technical words be spelled and used correctly in your oral and written reports. Your technical content must be clearly understandable.

1 Algorithm for Assigning Students to Articles

Please don't worry if you don't understand the algorithm, it's not examinable, and I'll be explaining it as we go!

1.1 Before the first round

All articles are in "Category 0". Formally, $\forall y : Category(y) \equiv 0$.

Category 0	no one assigned to this article.
Category 1	1 student is assigned to this article.
Category 2	2 students are assigned to this article.
Category 3	3 students are assigned to this article.

All students are "Type 0". Formally, $\forall x : Type(x) \equiv 0$.

Type 0	student is not assigned to any article.
Type 1	student is assigned to a category 1 article.
Type 2	student is assigned to a category 2 article.
Type 3	student is assigned to a category 3 article.

All students x have been assigned a unique integer Order(x) in the range $1..N$, where N is the number of students enrolled. The ordering integer determines a student's priority for article selection, and it also defines the order in which they will present their selected article to the class.

The maximum number of articles M has been fixed by the instructors at approximately $0.4N$, so that most articles will have three presenters. For example, if there are $N = 36$ students enrolled, $M = 16$.

Note 1: A student x who enrols at any time after the Order() is fixed will be assigned an ordering integer which is larger than that of any other student in the class.

Note 2: Articles may be added to the recommended list at any time by the instructors. Suggestions from students are welcome, however the suggested article must meet with the instructors' approval.

1.2 Round 1

Students send an email, as defined in the first section of this handout, to Clark (`mailto:cthombor@cs.auckland.ac.nz`), indicating their first, second, and third choice of article for their oral presentation.

For each article y on the list of suggested articles, let $X(y)$ be the set of students who want to present this article. Clark will perform the following computational steps.

```
foreach y in List
  if (|{z: Type(z) > 0}| < M)
    x1 = argmin( Order(X(y)) ); Article(x1) = y;
    Type(x1) = 1; Category(y) = 1;
  endif
  if (|{z: Type(z) > 0}| <= M) and (X>1)
    x2 = argmin( Order(X(y)/x1) ); Article(x2) = y;
    Type(x2) = 2; Category(y) = 2;
  endif
  if (|{z: Type(z) > 0}| <= M) and (X>2)
    x3 = argmin( Order(X(y)/{x1,x2}) ); Article(x3) = y;
    Type(x3) = 3; Category(y) = 3;
  endif
endfor
```

Postconditions for Round 1:

1. $\forall x : \text{Type}(x) \in \{0, 1, 2, 3\}$
2. $\forall p : \text{Category}(p) \in \{0, 1, 2, 3\}$
3. No more than M different articles will be presented: $|\{p : \text{Category}(p) > 0\}| \leq M$

1.3 Round 2: In class, Monday 5 August

Each student x, in assigned Order starting with student #1, must choose one of the following actions:

1. If student x has made a selection (i.e. if $\text{Type}(x) \equiv 1$), they may "hold" this article.
2. Student x may select any article y in category 1 or 2. Do:

   ```
   Category(y)++; Article(x) = y; Type(x) = Category(y);
   ```

3. If M articles haven't already been selected ($(|\{p : \text{Category} > 0\}| < M)$), then student x may select any article y in category 0. Do:

```
Category(y) = 1; Article(x) = y; Type(x) = 1;
```

Postcondition: all students who attended this lecture have chosen an article.

1.4 Round 3: In class, Monday 5 August

Students are allowed to "swap" their article with other students, in a controlled fashion. Also any students who haven't chosen an article must do so.

Each student x, in assigned order starting with student #1, must choose one of the following actions.

1. If (Type(x) $\equiv 0$), this student must select a article y in this round, either by choosing one of the Round-2 actions listed above or by choosing one of the actions below.

2. Student x may select a article y in Category 3, but *only if* one of the students z who is currently assigned to y is

   (a) willing to move to a different article w in category 1 or 2, or

   (b) if fewer than M articles have been selected and student z is willing to move to a article w in category 0.

Note 1: if more than one student is willing to move from y, then the student with the lowest number is the "volunteer" z.

Note 2: once a student (even one with the highest order #) is assigned a article, they cannot be forced to "move" to a different article.

