ELYSIUMPRO

INSPIRING THE LEADING EDGE TECHNOLOGIES

ELYSIUMPRO

A UNIT OF ELYSIUM GROUPS

FINAL YEAR PROJECTS

DEPENDABLE AND SECURE COMPUTING 2017-2018

TITLES WITH ABSTRACTS

CALL US @ 994-479-3398, (0452) 4390702

18 Years of Experience • Automated Services • 24/7 Help Desk Support

Advanced Technologies And Tools • Legitimate Members Of all Journals

ELYSIUM GROUPS

ELYSIUMPRO

A UNIT OF ELYSIUM GROUPS

www.elysiumpro.in

/elysiumtechnologiespvtltd /elysiumind /elysium-technologies-pvt

#227, Church Road, Anna Nagar, Madurai, Tamil Nadu, India

CALL US @ 994-479-3398, (0452) 4390702

ELYSIUM GROUPS

ETPL DSC - 001	A Novel Class of Robust Covert Channels Using Out-of-Order Packets

Covert channels are usually used to circumvent security policies and allow information leakage without being observed. In this paper, we propose a novel covert channel technique using the packet reordering phenomenon as a host for carrying secret communications. Packet reordering is a common phenomenon on the Internet. Moreover, it is handled transparently from the user and application-level processes. This makes it an attractive medium to exploit for sending hidden signals to receivers by dynamically manipulating packet order in a network flow. In our approach, specific permutations of successive packets are selected to enhance the reliability of the channel, while the frequency distribution of their usage is tuned to increase stealthiness by imitating real Internet traffic. It is very expensive for the adversary to discover the covert channel due to the tremendous overhead to buffer and sort the packets among huge amount of background traffic. A simple tool is implemented to demonstrate this new channel. We studied extensively the robustness and capabilities of our proposed channel using both simulation and experimentation over large varieties of traffic characteristics. The reliability and capacity of this technique have shown promising results. We also investigated a practical mechanism for distorting and potentially preventing similar novel channels.

ETPL DSC - 002	An Efficient Lattice Based Multi-Stage Secret Sharing Scheme

In this paper, we construct a lattice based (t, n) threshold multi-stage secret sharing (MSSS) scheme according to Ajtai's construction for one-way functions. In an MSSS scheme, the authorized subsets of participants can recover a subset of secrets at each stage while other secrets remain undisclosed. In this paper, each secret is a vector from a t-dimensional lattice and the basis of each lattice is kept private. A t-subset of n participants can recover the secret(s) using their assigned shares. Using a lattice based one-way function, even after some secrets are revealed, the computational security of the unrecovered secrets is provided against quantum computers. The scheme is multi-use in the sense that to share a new set of secrets, it is sufficient to renew some public information such that a new share distribution is no longer required. Furthermore, the scheme is verifiable meaning that the participants can verify the shares received from the dealer and the recovered secrets from the combiner, using public information.

ETPL DSC - 003	Defending Against Web Application Attacks: Approaches, Challenges and Implications

Some of the most dangerous web attacks, such as Cross-Site Scripting and SQL injection, exploit vulnerabilities in web applications that may accept and process data of uncertain origin without proper validation or filtering, allowing the injection and execution of dynamic or domain-specific language code. These attacks have been constantly topping the lists of various security bulletin providers despite the numerous countermeasures that have been proposed over the past 15 years. In this paper, we provide an analysis on various defense mechanisms against web code injection attacks. We propose a model that highlights the key weaknesses enabling these attacks, and that provides a common perspective for studying the available defenses. We then categorize and analyze a set of 41 previously proposed defenses based on their accuracy, performance, deployment, security, and availability characteristics. Detection accuracy is of particular importance, as our findings show that many defense mechanisms have been tested in a poor manner. In addition, we observe that some mechanisms can be bypassed by attackers with knowledge of how the mechanisms work. Finally, we discuss the results of our analysis, with emphasis on factors that may hinder the widespread adoption of defenses in practice.

