Research Statement

Boon Thau Loo
http://netdb.cis.upenn.edu

1 Summary

My research focuses on the development of new programming tools and analysis techniques that improve the process of designing, implementing, and securing large-scale distributed systems. At the University of Pennsylvania, I lead the NetDB@Penn (`http://netdb.cis.upenn.edu`) research group, where we work on a wide range of interdisciplinary research projects, combining techniques across varying subfields in computer science, including databases, formal methods, programming languages, with applications to networking, security, real-time systems, and distributed systems in general.

A unifying theme of my research is the use of novel programming frameworks such as *declarative networking* [36–40], *constraint-based programming* [26,27,30], and *scenario-based programming* [66,67], in combination with formal verification and optimization tools for automatically synthesizing provably correct distributed systems with performance guarantees. In declarative networking, distributed systems are specified and implemented using a declarative recursive query language. Declarative networking fundamentally changes the way network protocols are designed and implemented. My early work in declarative networking demonstrated its use as a domain specific language for fast prototyping, by compiling protocols written in a declarative language into efficient distributed implementations.

My work in the past decade has evolved declarative networking's original roots as a rapid prototyping framework, towards one that serves as an important bridge connecting formal theories (for reasoning about protocol correctness) and actual implementations. The ability to bridge this gap is a major step forward compared to traditional approaches, where formal specifications, proof of protocol correctness and implementations are decoupled from one another. This decoupling leads to increased development time, error-prone implementations, and tedious debugging. I have also successfully connected declarative networking work with other domains, for instance, data provenance, anonymous communications, and network forensics, and unified its declarative framework with similar paradigms from the security, real-time systems, and optimizations communities.

Given the inter-disciplinary nature of our work, my students and I publish across a wide range of conferences, including **databases** (SIGMOD [21, 43, 91], VLDB [33, 90], CIDR [44], ICDE [34, 89]), **networking** (SIGCOMM [2, 6, 65, 83], PODC [60], IEEE/ACM Transactions on Networking [27, 29, 58, 84], HotNets [5, 19, 57, 63, 64, 66], CoNEXT [42, 67], INFOCOM [24, 55], ICNP [3, 28], COMSNETS [26]), **security** (NDSS [53], HotSec [52], PETS [25, 49], SOSR [10]), **systems** (SOSP [87], SOCC [30], USENIX ATC [31,50], HPDC [13]), and **formal methods and programming languages** (TACAS [61], FORTE [8], PPDP [11,46], FMOODS [62], TPHOLs [59], PADL [39,54]). I have also published two books on declarative networking [41] (Morgan and Claypool publishers) and recursive query processing [17] (now publishers). My students and I have received three best paper awards in diverse venues in databases (ICDE 2009 [34]), autonomic computing (ICAC 2012 [78]), and high-performance computing (HPDC [13]). A fourth paper [67] was nominated for best paper at a networking venue (CoNEXT 2015).

Our research prototypes have been demonstrated at major conferences such as SIGMOD [88], SIGCOMM [16, 18, 45, 48, 56], NSDI [32, 86], and COMSNETS [26], and experimentally deployed on testbeds such as PlanetLab [42,89], ORBIT wireless [26,28,29], Amazon EC2 [79,80,87], and AT&T's ShadowNet [31]. Open-source implementations have been released, and elements of my research has also resulted in technology transfer via commercial adoption and startup companies.

In the rest of this statement, I will summarize the research work in my group, broadly classified under different stages of the distributed systems development cycle: *synthesizing provably correct distributed systems* (Section 2), *debugging distributed systems* (Section 3), *securing distributed systems* (Section 4), and *optimizing distributed systems* (Section 5). Section 6 provides some highlights on student development through research mentoring in my group. I will conclude in Section 7 with a discussion of our overall research impact, technology transfer, and future directions.

2 Thrust 1: Synthesizing Provably Correct Distributed Systems

Our first research thrust addresses a long-standing challenge in networking research – bridging the gap between formal theories and actual implementations. One of our first contributions is the *Formally Safe Routing* (FSR) toolkit [48, 57, 58], that attempts to bridge this gap in the context of interdomain routing. FSR is motivated by the complexity of policy configuration over the Internet's global routing system, whose convergence depends on how individual networks configure their Border Gateway Protocol (BGP) policies. Since protocol oscillations cause serious performance disruptions and router overhead, researchers devote significant attention to BGP stability (or "safety").

FSR unifies research in routing algebras with recent advances in declarative networking to produce provably correct distributed implementations. Specifically, FSR automates the process of analyzing routing configurations expressed in algebra for safety using SMT solvers, and automatically compiles routing algebra into declarative routing implementations.

In addition to the FSR toolkit, we have explored the use of theorem proving [54] and rewriting logic techniques [62] for *verifying* routing protocols. We further explored the use of reduction-based techniques [55, 56, 60, 61] to make formal analysis feasible on large Internet graphs, and performing safety analysis and routing reconfigurations in the presence of incomplete routing information [18, 20]. Through our work on this thrust, we have also formalized the operational semantics of declarative networking programs [46], developed automated ways [11] of checking safety properties in declarative networking programs using SMT solvers, and proposed new distributed recursive view maintenance techniques [34, 35].