Postconditions for Rounds > 2:

1. All registered students must present an article: $\forall x : \text{Type}(x) \in \{1, 2, 3\}$.

2. No more than M articles will be presented: $|\{p : \text{Category}(p) > 0\}| \leq M$

Termination Condition: Rounds will continue (to a maximum of 10) until a fixed-point is reached, i.e. until a Round makes no changes to the Article() assignments.

References

[1] D. Akhawe and A. P. Felt, "Alice in warningland: A large-scale field study of browser security warning effectiveness," in *USENIX Security Symposium*, 2013, author's preprint. [Online]. Available: http://www.cs.berkeley.edu/~devdatta/papers/alice-in-warningland.pdf

Abstract. We empirically assess whether browser security warnings are as ineffective as suggested by popular opinion and previous literature. We used Mozilla Firefox and Google Chrome's in-browser telemetry to observe over 25 million warning impressions in situ. During our field study, users continued through a tenth of Mozilla Firefox's malware and phishing warnings, a quarter of Google Chrome's malware and phishing warnings, and a third of Mozilla Firefox's SSL warnings. This demonstrates that security warnings can be effective in practice; security experts and system architects should not dismiss the goal of communicating security information to end users. We also find that user behavior varies across warnings. In contrast to the other warnings, users continued through 70.2% of Google Chrome's SSL warnings. This indicates that the user experience of a warning can have a significant impact on user behavior. Based on our findings, we make recommendations for warning designers and researchers.

[2] G. Ateniese, R. Burns, R. Curtmola, J. Herring, L. Kissner, Z. Peterson, and D. Song, "Provable data possession at untrusted stores," in *Proceedings of the 14th ACM conference on Computer and communications security*, ser. CCS '07. New York, NY, USA: ACM, 2007, pp. 598–609. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/1315245.1315318

Abstract. We introduce a model for provable data possession (PDP) that allows a client that has stored data at an untrusted server to verify that the server possesses the original data without retrieving it. The model generates probabilistic proofs of possession by sampling random sets of blocks from the server, which drastically reduces I/O costs. The client maintains a constant amount of metadata to verify the proof. The challenge/response protocol transmits a small, constant amount of data, which minimizes network communication. Thus, the PDP model for remote data checking supports large data sets in widely-distributed storage systems. We present two provably-secure PDP schemes that are more efficient than previous solutions, even when compared with schemes that achieve weaker guarantees. In particular, the overhead at the server is low (or even constant), as opposed to linear in the size of the data. Experiments using our implementation verify the practicality of PDP and reveal that the performance

of PDP is bounded by disk I/O and not by cryptographic computation.

[3] A. Birgisson, M. Dhawan, U. Erlingsson, V. Ganapathy, and L. Iftode, "Enforcing authorization policies using transactional memory introspection," in *Proceedings of the 15th ACM conference on Computer and communications security*, ser. CCS '08. New York, NY, USA: ACM, 2008, pp. 223–234. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/1455770.1455800

Abstract. Correct enforcement of authorization policies is a difficult task, especially for multi-threaded software. Even in carefully-reviewed code, unauthorized access may be possible in subtle corner cases. We introduce Transactional Memory Introspection (TMI), a novel reference monitor architecture that builds on Software Transactional Memory—a new, attractive alternative for writing correct, multi-threaded software.

TMI facilitates correct security enforcement by simplifying how the reference monitor integrates with software functionality. TMI can ensure complete mediation of security-relevant operations, eliminate race conditions related to security checks, and simplify handling of authorization failures. We present the design and implementation of a TMI-based reference monitor and experiment with its use in enforcing authorization policies on four significant servers. Our experiments confirm the benefits of the TMI architecture and show that it imposes an acceptable runtime overhead.

[4] L. Brandimarte, A. Acquisti, and G. Loewenstein, "Misplaced confidences: Privacy and the control paradox," *Social Psychological and Personality Science*, vol. 4, no. 3, pp. 340–347, 2013. [Online]. Available: http://spp.sagepub.com.ezproxy.auckland.ac.nz/content/4/3/340.short

Abstract. We test the hypothesis that increasing individuals' perceived control over the release and access of private information–even information that allows them to be personally identified–will increase their willingness to disclose sensitive information. If their willingness to divulge increases sufficiently, such an increase in control can, paradoxically, end up leaving them more vulnerable. Our findings highlight how, if people respond in a sufficiently offsetting fashion, technologies designed to protect them can end up exacerbating the risks they face.