ETPL DSC - 004	Design and Implementation of the Ascend Secure Processor

This paper presents hardware implementations of the Ascend secure processor, prototyped on an FPGA and taped out in a 32 nm SOI process. Ascend prevents information leakage over a processor's digital I/O pins — in particular, the processor's requests to external memory — and certifies the program's execution by integrity-verifying the external memory. In secure processor design, encrypting main memory is not sufficient for security because where and when memory is accessed reveals secret information. To this end, Ascend is equipped with a hardware Oblivious RAM (ORAM) controller, which obfuscates the address bus by reshuffling memory as it is accessed. To our knowledge, Ascend is the first prototyping of ORAM in custom silicon. Ascend has also been carefully engineered to ensure its timing behaviors are independent of user private data. We describe our open-source FPGA prototype and the different design considerations that were made when optimizing for an FPGA vs. the ASIC. In 32 nm silicon, all security components combined (the ORAM controller, which includes 12 AES rounds and one SHA-3 hash unit) impose a moderate area overhead of 1 mm2. Post tape-out, the Ascend chip has been successfully tested at 500 MHz.

ETPL DSC - 005	Efficient and Privacy-Preserving Min and k th Min Computations in Mobile Sensing Systems

Protecting the privacy of mobile phone user participants is extremely important for mobile phone sensing applications. In this paper, we study how an aggregator can expeditiously compute the minimum value or the kth minimum value of all users' data without knowing them. We construct two secure protocols using probabilistic coding schemes and a cipher system that allows homomorphic bitwise XOR computations for our problems. Following the standard cryptographic security definition in the semi-honest model, we formally prove our protocols' security. The protocols proposed by us can support time-series data and need not to assume the aggregator is trusted. Moreover, different from existing protocols that are based on secure arithmetic sum computations, our protocols are based on secure bitwise XOR computations, thus are more efficient.

ETPL DSC - 006	Detecting and Preventing Kernel Rootkit Attacks with Bus Snooping

To protect the integrity of operating system kernels, we present Vigilare system, a kernel integrity monitor that is architected to snoop the bus traffic of the host system from a separate independent hardware. This snoop-based monitoring enabled by the Vigilare system, overcomes the limitations of the snapshot-based monitoring employed in previous kernel integrity monitoring solutions. Being based on inspecting snapshots collected over a certain interval, the previous hardware-based monitoring solutions cannot detect transient attacks that can occur in between snapshots, and cannot protect the kernel against permanent damage. We implemented three prototypes of the Vigilare system by adding Snooper hardware connections module to the host system for bus snooping, and a snapshot-based monitor to be compared with, in order to evaluate the benefit of snoop-based monitoring. The prototypes of Vigilare system detected all the transient attacks and the second one protected the kernel with negligible performance degradation while the snapshot-based monitor could not detect all the attacks and induced considerable performance degradation as much as 10 percent in our tuned STREAM benchmark test.

ETPL DSC - 007	Exact Inference Techniques for the Analysis of Bayesian Attack Graphs

Attack graphs are a powerful tool for security risk assessment by analysing network vulnerabilities and the paths attackers can use to compromise network resources. The uncertainty about the attacker's behaviour makes Bayesian networks suitable to model attack graphs to perform static and dynamic analysis. Previous approaches have focused on the formalization of attack graphs into a Bayesian model rather than proposing mechanisms for their analysis. In this paper we propose to use efficient algorithms to make exact inference in Bayesian attack graphs, enabling the static and dynamic network risk assessments. To support the validity of our approach we have performed an extensive experimental evaluation on synthetic Bayesian attack graphs with different topologies, showing the computational advantages in terms of time and memory use of the proposed techniques when compared to existing approaches.

ETPL DSC - 008	Fuzzy Identity-Based Data Integrity Auditing for Reliable Cloud Storage Systems

Data integrity, a core security issue in reliable cloud storage, has received much attention. Data auditing protocols enable a verifier to efficiently check the integrity of the outsourced data without downloading the data. A key research challenge associated with existing designs of data auditing protocols is the complexity in key management. In this paper, we seek to address the complex key management challenge in cloud data integrity checking by introducing fuzzy identity-based auditing, the first in such an approach, to the best of our knowledge. More specifically, we present the primitive of fuzzy identity-based data auditing, where a user's identity can be viewed as a set of descriptive attributes. We formalize the system model and the security model for this new primitive. We then present a concrete construction of fuzzy identity-based auditing protocol by utilizing biometrics as the fuzzy identity. The new protocol offers the property of error-tolerance, namely, it binds with private key to one identity which can be used to verify the correctness of a response generated with another identity, if and only if both identities are sufficiently close. We prove the security of our protocol based on the computational Diffie-Hellman assumption and the discrete logarithm assumption in the selective-ID security model. Finally, we develop a prototype implementation of the protocol which demonstrates the practicality of the proposal.

