

Research Statement

Wenchao Zhou

<http://www.cis.upenn.edu/~wenchaoz/>

1 Introduction

My research aims to provide *formally provable* safety and security guarantees in developing distributed systems. My work is motivated by the popularity of distributed systems deployed for a variety of applications, ranging from cloud computing platforms located at data centers, to global-scale interdomain routing deployed across multiple administrative and geographical domains. Given their pervasive usage that is closely coupled with daily lives, consequences of faults in distributed systems can be catastrophic. For instance, faults in distributed systems can expose sensitive information or disrupt of mission-critical transactions. However, designing and implementing robust distributed systems remains a daunting task, in part because of the massive scale of typical deployments, the complexity and unpredictability of system executions, but also due to emerging security threats.

To address this challenge, my research adopts a multi-disciplinary approach, by combining techniques from databases, security, networking and formal methods. I consider enhancement of system reliability and security in different phases of the system development cycle. First, in the *design and development* phase, I use programming language and formal method techniques to enable automated code synthesis for provably safe and secure implementation that guards against foreseeable faults or vulnerabilities. Second, as a treatment for unexpected faults occurred in the *deployment* phase, I adopt and extend the database notion of provenance to enable efficient maintenance and querying of system dependencies for fault diagnosis and debugging. In addition, I work with researchers in applied cryptography and computer systems to develop security techniques for providing guarantees even with the existence of Byzantine faults.

I actively involve researchers with diverse background in my research – I have co-authored publications with eight professors from four universities, and nine researchers from five industry labs. My work has been published at premier conferences in a variety of research areas, such as SOSP [12], SIGMOD [7, 17], ICDE [3, 14], NDSS [6], and CIDR [4]. In all my research projects, I work towards full-fledged implementations that are released as open-source [5], as well as live-demonstrations at SIGMOD [13] and SIGCOMM [16]. I led a demonstration at SIGCOMM, which received the 2nd prize for the ACM Student Research Competition.

2 Dissertation Work: Secure Network Provenance

My dissertation research on *Secure Network Provenance (SNP)* [12] presents an approach that provides the fundamental functionality required for performing fault analysis and debugging – the capability to “explain” the existence (or change) of system state in a potentially *adversarial* environment. SNP reveals the dependencies between system states, and permits system operators to transitively tie observed faults to their potential causes, and to assess the damage that these faults may have caused to the rest of the system. I have identified several practical challenges in deploying SNP, and have presented the solutions that addressed each of the main challenges:

Distribution. A key challenge of supporting provenance in distributed system is to develop an abstract system model in which provenance data can be maintained efficiently. I demonstrated that it is achievable by modeling the system state as a set of distributed databases, and by extracting logical dependencies from system specifications and runtime [10, 13, 17]. Enabled by the distributed query processing capabilities, provenance information is then *incrementally* maintained as views of system state during the execution. I analytically and empirically showed that the overhead incurred by provenance maintenance is linear in the cost of the base system, and, therefore, does not affect its scalability.

Time-awareness. Another challenge is to track state changes over time in a relaxed system model, in which clocks are not synchronized and messages can be delayed, reordered or lost. To address this challenge, I examined the fundamental correlation between provenance and (observable) event ordering in distributed systems. I then presented an enhanced provenance model that provides a sound and complete representation that correctly captures the system dependencies [11, 15].

Security. A final challenge is to provide security guarantees in *completely untrusted* environments, in which the adversary may have compromised an arbitrary subset of the nodes, and that he may have complete control over these nodes. I showed that, despite the conservative threat model, our security enhancement

in SNP still provides strong, provable guarantees [12]: it ensures that an observable symptom of a fault or an attack can always be traced to a specific event—passive evasion or active misbehavior—on at least one faulty node, even when the adversary attempts to prevent this.

To demonstrate SNP’s practicality and generality, I have applied it to a variety of different systems, including the Internet’s interdomain routing system, the Chord distributed hash table, and the Hadoop MapReduce system. My evaluation have demonstrated that SNP is able to detect a number of different problems that had been previously described in the literature, and that SNP is practical, both in terms of its run-time overhead and in terms of the effort required to deploy it.