[5] S. Butt, V. Ganapathy, A. Baliga, and M. Christodorescu, "Monitoring data structures using hardware transactional memory," in *Runtime Verification*, ser. Lecture Notes in Computer Science, S. Khurshid and K. Sen, Eds. Springer Berlin Heidelberg, 2012, vol. 7186, pp. 345–359. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1007/978-3-642-29860-8_26

Abstract. The robustness of software systems is adversely affected by programming errors and security exploits that corrupt heap data structures. In this paper, we present the design and implementation of TxMon, a system to detect such data structure corruptions. TxMon leverages the concurrency control machinery implemented by hardware transactional memory (HTM) systems to additionally enforce programmer-specified consistency properties on data structures at runtime. We implemented a prototype version of TxMon using an HTM system (LogTM-SE) and studied the feasibility of applying TxMon to enforce data structure consistency properties on several benchmarks. Our experiments show that TxMon is effective at monitoring data structure properties, imposing tolerable runtime performance overheads.

[6] S. Butt, H. A. Lagar-Cavilla, A. Srivastava, and V. Ganapathy, "Self-service cloud computing," in *Proceedings of the 2012 ACM conference on Computer and communications security*, ser. CCS '12. New York, NY, USA: ACM, 2012, pp. 253–264. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/2382196.2382226

Abstract. Modern cloud computing infrastructures use virtual machine monitors (VMMs) that often include a large and complex administrative domain with privileges to inspect client VM state. Attacks against or misuse of the administrative domain can compromise client security and privacy. Moreover, these VMMs provide clients inflexible control over their own VMs, as a result of which clients have to rely on the cloud provider to deploy useful services, such as VM introspection-based security tools.

We introduce a new self-service cloud (SSC) computing model that addresses these two shortcomings. SSC splits administrative privileges between a system-wide domain and per-client administrative domains. Each client can manage and perform privileged system tasks on its own VMs, thereby providing flexibility. The system-wide administrative domain cannot inspect the code, data or computation of client VMs, thereby ensuring security and privacy. SSC also allows

providers and clients to establish mutually trusted services that can check regulatory compliance while respecting client privacy. We have implemented SSC by modifying the Xen hypervisor. We demonstrate its utility by building user domains to perform privileged tasks such as memory introspection, storage intrusion detection, and anomaly detection.

[7] S. Chakradeo, B. Reaves, P. Traynor, and W. Enck, "MAST: Triage for market-scale mobile malware analysis," in *Sixth ACM Conference on Security and Privacy in Wireless and Mobile Networks*, L. Buttyán, A.-R. Sadeghi, and M. Gruteser, Eds. ACM, 2013, pp. 13–24. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/2462096.2462100

Abstract. Malware is a pressing concern for mobile application market operators. While current mitigation techniques are keeping pace with the relatively infrequent presence of malicious code, the rapidly increasing rate of application development makes manual and resourceintensive automated analysis costly at market-scale. To address this resource imbalance, we present the Mobile Application Security Triage (MAST) architecture, a tool that helps to direct scarce malware analysis resources towards the applications with the greatest potential to exhibit malicious behavior. MAST analyzes attributes extracted from just the application package using Multiple Correspondence Analysis (MCA), a statistical method that measures the correlation between multiple categorical (i.e., qualitative) data. We train MAST using over 15,000 applications from Google Play and a dataset of 732 known-malicious applications. We then use MAST to perform triage on three third-party markets of different size and malware composition 36,710 applications in total. Our experiments show that MAST is both effective and performant. Using MAST ordered ranking, malware-analysis tools can find 95% of malware at the cost of analyzing 13% of the non-malicious applications on average across multiple markets, and MAST triage processes markets in less than a quarter of the time required to perform signature detection. More importantly, we show that successful triage can dramatically reduce the costs of removing malicious applications from markets.