ETPL DSC - 009	High Performance and High Scalable Packet Classification Algorithm for Network Security Systems

Packet classification is a core function in network and security systems; hence, hardware-based solutions, such as packet classification accelerator chips or Ternary Content Addressable Memory (T-CAM), have been widely adopted for high-performance systems. With the rapid improvement of general hardware architectures and growing popularity of multi-core multi-threaded processors, software-based packet classification algorithms are attracting considerable attention, owing to their high flexibility in satisfying various industrial requirements for security and network systems. For high classification speed, these algorithms internally use large tables, whose size increases exponentially with the ruleset size; consequently, they cannot be used with a large rulesets. To overcome this problem, we propose a new software-based packet classification algorithm that simultaneously supports high scalability and fast classification performance by merging partition decision trees in a search table. While most partitioning-based packet classification algorithms show good scalability at the cost of low classification speed, our algorithm shows very high classification speed, irrespective of the number of rules, with small tables and short table building time. Our test results confirm that the proposed algorithm enables network and security systems to support heavy traffic in the most effective manner.

ETPL DSC - 010	Leveraging Compression-based Graph Mining for Behavior-based Malware Detection

Behavior-based detection approaches commonly address the threat of statically obfuscated malware. Such approaches often use graphs to represent process or system behavior and typically employ frequency-based graph mining techniques to extract characteristic patterns from collections of malware graphs. Recent studies in the molecule mining domain suggest that frequency-based graph mining algorithms often perform sub-optimally in finding highly discriminating patterns. We propose a novel malware detection approach that uses so-called compression-based mining on quantitative data flow graphs to derive highly accurate detection models. Our evaluation on a large and diverse malware set shows that our approach outperforms frequency-based detection models in terms of detection effectiveness by more than 600%.

ETPL DSC - 011	My Privacy My Decision: Control of Photo Sharing on Online Social Networks

Photo sharing is an attractive feature which popularizes online social networks (OSNs). Unfortunately, it may leak users' privacy if they are allowed to post, comment, and tag a photo freely. In this paper, we attempt to address this issue and study the scenario when a user shares a photo containing individuals other than himself/herself (termed co-photo for short). To prevent possible privacy leakage of a photo, we design a mechanism to enable each individual in a photo be aware of the posting activity and participate in the decision making on the photo posting. For this purpose, we need an efficient facial recognition (FR) system that can recognize everyone in the photo. However, more demanding privacy setting may limit the number of the photos publicly available to train the FR system. To deal with this dilemma, our mechanism attempts to utilize users' private photos to design a personalized FR system specifically trained to differentiate possible photo co-owners without leaking their privacy. We also develop a distributed consensus-based method to reduce the computational complexity and protect the private training set. We show that our system is superior to other possible approaches in terms of recognition ratio and efficiency. Our mechanism is implemented as a proof of concept Android application on Facebook's platform.

ETPL DSC - 012	Optimized Identity-Based Encryption from Bilinear Pairing for Lightweight Devices

Lightweight devices such as smart cards and RFID tags have a very limited hardware resource, which could be too weak to cope with asymmetric-key cryptography. It would be desirable if the cryptographic algorithm could be optimized in order to better use hardware resources. In this paper, we demonstrate how identity-based encryption algorithms from bilinear pairing can be optimized so that hardware resources can be saved. We notice that the identity-based encryption algorithms from bilinear pairing in the literature must perform both elliptic curve group operations and multiplicative group operations, which consume a lot of hardware resources. We manage to eliminate the need of multiplicative group operations for encryption. This is a significant discovery since the hardware structure can be simplified for implementing pairing-based cryptography. Our experimental results show that our encryption algorithm saves up to 47 percent memory (27,239 RAM bits) in FPGA implementation.