3 Generating Formally Safe and Secure Implementations

An approach that complements the reactive fault analysis is the development of formal analysis frameworks and domain specific languages that *automates* the process of generating formally safe and secure implementations. Here, the goal is to use formal tools and programming language techniques to generate implementations that maintain specified invariants throughout system executions. I have explored this general idea in the following two settings:

Formally Safe Routing (FSR). The FSR [8] system automates the process of safety (i.e., convergence) analysis for routing configurations using a constraint solver, and compiles the verified configuration into distributed implementations. In a recent prototype demonstration [16], FSR was used to detect problems in an autonomous system’s iBGP configuration (using realistic topologies), and to prove sufficient conditions for BGP safety.

Secure Network Datalog (SeNDlog). SeNDlog [14] unifies logic-based access languages and distributed query languages for declarative networks. Given its root in logic-based languages, SeNDlog allows rigorous reasoning and static verification for safety and secure properties (e.g., authenticity of communications). It bridges the gap between formally verified specifications and implementation in practice using a provably correct compilation process.

In collaboration with LogicBlox Inc., I further extended SeNDlog to support *extensible trust management* [4], where various security constructs can be customized and composed in a declarative fashion. It supports a variety of security primitives (e.g., speaks-for and restricted delegation) for authentication, integrity and confidentiality, which enables a wide range of distributed applications with customizable security policies, while still benefiting from its roots in logic-based languages.

As a feasibility demonstration of my methods for developing secure distributed systems in practice, the *Application-Aware Anonymity (A³)* [6] system provides an extensible platform for applications to deploy anonymity-based services on the Internet. A³ uses SeNDlog for customizing how relays in a routing path should be selected, and how the routing path should be instantiated. It allows applications to tailor the anonymity properties and performance characteristics according to their communication requirements.

4 Future Work

Below, I briefly outline my future research agenda, describing short-term extensions to my dissertation work as well as my longer-term research plan.

4.1 Extensions to Dissertation

SNP automates fault diagnosis and debugging, by systematically maintaining and querying state dependencies as the system execution progresses. I intend to further improve its usability to encourage its adoption in academia and industrial development settings.

One important aspect for future research is to enhance the “readability” of the returned provenance results. Provenance information could be overwhelmingly large in systems with complex dependency logic. To address this challenge, I intend to explore the following two complementary approaches: the first approach focuses on developing an expressive yet easy-to-use interface (e.g., a SQL-like declarative query language tailored for the SNP model), for users to annotate and prune provenance data based on a customizable pattern; alternatively, the size and complexity of provenance information can be controlled by introducing *layering* into the provenance system, in which case provenance data can be captured at a variety of granularities, and be interactively expanded.

SNP mainly explores the authenticity and integrity aspects of security in provenance systems. I plan to extend the exploration to their counterparts, privacy and confidentiality. It is intriguing to study the

tension between privacy and verifiability, two seemingly contradictory properties. As a first step, the private and verifiable routing (PVR) [1] provides initial evidence that strong privacy guarantees can be achieved in the interdomain routing, where the functionality of each node is well-restricted to route selection and advertisement based on a customized ranking function. I intend to further understand the performance implications or limits when extending the guarantees to more general systems.

4.2 Looking Beyond

Looking beyond my dissertation work, my long-term research goal is to facilitate the development of provably secure and reliable distributed systems. Specifically, I intend to accomplish this by enhancing the iterative development cycle of distributed systems, with research advances in the following aspects:

Rigorous verification on design and vulnerabilities. Formal verification provides sound and complete guarantees by checking that a certain set of properties holds in *all* possible execution traces. My current research focuses on static analysis, by leveraging prove-by-construction capability enabled by logic-based programming languages. I intend to extend the exploration to dynamic verification techniques [2], such as model checking, for verifying more complex properties and performing “what-if” analysis to discover potential vulnerabilities given certain assumptions on the attack model.