[8] S. Creese, M. Goldsmith, J. Nurse, and E. Phillips, "A data-reachability model for elucidating privacy and security risks related to the use of online social networks," in *2012 IEEE 11th International*

Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), 2012, pp. 1124–1131. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1109/TrustCom.2012.22

Abstract. Privacy and security within Online Social Networks (OSNs) has become a major concern over recent years. As individuals continue to actively use and engage with these mediums, one of the key questions that arises pertains to what unknown risks users face as a result of unchecked publishing and sharing of content and information in this space. There are numerous tools and methods under development that claim to facilitate the extraction of specific classes of personal data from online sources, either directly or through correlation across a range of inputs. In this paper we present a model which specifically aims to understand the potential risks faced should all of these tools and methods be accessible to a malicious entity. The model enables easy and direct capture of the data extraction methods through the encoding of a data-reachability matrix for which each row represents an inference or data-derivation step. Specifically, the model elucidates potential linkages between data typically exposed on social-media and networking sites, and other potentially sensitive data which may prove to be damaging in the hands of malicious parties, i.e., fraudsters, stalkers and other online and offline criminals. In essence, we view this work as a key method by which we might make cyber risk more tangible to users of OSNs.

[9] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, 2013. [Online]. Available: http://www.nature.com/srep/2013/130325/srep01376/pdf/srep01376.pdf

Abstract. We study fifteen months of human mobility data for one and a half million individuals and find that human mobility traces are highly unique. In fact, in a dataset where the location of an individual is specified hourly, and with a spatial resolution equal to that given by the carrier's antennas, four spatio-temporal points are enough to uniquely identify 95% of the individuals. We coarsen the data spatially and temporally to find a formula for the uniqueness of human mobility traces given their resolution and the available outside information. This formula shows that the uniqueness of mobility traces decays approximately as the 1/10 power of their resolution. Hence, even coarse datasets provide little

anonymity. These findings represent fundamental constraints to an individual's privacy and have important implications for the design of frameworks and institutions dedicated to protect the privacy of individuals.

[10] N. Z. Gong, W. Xu, L. Huang, P. Mittal, E. Stefanov, V. Sekar, and D. Song, "Evolution of social-attribute networks: Measurements, modeling, and implications using Google+," in *Proceedings of the 2012 ACM Conference on Internet Measurement*, ser. IMC '12. New York, NY, USA: ACM, 2012, pp. 131–144. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/2398776.2398792

Abstract. Understanding social network structure and evolution has important implications for many aspects of network and system design including provisioning, bootstrapping trust and reputation systems via social networks, and defenses against Sybil attacks. Several recent results suggest that augmenting the social network structure with user attributes (e.g., location, employer, communities of interest) can provide a more fine-grained understanding of social networks. However, there have been few studies to provide a systematic understanding of these effects at scale.

We bridge this gap using a unique dataset collected as the Google+ social network grew over time since its release in late June 2011. We observe novel phenomena with respect to both standard social network metrics and new attribute-related metrics (that we define). We also observe interesting evolutionary patterns as Google+ went from a bootstrap phase to a steady invitation-only stage before a public release.

Based on our empirical observations, we develop a new generative model to jointly reproduce the social structure and the node attributes. Using theoretical analysis and empirical evaluations, we show that our model can accurately reproduce the social and attribute structure of real social networks. We also demonstrate that our model provides more accurate predictions for practical application contexts.

[11] A. Houmansadr, C. Brubaker, and V. Shmatikov, "The parrot is dead: Observing unobservable network communications," *2013 IEEE Symposium on Security and Privacy*, vol. 0, pp. 65–79, 2013. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1109/SP.2013.14

Abstract. In response to the growing popularity of Tor and other censorship circumvention systems, censors in non-

democratic countries have increased their technical capabilities and can now recognize and block network traffic generated by these systems on a nationwide scale. New censorship-resistant communication systems such as Skype Morph, Stego Torus, and Censor Spoofer aim to evade censors' observations by imitating common protocols like Skype and HTTP.

We demonstrate that these systems completely fail to achieve unobservability. Even a very weak, local censor can easily distinguish their traffic from the imitated protocols. We show dozens of passive and active methods that recognize even a single imitated session, without any need to correlate multiple network flows or perform sophisticated traffic analysis.