ETPL DSC - 013	Privacy-Preserving Multi-keyword Top-k Similarity Search Over Encrypted Data

Cloud computing provides individuals and enterprises massive computing power and scalable storage capacities to support a variety of big data applications in domains like health care and scientific research, therefore more and more data owners are involved to outsource their data on cloud servers for great convenience in data management and mining. However, data sets like health records in electronic documents usually contain sensitive information, which brings about privacy concerns if the documents are released or shared to partially untrusted third-parties in cloud. A practical and widely used technique for data privacy preservation is to encrypt data before outsourcing to the cloud servers, which however reduces the data utility and makes many traditional data analytic operators like keyword-based top-k document retrieval obsolete. In this paper, we investigate the multi-keyword top-k search problem for big data encryption against privacy breaches, and attempt to identify an efficient and secure solution to this problem. Specifically, for the privacy concern of query data, we construct a special tree-based index structure and design a random traversal algorithm, which makes even the same query to produce different visiting paths on the index, and can also maintain the accuracy of queries unchanged under stronger privacy. For improving the query efficiency, we propose a group multi-keyword top-k search scheme based on the idea of partition, where a group of tree-based indexes are constructed for all documents. Finally, we combine these methods together into an efficient and secure approach to address our proposed top-k similarity search. Extensive experimental results on real-life data sets demonstrate that our proposed approach can significantly improve the capability of defending the privacy breaches, the scalability and the time efficiency of query processing over the state-of-the-art methods.

ETPL DSC - 014	Secure and Private Data Aggregation for Energy Consumption Scheduling in Smart Grids

The recent proposed solutions for demand side energy management leverage the two-way communication infrastructure provided by modern smart-meters and sharing the usage information with the other users. In this paper, we first highlight the privacy and security issues involved in the distributed demand management protocols. We propose a novel protocol to share required information among users providing privacy, confidentiality, and integrity. We also propose a new clustering-based, distributed multi-party computation (MPC) protocol. Through simulation experiments we demonstrate the efficiency of our proposed solution. The existing solutions typically usually thwart selfish and malicious behavior of consumers by deploying billing mechanisms based on total consumption during a few time slots. However, the billing is typically based on the total usage in each time slot in smart grids. In the second part of this paper, we formally prove that under the per-slot based charging policy, users have incentive to deviate from the proposed protocols. We also propose a protocol to identify untruthful users in these networks. Finally, considering a repeated interaction among honest and dishonest users, we derive the conditions under which the smart grid can enforce cooperation among users and prevents dishonest declaration of consumption.

ETPL DSC - 015	Using Virtual Machine Allocation Policies to Defend against Co-Resident Attacks in Cloud Computing

Cloud computing enables users to consume various IT resources in an on-demand manner, and with low management overhead. However, customers can face new security risks when they use cloud computing platforms. In this paper, we focus on one such threat-the co-resident attack, where malicious users build side channels and extract private information from virtual machines co-located on the same server. Previous works mainly attempt to address the problem by eliminating side channels. However, most of these methods are not suitable for immediate deployment due to the required modifications to current cloud platforms. We choose to solve the problem from a different perspective, by studying how to improve the virtual machine allocation policy, so that it is difficult for attackers to co-locate with their targets. Specifically, we (1) define security metrics for assessing the attack; (2) model these metrics, and compare the difficulty of achieving co-residence under three commonly used policies; (3) design a new policy that not only mitigates the threat of attack, but also satisfies the requirements for workload balance and low power consumption; and (4) implement, test, and prove the effectiveness of the policy on the popular open-source platform Open Stack.

ETPL DSC - 016	Efficient Delegated Private Set Intersection on Outsourced Private Datasets

Private set intersection (PSI) is an essential cryptographic protocol that has many real world applications. As cloud computing power and popularity have been swiftly growing, it is now desirable to leverage the cloud to store private datasets and delegate PSI computation to it. Although a set of efficient PSI protocols have been designed, none support outsourcing of the datasets and the computation. In this paper, we propose two protocols for delegated PSI computation on outsourced private datasets. Our protocols have a unique combination of properties that make them particularly appealing for a cloud computing setting. Our first protocol, O-PSI, satisfies these properties by using additive homomorphic encryption and point-value polynomial representation of a set. Our second protocol, EO-PSI, is mainly based on a hash table and point-value polynomial representation and it does not require public key encryption; meanwhile, it retains all the desirable properties and is much more efficient than the first one. We also provide a formal security analysis of the two protocols in the semi-honest model and we analyse their performance utilizing prototype implementations we have developed. Our performance analysis shows that EO-PSI scales well and is also more efficient than similar state-of-the-art protocols for large set sizes.