Systematic fault diagnosis and recovery. SNP can be used to systematically diagnose faults, where “explanations” of a suspicious symptom are compiled as a set of state dependencies that recursively trace back to the root causes. As the basis for inferring state dependencies, the high-level dependency logic (captured as derivation rules in SNP) is of critical importance. To generalize and further automate the extraction of such dependency logic from a target application, one potential avenue that I intend to explore is to employ programming language techniques that perform static (or dynamic) analysis on the information flow of the target system.

In addition to debugging, one intriguing direction is the use of SNP for *provenance-based* recovery. SNP maintains sufficient information to reproduce the system execution trace individually for each node. This brings opportunities to *undo* the damages caused by an exposed system fault, by applying the inverse operations in the reverse order. For example, a mistakenly deleted system state can be restored by the corresponding insertion. In addition, provenance keeps the dependency information and thus allows *minimal recovery*, i.e., the recovery only impacts the nodes that are *actually* affected by the fault.

Provenance-driven invariant generation. One important challenge in formal verification is for system designers to discover the safety properties (or invariants). The quality of these safety properties directly affects the quality of the verification results, however, there lacks a systematic approach to extract the safety properties, and the process largely relies on manual efforts today. To address this problem, I am interested in feeding the design bugs or security vulnerabilities exposed in fault diagnosis (as hints for the safety properties) to refine the invariants in the design and development phase.

5 Other Work

In addition to the research work around the general theme of developing provable secure and reliable distributed systems, I have worked on a variety of other projects in both academic and industrial settings. I provide a brief summary of each project:

Query optimization in streaming systems. With a proliferation of edge devices such as smartphones, and the global availability of Cloud computing resources, the Cloud-Edge topology – where multiple edge devices are inter-connected through the Cloud – is becoming commonplace. Edge devices generate real-time data such as GPS location, battery consumption, etc. There exists a broad class of distributed applications that involve data correlation across multiple edge devices and the Cloud.

With collaborators from Microsoft Research, I proposed the RACE (Real-time Applications over Cloud-Edge topologies) system to support such applications [9]. The key challenge in RACE lies in the generation of a distributed query plan that optimizes global overhead. To address this challenge, I developed novel operator placement algorithms that are provably optimal. An evaluation over real datasets indicated that RACE is highly scalable, and is orders of magnitude more efficient than current placement strategies.

Cloud serving and batch systems. Cloud data management systems can largely be divided into batch-oriented and serving-oriented systems. However, a number of applications (e.g., targeted advertising) cannot be cleanly mapped to one of these types of systems, but instead require hybrid functionality. In combining batch and serving systems [7], my collaborators at Yahoo! Research and I identified several key areas where

the fundamental properties, such as failure and recovery granularity, performance optimization metrics, and consistency model, of each system are mismatched. Techniques are then introduced to mitigate these conflicts, and to provide hybrid functionality in an efficient manner.

Distributed monitoring and checking. In addition to the research on the static verification enabled by the logic-based languages (Section 3), I demonstrated that runtime monitoring and checking are also achievable within the same framework. I have conceptualized and developed DMaC [18], a system that monitors and checks running systems against formally specified properties in a declarative fashion, where high-level safety properties are automatically compiled from platform-independent formal specifications into distributed monitoring queries for execution.