We enumerate the requirements that a censorship-resistant system must satisfy to successfully mimic another protocol and conclude that "unobservability by imitation" is a fundamentally flawed approach. We then present our recommendations for the design of unobservable communication systems.

[12] C. Jackson, D. Simon, D. Tan, and A. Barth, "An evaluation of extended validation and picture-in-picture phishing attacks," in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, S. Dietrich and R. Dhamija, Eds. Springer Berlin / Heidelberg, 2007, vol. 4886, pp. 281–293. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1007/978-3-540-77366-5_27

Abstract. In this usability study of phishing attacks and browser anti-phishing defenses, 27 users each classified 12 web sites as fraudulent or legitimate. By dividing these users into three groups, our controlled study measured both the effect of extended validation certificates that appear only at legitimate sites and the effect of reading a help file about security features in Internet Explorer 7. Across all groups, we found that picture-in-picture attacks showing a fake browser window were as effective as the best other phishing technique, the homograph attack. Extended validation did not help users identify either attack. Additionally, reading the help file made users more likely to classify both real and fake web sites as legitimate when the phishing warning did not appear.

[13] S. Jana, D. Porter, and V. Shmatikov, "TxBox: Building secure, efficient sandboxes with system transactions," in *2011 IEEE Symposium on Security and Privacy (SP)*, 2011, pp. 329–344. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1109/SP.2011.33

Abstract. TxBox is a new system for sandboxing untrusted applications. It speculatively executes the application in a system transaction, allowing security checks to be parallelized and yielding significant performance gains for techniques such as on-access anti-virus scanning. TxBox is not vulnerable to TOCTTOU attacks and incorrect mirroring of kernel state. Furthermore, TxBox supports automatic recovery: if a violation is detected, the sandboxed program is terminated and all of its effects on the host are rolled back. This enables effective enforcement of security policies that span multiple system calls.

[14] S. Kim, M. Z. Lee, A. M. Dunn, O. S. Hofmann, X. Wang, E. Witchel, and D. E. Porter, "Improving server applications with system transactions," in *Proceedings of the 7th ACM European Conference on Computer Systems (EuroSys)*, Bern, Switzerland, April 2012. [Online]. Available: http://dl.acm.org.ezproxy.auckland.ac.nz/citation.cfm?id=2168839

Abstract. Server applications must process requests as quickly as possible. Because some requests depend on earlier requests, there is often a tension between increasing throughput and maintaining the proper semantics for dependent requests. Operating system transactions make it easier to write reliable, high-throughput server applications because they allow the application to execute non-interfering requests in parallel, even if the requests operate on OS state, such as file data.

By changing less than 200 lines of application code, we improve performance of a replicated Byzantine Fault Tolerant (BFT) system by up to 88% using server-side speculation, and we improve concurrent performance up to 80% for an IMAP email server by changing only 40 lines. Achieving these results requires substantial enhancements to system transactions, including the ability to pause and resume transactions, and an API to commit transactions in a pre-defined order.

[15] T. Knall, A. Tauber, T. Zefferer, B. Zwattendorfer, A. Axfjord, and H. Bjarnason, "Secure and privacy-preserving cross-border authentication: The STORK pilot 'SaferChat'," in *Electronic Government and the Information Systems Perspective*, ser. Lecture Notes in Computer Science, K. N. Andersen, E. Francesconi, A. Grönlund, and T. M. Engers, Eds. Springer Berlin Heidelberg, 2011, vol. 6866, pp. 94–106. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1007/978-3-642-22961-9_8

Abstract. Secure user authentication, provision of identity attributes, privacy preservation, and cross-border applicability are key requirements of security and privacy sensitive ICT based services. The EU large scale pilot STORK provides a European cross-border authentication framework that satisfies these requirements by establishing interoperability between existing national eID infrastructures. To allow for privacy preservation, the developed framework supports the provision of partial identity information and pseudonymization. In this paper we present the pilot application SaferChat that has been developed to evaluate and demonstrate the functionality of the STORK authentication framework. SaferChat makes use of age claim based authentication mechanisms that allow for an online environment where kids and teenagers are able to communicate with their peers in a safe way. We first identify relevant prerequisites for the SaferChat pilot application and then give an introduction to the basic architecture of the STORK authentication framework. We finally show how this framework has been integrated into the SaferChat pilot application to meet the identified requirements and to implement a secure and privacy preserving cross-border user authentication mechanism.