ETPL DSC - 017	Wormhole: The Hidden Virus Propagation Power of a Search Engine in Social Networks

Today search engines are tightly coupled with social networks, and present users with a double-edged sword: they are able to acquire information interesting to users but are also capable of spreading viruses introduced by hackers. It is challenging to characterize how a search engine spreads viruses, since the search engine serves as a virtual virus pool and creates propagation paths over the underlying network structure. In this paper, we quantitatively analyse virus propagation effects and the stability of the virus propagation process in the presence of a search engine in social networks. First, although social networks have a community structure that impedes virus propagation, we find that a search engine generates a propagation wormhole. Second, we propose an epidemic feedback model and quantitatively analyse propagation effects employing four metrics: infection density, the propagation wormhole effect, the epidemic threshold, and the basic reproduction number. Third, we verify our analyses on four real-world data sets and two simulated data sets. Moreover, we prove that the proposed model has the property of partial stability. Evaluation results show that, compablack with to a case without a search engine present, virus propagation with the search engine has a higher infection density, shorter network diameter, greater propagation velocity, lower epidemic threshold, and larger basic reproduction number.

ETPL DSC - 018	A Scalable Approach to Joint Cyber Insurance and Security-as-a-Service Provisioning in Cloud Computing

As computing services are increasingly cloud-based, corporations are investing in cloud-based security measures. The Security-asa- Service (SECaaS) paradigm allows customers to outsource security to the cloud, through the payment of a subscription fee. However, no security system is bulletproof, and even one successful attack can result in the loss of data and revenue worth millions of dollars. To guard against this eventuality, customers may also purchase cyber insurance to receive recompense in the case of loss. To achieve cost effectiveness, it is necessary to balance provisioning of security and insurance, even when future costs and risks are uncertain. To this end, we introduce a stochastic optimization model to optimally provision security and insurance services in the cloud. Since the model we design is a mixed integer problem, we also introduce a partial Lagrange multiplier algorithm that takes advantage of the total unimodularity property to find the solution in polynomial time. We also apply sensitivity analysis to find the exact tolerance of decision variables to parameter changes. We show the effectiveness of these techniques using numerical results based on real attack data to demonstrate a realistic testing environment, and find that security and insurance are interdependent.

ETPL DSC - 019	Privacy-Preserving Aggregate Queries for Optimal Location Selection

Private set intersection (PSI) is an essential cryptographic protocol that has many real world applications. As cloud computing power and popularity have been swiftly growing, it is now desirable to leverage the cloud to store private datasets and delegate PSI computation to it. Although a set of efficient PSI protocols have been designed, none support outsourcing of the datasets and the computation. In this paper, we propose two protocols for delegated PSI computation on outsourced private datasets. Our protocols have a unique combination of properties that make them particularly appealing for a cloud computing setting. Our first protocol, O-PSI, satisfies these properties by using additive homomorphic encryption and point-value polynomial representation of a set. Our second protocol, EO-PSI, is mainly based on a hash table and point-value polynomial representation and it does not require public key encryption; meanwhile, it retains all the desirable properties and is much more efficient than the first one. We also provide a formal security analysis of the two protocols in the semi-honest model and we analyse their performance utilizing prototype implementations we have developed. Our performance analysis shows that EO-PSI scales well and is also more efficient than similar state-of-the-art protocols for large set sizes.