References

- [1] Alexander J. T. Gurney, Andreas Haeberlen, **Wenchao Zhou**, Micah Sherr, and Boon Thau Loo. Having your Cake and Eating it too: Routing Security with Privacy Protections. In *Proc. ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets)*, 2011.
- [2] Limin Jia, Chen Chen, Sangeetha A. Jyothi, **Wenchao Zhou**, Suyog Mapara, and Boon Thau Loo. Towards a Secure and Verifiable Future Internet. In *Proc. Workshop on Off the Beaten Track (OBT), co-located with POPL*, 2012.
- [3] Mengmeng Liu, Nicholas Taylor, **Wenchao Zhou**, Zachary Ives, and Boon Thau Loo. Recursive Computation of Regions and Connectivity in Networks. In *Proc. IEEE Int. Conf. on Data Engineering (ICDE)*, 2009. **Best Papers of ICDE 2009.**
- [4] William R. Marczak, David Zook, **Wenchao Zhou**, Molham Aref, and Boon Thau Loo. Declarative Reconfigurable Trust Management. In *Proc. Biennial Conf. on Innovative Data Systems Research (CIDR)*, 2009.
- [5] RapidNet. RapidNet Declarative Networking Engine. <http://netdb.cis.upenn.edu/rapidnet/>.
- [6] Micah Sherr, Andrew Mao, William R. Marczak, **Wenchao Zhou**, Boon Thau Loo, and Matt Blaze. A3: An Extensible Platform for Application-Aware Anonymity. In *Proc. Network and Distributed Systems Security Symp. (NDSS)*, 2010.
- [7] Adam Silberstein, Russell Sears, **Wenchao Zhou**, and Brian F. Cooper. A Batch of Pnuts: Experiences Connecting Cloud Batch and Serving Systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data (SIGMOD)*, 2011.
- [8] Anduo Wang, Limin Jia, **Wenchao Zhou**, Yiqing Ren, Boon Thau Loo, Jennifer Rexford, Vivek Nigam, Andre Scedrov, and Carolyn Talcott. FSR: Formal Analysis and Implementation Toolkit for Safe Inter-domain Routing. In *submission to IEEE/ACM Trans. Networking*, 2011.
- [9] **Wenchao Zhou**, Badrish Chandramouli, and Suman Nath. Specification and Optimization of Real-Time Cloud-Edge Applications. In *submission*, 2011.
- [10] **Wenchao Zhou**, Eric Cronin, and Boon Thau Loo. Provenance-aware Secure Networks. In *Proc. Int. Conf. on Data Engineering Workshops (NetDB)*, 2008.
- [11] **Wenchao Zhou**, Ling Ding, Andreas Haeberlen, Zachary Ives, and Boon Thau Loo. Time-aware Provenance for Distributed Systems. In *Proc. USENIX Workshop on Theory & Practice of Provenance (TaPP)*, 2011.
- [12] **Wenchao Zhou**, Qiong Fei, Arjun Narayan, Andreas Haeberlen, Boon Thau Loo, and Micah Sherr. Secure Network Provenance. In *Proc. ACM Symp. on Operating System Principles (SOSP)*, 2011.
- [13] **Wenchao Zhou**, Qiong Fei, Shengzhi Sun, Tao Tao, Andreas Haeberlen, Zachary Ives, Boon Thau Loo, and Micah Sherr. NetTrails: A Declarative Platform for Provenance Maintenance and Querying in Distributed Systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data (SIGMOD) - Demonstration*, 2011.
- [14] **Wenchao Zhou**, Yun Mao, Boon Thau Loo, and Martín Abadi. Unified Declarative Platform for Secure Networked Information Systems. In *Proc. IEEE Int. Conf. on Data Engineering (ICDE)*, 2009.
- [15] **Wenchao Zhou**, Suyog Mapara, Yang Li, Andreas Haeberlen, Zachary Ives, Boon Thau Loo, and Micah Sherr. TapTrails: Provenance of Changes for Distributed Systems. In *submission*, 2011.
- [16] **Wenchao Zhou**, Yiqing Ren, Anduo Wang, Limin Jia, Alexander J.T. Gurney, Boon Thau Loo, and Jennifer Rexford. FSR: Formal Analysis and Implementation Toolkit for Safe Inter-domain Routing. In *Proc. ACM SIGCOMM Int. Conf. on Data Communication (SIGCOMM) - Demonstration*, 2011. **Runner-up for the ACM Student Research Competition.**
- [17] **Wenchao Zhou**, Micah Sherr, Tao Tao, Xiaozhou Li, Boon Thau Loo, and Yun Mao. Efficient Querying and Maintenance of Network Provenance at Internet-Scale. In *Proc. ACM SIGMOD Int. Conf. on Management of Data (SIGMOD)*, 2010.
- [18] **Wenchao Zhou**, Oleg Sokolsky, Boon Thau Loo, and Insup Lee. DMaC: Distributed Monitoring and Checking. In *Proc. Int. Conf. on Runtime Verification (RV)*, 2009.