[16] A. Kurtz, F. Freiling, and D. Metz, "Usability vs. security: The everlasting trade-off in the context of Apple iOS mobile hotspots," University of Erlangen, Department of Computer Science, Tech. Rep. CS-2013-02 (author's preprint), June 2013, 10 pp. [Online]. Available: https://www1.cs.fau.de/filepool/projects/hotspot/hotspot.pdf

Abstract. Passwords have to be secure and usable at the same time, a trade-off that is long known. There are many approaches to avoid this trade-off, e.g., to advice users on generating strong passwords and to reject user passwords that are weak. The same usability/security trade-off arises in scenarios where passwords are generated by machines but exchanged by humans, as is the case in pre-shared key (PSK) authentication. We investigate this trade-off by analyzing the PSK authentication method used by Apple iOS to set up a secure WPA2 connection when using an iPhone as a Wi-Fi mobile hotspot. We show that Apple iOS generates weak default passwords which makes the mobile hotspot feature of Apple iOS susceptible to brute force attacks on the WPA2 handshake. More precisely, we observed that the generation of default passwords is based on a word list, of which only 1.842

entries are taken into consideration. In addition, the process of selecting words from that word list is not random at all, resulting in a skewed frequency distribution and the possibility to compromise a hotspot connection in less than 50 seconds. Spot tests show that other mobile platforms are also affected by similar problems. We conclude that more care should be taken to create secure passwords even in PSK scenarios.

[17] G. Miller and L. Williams, "Personas: Moving beyond role-based requirements engineering," Internet manuscript, circa 2006, 10 pp. [Online]. Available: http://agile.csc.ncsu.edu/SEMaterials/Personas.pdf

Abstract. A primary vehicle for understanding the user in the context of the requirements for a system has been the role. For example, the role is captured through the use of actors in the use case diagram and use case descriptions. Recently, personas have been used in conjunction with scenarios in participatory design to go deeper into examining the different types of people who could play a role. A persona is an archetype of a fictional user representing a specific group of typical users. This paper expands the use of personas to scenario-based requirements engineering. Personas and scenarios are being used together for specifying requirements at Microsoft. The result of this combination has been a more comprehensive understanding of the target customers' behaviors to drive and refine our scenarios and subsequently our product development.

[18] M. O. Rabin, "Efficient dispersal of information for security, load balancing, and fault tolerance," *J. ACM*, vol. 36, no. 2, pp. 335–348, Apr. 1989. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1145/62044.62050

Abstract. An Information Dispersal Algorithm (IDA) is developed that breaks a file F of length $L = |F|$ into n pieces F_i, $l \leq i \leq n$, each of length $|F_i| = L/m$, so that every m pieces suffice for reconstructing F. Dispersal and reconstruction are computationally efficient. The sum of the lengths $|F_i|$ is $(n/m)L$. Since n/m can be chosen to be close to 1, the IDA is space efficient. IDA has numerous applications to secure and reliable storage of information in computer networks and even on single disks, to fault-tolerant and efficient transmission of information in networks, and to communications between processors in parallel computers. For the latter prob-

lem provably time-efficient and highly fault-tolerant routing on the n-cube is achieved, using just constant size buffers.

[19] S. Ransbotham and S. Mitra, "Choice and chance: A conceptual model of paths to information security compromise," *Information Systems Research*, vol. 20, no. 1, pp. 121–139, 2009. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1287/isre.1080.0174

Abstract. No longer the exclusive domain of technology experts, information security is now a management issue. Through a grounded approach using interviews, observations, and secondary data, we advance a model of the information security compromise process from the perspective of the attacked organization. We distinguish between deliberate and opportunistic paths of compromise through the Internet, labeled choice and chance, and include the role of countermeasures, the Internet presence of the firm, and the attractiveness of the firm for information security compromise. Further, using one year of alert data from intrusion detection devices, we find empirical support for the key contributions of the model. We discuss the implications of the model for the emerging research stream on information security in the information systems literature.