ETPL DSC - 020	Privacy-Preserving Multi-keyword Top-k Similarity Search Over Encrypted Data

Cloud computing provides individuals and enterprises massive computing power and scalable storage capacities to support a variety of big data applications in domains like health care and scientific research, therefore more and more data owners are involved to outsource their data on cloud servers for great convenience in data management and mining. However, data sets like health records in electronic documents usually contain sensitive information, which brings about privacy concerns if the documents are released or shared to partially untrusted third-parties in cloud. A practical and widely used technique for data privacy preservation is to encrypt data before outsourcing to the cloud servers, which however reduces the data utility and makes many traditional data analytic operators like keyword-based top-k document retrieval obsolete. In this paper, we investigate the multi-keyword top-k search problem for big data encryption against privacy breaches, and attempt to identify an efficient and secure solution to this problem. Specifically, for the privacy concern of query data, we construct a special tree-based index structure and design a random traversal algorithm, which makes even the same query to produce different visiting paths on the index, and can also maintain the accuracy of queries unchanged under stronger privacy. For improving the query efficiency, we propose a group multi-keyword top-k search scheme based on the idea of partition, where a group of tree-based indexes are constructed for all documents. Finally, we combine these methods together into an efficient and secure approach to address our proposed top-k similarity search. Extensive experimental results on real-life data sets demonstrate that our proposed approach can significantly improve the capability of defending the privacy breaches, the scalability and the time efficiency of query processing over the state-of-the-art methods.

ETPL DSC - 021	Fault-Tolerant Adaptive Routing in Dragonfly Networks

Dragonfly networks have been widely used in the current high-performance computers or high-end servers. Fault-tolerant routing in dragonfly networks is essential. The rich interconnects provide good fault-tolerance ability for the network. A new deadlock free adaptive fault-tolerant routing algorithm based on a new two-layer safety information model, is proposed by mapping routers in a group, and groups of the dragonfly network into two separate hyper cubes. The new fault-tolerant routing algorithm tolerates static and dynamic faults. Our method can determine whether a packet can reach the destination at the source by using the new safety information model, which avoids dead-ends and aimless misrouting. Sufficient simulation results show that the proposed fault-tolerant routing algorithm even outperforms the previous minimal routing algorithm in fault-free networks in many cases.

ETPL DSC - 022	On the Security of a Variant of ElGamal Encryption Scheme

Recently, based on the Paillier cryptosystem [1], Yi et al. [2] proposed a distributed ElGamal cryptosystem which allows for both a much simpler distributed key generation procedure and distributed decryption of messages from a large plaintext domain. In this paper, we analyze the security of their proposed variant of ElGamal encryption scheme and demonstrate that their construction is not secure as claimed.

ETPL DSC - 023	A Scalable Approach to Joint Cyber Insurance and Security-as-a-Service Provisioning in Cloud Computing

As computing services are increasingly cloud-based, corporations are investing in cloud-based security measures. The Security-asa- Service (SECaaS) paradigm allows customers to outsource security to the cloud, through the payment of a subscription fee. However, no security system is bulletproof, and even one successful attack can result in the loss of data and revenue worth millions of dollars. To guard against this eventuality, customers may also purchase cyber insurance to receive recompense in the case of loss. To achieve cost effectiveness, it is necessary to balance provisioning of security and insurance, even when future costs and risks are uncertain. To this end, we introduce a stochastic optimization model to optimally provision security and insurance services in the cloud. Since the model we design is a mixed integer problem, we also introduce a partial Lagrange multiplier algorithm that takes advantage of the total unimodularity property to find the solution in polynomial time. We also apply sensitivity analysis to find the exact tolerance of decision variables to parameter changes. We show the effectiveness of these techniques using numerical results based on real attack data to demonstrate a realistic testing environment, and find that security and insurance are interdependent.

ETPL DSC - 024	Efficient Multi-Factor Authenticated Key Exchange Scheme for Mobile Communications

Authenticated key exchange (AKE) is one of the most important applications in applied cryptography, where a user interacts with a server to set up a session key where pre-registered information (aka. authentication factor), such as a password or biometrics, of the user is stored. While single-factor AKE is widely used in practice, higher security concerns call for multi-factor AKE (MFAKE) schemes, e.g. combining both passwords and biometrics simultaneously. However, in some casually designed schemes, security is even weakened in the sense that leakage of one authentication factor will defeat the whole MFAKE protocol. Furthermore, an inevitable by-product arise that the usability of the protocol often drop greatly. To summarize, the existing multi-factor protocols did not provide enough security and efficiency simultaneously. In this paper, we make one step ahead by proposing a very efficient MFAKE protocol. We define the security model and give the according security analysis. We also implement our protocol on a smartphone and a cloud server. The theoretic comparisons and the experimental results show that our scheme achieves both security and usability.

