

Personal Statement of Interest

My research work aims to ease the implementation, management, and analysis of secure distributed systems. I adopt a multi-disciplinary approach towards address this problem by unifying three bodies of work: (a) logic-based trust management systems, (b) declarative networking [1][2] that provides a programming framework that enable compact declarative specifications of network protocols, and (c) database techniques for analyzing data computations via the concept of provenance. I seek to ground my research work in practical problems by collaborating closely with industry (e.g. Microsoft Research Asia and LogicBlox) with an eye towards solving realistic problems.

In recent years, we have witnessed a proliferation of networked information systems deployed at Internet-scale for a variety of application domains. Despite the widespread usage, designing and implementing these large-scale systems remains a challenge, in part because of the emerging security threats. Though there have been several proposals attacking this problem, they are typically designed to tackle specific security threats at the underlying network, without taking into account information processing at higher layers. Moreover, they are often implemented and enforced in different languages or environments, hence having different tradeoffs in expressiveness, complexity and performance.

Given these challenges, we have been striving to develop a unified and re-configurable platform on which user can specify a variety of network infrastructures together with security specifications. From a practical standpoint, this integration has several benefits, ranging from fewer languages to learn, fewer sets of optimizations, finer-grain control over the interaction between security and network protocols, and the potential of cross-layer analysis and optimizations.

First, we present the *Secure Network Datalog (SeNDlog)* [3] language that unifies Binder, a logic-based access control language, and *Network Datalog*, a distributed recursive query language originally proposed for declarative networks. *SeNDlog* enables network routing, information systems, and their security policies to be specified and implemented within a common declarative framework. In addition, we extend existing distributed recursive query processing techniques to execute *SeNDlog* programs that incorporate authenticated communications. Besides authentication, we are also exploring in extending the declarative language to incorporate security constructs for secrecy and encrypted facts. Other language features in our research radar involves vote-based security protocols and privilege revocation.

Second, in collaboration with Bill Marczak (Penn undergrad) and LogicBlox Inc., we extend our exploration to a declarative system for reconfigurable trust management, which enables customizable cryptographic, partitioning and distribution strategies based on the execution environment. We propose the *LBTrust* [4] based on the *LogicBlox* [5] language which is a variant of Datalog with the support of constraints, meta-programming, and the programmer-defined constraints which read and modify the meta-model itself.

Third, we propose the use of network provenance for explaining the derivation of network states, and explore its applicability to real network security applications. To ensure scalability,

we explore the use of traditional distributed database optimization techniques [6] to significantly reduce communication overhead of such provenance computations.

References

- [1] B. T. Loo, T. Condie, J. M. Hellerstein, P. Maniatis, T. Roscoe, I. Stoica, Implementing Declarative Overlays, In *20th ACM Symposium on Operating Systems Principles (SOSP)*, Brighton, UK, October 2005.
- [2] B. T. Loo, J. M. Hellerstein, I. Stoica, R. Ramakrishnan, Declarative Routing: Extensible Routing with Declarative Queries, In *ACM SIGCOMM Conference on Data Communication*, Philadelphia, PA, Aug 2005.
- [3] W. Zhou, Y. Mao, B. T. Loo, M. Abadi, Unified Declarative Platform for Secure Networked Information Systems. In *25th International Conference on Data Engineering (ICDE)*, Shanghai, China, Apr 2009.
- [4] W. R. Marczak, D. Zook, W. Zhou, M. Aref, B. T. Loo, Declarative Reconfigurable Trust Management. In *4th Biennial Conference on Innovative Data Systems Research (CIDR) Proceedings Track*, Pacific Grove, CA, Jan 2009.
- [5] LogicBlox. <http://www.logicblox.com/>
- [6] M. Liu, N. E. Taylor, W. Zhou, Z. Ives, B. T. Loo. Recursive Computation of Regions and Connectivity in Networks. In *25th International Conference on Data Engineering (ICDE)*, Shanghai, China, Apr 2009.