[20] C. J. Rossbach, J. Currey, M. Silberstein, B. Ray, and E. Witchel, "PTask: Operating system abstractions to manage gpus as compute devices," in *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, October 2011. [Online]. Available: http://dl.acm.org.ezproxy.auckland.ac.nz/citation.cfm?id=2043579

Abstract. We propose a new set of OS abstractions to support GPUs and other accelerator devices as first class computing resources. These new abstractions, collectively called the *PTask API*, support a dataflow programming model. Because a PTask graph consists of OS-managed objects, the kernel has sufficient visibility and control to provide systemwide guarantees like fairness and performance isolation, and can streamline data movement in ways that are impossible under current GPU programming models.

Our experience developing the PTask API, along with a gestural interface on Windows 7 and a FUSE-based encrypted file system on Linux show that the PTask API can provide important systemwide guarantees where there were previously none, and can enable significant performance improvements,

for example gaining a 5 improvement in maximum throughput for the gestural interface.

[21] M. I. Sharif, A. Lanzi, J. T. Giffin, and W. Lee, "Impeding malware analysis using conditional code obfuscation," in *Proceedings of the Network and Distributed System Security Symposium*. The Internet Society, 2008. [Online]. Available: http://www.isoc.org/isoc/conferences/ndss/08/papers/19_impeding_malware_analysis.pdf

Abstract. Malware programs that incorporate trigger-based behavior initiate malicious activities based on conditions satisfied only by specific inputs. State-of-the-art malware analyzers discover code guarded by triggers via multiple path exploration, symbolic execution, or forced conditional execution, all without knowing the trigger inputs. We present a malware obfuscation technique that automatically conceals specific trigger-based behavior from these malware analyzers. Our technique automatically transforms a program by encrypting code that is conditionally dependent on an input value with a key derived from the input and then removing the key from the program. We have implemented a compiler-level tool that takes a malware source program and automatically generates an obfuscated binary. Experiments on various existing malware samples show that our tool can hide a significant portion of trigger based code. We provide insight into the strengths, weaknesses, and possible ways to strengthen current analysis approaches in order to defeat this malware obfuscation technique.

[22] D. Shin and R. Lopes, "An empirical study of visual security cues to prevent the SSLstripping attack," in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC '11. New York, NY, USA: ACM, 2011, pp. 287–296. [Online]. Available: http://doi.acm.org.ezproxy.auckland.ac.nz/10.1145/2076732.2076773

Abstract. One of the latest attacks on secure socket layer (SSL), called the SSLstripping attack, was reported at the Blackhat conference in 2009. As a type of man-in-the-middle (MITM) attack, it has the potential to affect tens of millions of users of popular online social networking and financial websites protected by SSL. Interestingly, the attack exploits users' browsing habits, rather than a technical flaw in the protocol, to defeat the SSL security. In this paper we present a novel approach to addressing this attack by using visually augmented security. Specifically, motivated by typi-

cal traffic lights, we introduce a set of visual cues aimed at thwarting the attack. The visual cues, called security status light (SSLight), can be used to help users make better, more informed decisions when their sensitive information need to be submitted to the websites. A user study was conducted to investigate the effectiveness of our scheme, and its results show that our approach is more promising than the traditional pop-up method adopted by major web browsers.

[23] A. van Overeem and J. van Oosten, "Towards a pan European e-ID interoperability infrastructure," in *42nd Hawaii International Conference on System Sciences (HICSS '09)*. IEEE Computer Society, Jan. 2009, pp. 1–10. [Online]. Available: http://dx.doi.org.ezproxy.auckland.ac.nz/10.1109/HICSS.2009.466

Abstract. The proliferation of e-Services in most European Countries has been favorable to the emergence of common identity providers and national identity management infrastructures in these countries. The STORK project aims to interconnect all of these identity management infra-structures to form a Pan-European federated e-Identity space. In this paper we show that due to two different identity concepts in use by the European countries, this objective is a far from trivial challenge. Based on our analysis we present two scenarios: homogeneous interoperability for countries with alike identity concepts and heterogeneous interoperability for countries with different identity concepts. For the latter case we present three solution directions to overcome technical limitations and challenges. The STORK project is co-funded by the European Union and will deliver real solutions by implementing five demo projects.