ETPL DSC - 025	Privacy-Preserving Aggregate Queries for Optimal Location Selection

A practical and widely used technique for data privacy preservation is to encrypt data before outsourcing to the cloud servers, which however reduces the data utility and makes many traditional data analytic operators like keyword-based top-k document retrieval obsolete. In this paper, we investigate the multi-keyword top-k search problem for big data encryption against privacy breaches, and attempt to identify an efficient and secure solution to this problem. Specifically, for the privacy concern of query data, we construct a special tree-based index structure and design a random traversal algorithm, which makes even the same query to produce different visiting paths on the index, and can also maintain the accuracy of queries unchanged under stronger privacy. For improving the query efficiency, we propose a group multi-keyword top-k search scheme based on the idea of partition, where a group of tree-based indexes are constructed for all documents.

ETPL DSC - 026	Secure and Efficient Initialization and Authentication Protocols for SHIELD

With the globalization of semiconductor production, out-sourcing IC fabrication has become a trend in various aspects. This, however, introduces serious threats from the entire untrusted supply chain. To combat these threats, DARPA (Defense Advanced Research Projects Agency) proposed in 2014 the SHIELD (Supply Chain Hardware Integrity for Electronics Defense) program to design a secure hardware root-of-trust, called dielet, to be inserted into the host package of legitimately produced ICs. Dielets are RF powered and communicate with the outside world through their RF antennas. They have sensors which allow them to passively (without the need for power) record malicious events which can later be read out during an authentication protocol between the dielet and server with a smartphone as intermediary. This paper introduces a general framework for the initialization and authentication protocols in SHIELD with different adversarial models based on formally-defined security games. We introduce a “try-and-check” attack against DARPA’s example authentication protocol in their call for SHIELD proposals which nullifies the effectiveness of SHIELD’s main goal of being able to detect and trace adversarial activities with significant probability. We introduce the first concrete initialization protocol and, compared to DARPA’s example authentication protocol, introduce an improved authentication protocol which resists the try-and-check attack. The area overhead of our authentication and initialization protocols together is only 64-bit NVM, one 8-bit counter and a TRNG based on a single SRAM-cell together with corresponding control logic. Our findings and rigorous analysis are of utmost importance for the teams which received DARPA’s funding for implementing SHIELD.

ETPL DSC - 027	Cryptographic Solutions for Credibility and Liability Issues of Genomic Data

In this work, we consider a scenario that includes an individual sharing his genomic data (or results obtained from his genomic data) with a service provider. In this scenario, (i) the service provider wants to make sure that received genomic data (or results) in fact belongs to the corresponding individual (and computed correctly), (ii) the individual wants to provide a digital consent along with his data specifying whether the service provider is allowed to further share his data, and (iii) if his data is shared without his consent, the individual wants to determine the service provider that is responsible for this leakage. We propose two schemes based on homomorphic signature and aggregate signature that links the information about the legitimacy of the data to the consent and the phenotype of the individual. Thus, to verify the data, each party also needs to use the correct consent and phenotype of the individual who owns the data.

ETPL DSC - 028	Design and Implementation of the Ascend Secure Processor

This paper presents hardware implementations of the Ascend secure processor, prototyped on an FPGA and taped out in a 32 nm SOI process. Ascend prevents information leakage over a processor's digital I/O pins — in particular, the processor's requests to external memory — and certifies the program's execution by integrity-verifying the external memory. In secure processor design, encrypting main memory is not sufficient for security because where and when memory is accessed reveals secret information. To this end, Ascend is equipped with a hardware Oblivious RAM (ORAM) controller, which obfuscates the address bus by reshuffling memory as it is accessed. To our knowledge, Ascend is the first prototyping of ORAM in custom silicon. Ascend has also been carefully engineered to ensure its timing behaviors are independent of user private data. We describe our open-source FPGA prototype and the different design considerations that were made when optimizing for an FPGA vs. the ASIC. In 32 nm silicon, all security components combined (the ORAM controller, which includes 12 AES rounds and one SHA-3 hash unit) impose a moderate area overhead of 1 mm2. Post tape-out, the Ascend chip has been successfully tested at 500 MHz.

ETPL DSC - 029	Privacy-Aware Caching in Information-Centric Networking

Information-Centric Networking (ICN) is an emerging networking paradigm where named and routable data (content) is the focal point. Users send explicit requests (interests) which specify content by name, and the network handles routing these interests to some entity capable of satisfying them with the appropriate data response (producer). One key feature of ICN is opportunistic in-network content caching. This property facilitates efficient content distribution by reducing bandwidth consumption, lessening network congestion, and improving the content retrieval latency by users (consumers). Unfortunately, the same feature is also detrimental to privacy of content consumers and producers. Simple to implement, and difficult to detect, timing attacks can exploit ICN routers as "oracles" and allow an adversary to learn whether a nearby consumer recently requested certain content. The attack leverages a timing side channel that relies on router caches and is implemented by requesting a few packets from each piece of content being probed. Similarly, probing attacks that target content producers can be used to discover whether certain content has been recently distributed. After analyzing the scope and feasibility of such attacks, we propose and evaluate some efficient countermeasures that offer quantifiable privacy guarantees while retaining the benefits of ICN.

ETPL DSC - 030	Efficient and Private Scoring of Decision Trees, Support Vector Machines and Logistic Regression Models based on Pre-Computation

Many data-driven personalized services require that private data of users is scored against a trained machine learning model. In this paper we propose a novel protocol for privacypreserving classification of decision trees, a popular machine learning model in these scenarios. Our solutions is composed out of building blocks, namely a secure comparison protocol, a protocol for obliviously selecting inputs, and a protocol for multiplication. By combining some of the building blocks for our decision tree classification protocol, we also improve previously proposed solutions for classification of support vector machines and logistic regression models. Our protocols are information theoretically secure and, unlike previously proposed solutions, do not require modular exponentiations. We show that our protocols for privacy-preserving classification lead to more efficient results from the point of view of computational and communication complexities. We present accuracy and runtime results for 7 classification benchmark datasets from the UCI repository.

ETPL DSC - 031	Design and Implementation of SecPod, A Framework for Virtualization-based Security Systems

The OS kernel is critical to the security of a computer system. Many systems have been proposed to improve its security. A fundamental weakness of those systems is that page tables, the data structures that control the memory protection, are not isolated from the vulnerable kernel, and thus subject to tampering. To address that, researchers have relied on virtualization for reliable kernel memory protection. Unfortunately, such memory protection requires to monitor every update to the guest's page tables. This fundamentally conflicts with the recent advances in the hardware virtualization support. In this paper, we present the design and implementation of SecPod, a practical and extensible framework for virtualization-based security systems that can provide both strong isolation and the compatibility with modern hardware. SecPod has two key techniques: paging delegation delegates and audits the kernel's paging operations to a secure space; execution trapping intercepts the (compromised) kernel's attempts to subvert SecPod by misusing privileged instructions. We have implemented a prototype of SecPod based on KVM. Our experiments show that SecPod is both effective and efficient.

ETPL DSC - 032	Fuzzy Identity-Based Data Integrity Auditing for Reliable Cloud Storage Systems

Data integrity, a core security issue in reliable cloud storage, has received much attention. Data auditing protocols enable a verifier to efficiently check the integrity of the outsourced data without downloading the data. A key research challenge associated with existing designs of data auditing protocols is the complexity in key management. In this paper, we seek to address the complex key management challenge in cloud data integrity checking by introducing fuzzy identity-based auditing, the first in such an approach, to the best of our knowledge. More specifically, we present the primitive of fuzzy identity-based data auditing, where a user's identity can be viewed as a set of descriptive attributes. We formalize the system model and the security model for this new primitive. We then present a concrete construction of fuzzy identity-based auditing protocol by utilizing biometrics as the fuzzy identity. The new protocol offers the property of error-tolerance, namely, it binds with private key to one identity which can be used to verify the correctness of a response generated with another identity, if and only if both identities are sufficiently close. We prove the security of our protocol based on the computational Diffie-Hellman assumption and the discrete logarithm assumption in the selective-ID security model. Finally, we develop a prototype implementation of the protocol which demonstrates the practicality of the proposal.

ELYSIUMPRO

A UNIT OF ELYSIUM GROUPS

Thank you!

ELYSIUM GROUPS