FSR uses declarative networking as an intermediary domain-specific language for generating provably correct distributed systems. Domain specific languages, while useful, are perhaps not the most intuitive approach for network operators, who often times are the consumers of router solutions. In a NSF Innovation Corps program, we interviewed more than 100 network operators and engineers. We observe that while operators are experts at configuring routers, many of them are not be trained in programming, let alone new domain specific languages that have significant learning curves. To address this need, we recently explore a novel alternative approach based on *scenario-based programming*. The resulting NetEgg [66, 67] tool allows network operators to program network policies by describing representative example behaviors, significantly easing the process of synthesizing new protocols.

3 Thrust 2: Debugging Distributed Systems

The second research thrust focuses on providing runtime support for network accountability, forensic analysis, and failure diagnosis. These capabilities are becoming increasingly important for network management and security. While our work on FSR provides mechanisms for formally proving convergence, often times, deployed systems exhibit unexpected interactions that are not easy to catch solely through the use of formal tools at design time. Mechanisms for runtime debugging would typically be required to complement formal techniques adopted at design time.

To address the above challenges, we have developed *NetTrails* [88, 91], a declarative engine that allows a network operator to issue queries explaining the derivation and change of network state at any given node in a distributed system. NetTrails provides support for *network provenance* – the ability to issue queries over network meta-data. For example, network provenance may be used to trace the path a message traverses on the network, as well as to explain how a given network state was derived and which parties were involved in its derivation. NetTrails uses declarative networking as a basis for maintaining and querying network provenance at Internet-scale. In developing NetTrails, in additional to introducing declarative general-purpose capabilities for network forensics, we have made new contributions to the data provenance literature, by extending existing data provenance with notions of distribution [91] and time-awareness [85,90].

As long as there are no faulty nodes, network provenance techniques can construct such explanations. However, if some of the nodes are faulty or have been compromised by an adversary, the situation is complicated by the fact that the adversary can cause the nodes under his control to lie, suppress information, tamper with existing data, or report nonexistent events. This can cause the provenance system to turn from an advantage into a liability: its answers may cause operators to stop investigating an ongoing attack because everything looks fine. To address this challenge, we have developed *secure network provenance* (SNP) [87], which allows network provenance to be securely stored and queried even in the presence of Byzantine faults. Demonstrating its wide applicability, we have evaluated SNP in a variety of concrete applications, such as Hadoop MapReduce, declarative Chord distributed hash table, and BGP interdomain routing.

Beyond SNP, we have extended the network provenance work in several directions, including combining privacy and verifiability [19, 83, 84] in the context of BGP routing systems, *negative provenance* [64, 65] for explaining why certain events did not occur, *differential provenance* [5, 6] for diagnosing root causes of faults using a correct reference configuration, and *meta-provenance* [63] for automatically repairing faulty configurations using meta-programming techniques. As an alternative to the use of provenance for debugging, we have also started combining the use of declarative networking with multicore processing techniques, in order to enable high-performance extensible deep packet analysis [13–16] at the data plane. The resulting analytics platform has resulted in a spinoff company (Section 7). Finally, in collaboration with Microsoft, we developed NetPoirot [2], a debugging tool that can identify root cause analysis of failures in data centers solely by using TCP statistics at the end-points. NetPoirot uses machine-learning based classification techniques, is light-weight and non-intrusive, requiring only aggregate TCP statistics to be periodically collected and measured.

4 Thrust 3: Securing Distributed Systems

Our third research thrust explores high-level programming abstractions for implementing secure distributed systems, achieved by unifying declarative networking and logic-based access control specifications [1, 89, 92]. Our work is motivated by the proliferation of large-scale network information systems currently deployed for a variety of application domains including network monitoring infrastructures, publish-subscribe systems, cloud computing, content distribution networks, and network routing. Despite their widespread usage, designing and implementing these large-scale systems remains a challenge, in part because of the sheer scale of deployment, but also due to emerging security threats.

We aim to address two challenges: (1) *securing* these distributed data-centric computations, particularly when the computations span administrative boundaries, and (2) *analyzing* the security properties of these systems. Achieving secure distributed computations may require authentication, encryption, and integrity-checking of inter-node communications, as well as intra-node access control. Although it may be possible to achieve a high level of security using a "one-size-fits-all" solution that imposes particular cryptographic protocols on top of network communication, we argue that providing programmers with a flexible security policy framework enables security that is better suited for a network, given its resource constraints and security requirements.

We have developed the *Secure Network Datalog* (SeNDlog) language [89, 92] that unifies the Network Datalog (NDlog) language used in declarative networking and Binder, a logic-based language for access control in distributed systems. DS2 has a wide range of applications, including reconfigurable trust management [44], secure distributed data processing [43], and tunable anonymity [51–53]. In combining thrust 1 and thrust 3, we have applied a wide range of formal analysis techniques for analyzing security properties of distributed systems expressed in SeNDlog [7–9, 22].

As one example use case, the *Application-Aware Anonymity (A3)* system [49–53] is an extensible platform for applications to deploy anonymity-based services on the Internet. A3 allows applications to tailor their anonymity properties and performance characteristics according to their specific communication requirements. To support flexible path construction, A3 exposes a declarative language similar to SeNDlog that enables applications to compactly specify path selection and instantiation policies which are then executed using a declarative networking engine. In addition to specifying relay selection strategies, senders are able to use our declarative techniques to construct anonymous tunnels according to their specifications (for example, via *Onion Routing* or *Crowds*). We have integrated the A3 system with the popular open-source Tor anonymity software, and have deployed our system on PlanetLab and the DETER security testbed. We have also extended A3's approach to work for group communication protocols [25].

5 Thrust 4: Optimizing Distributed Systems

In distributed systems management, operators often configure system parameters that optimize performance objectives, given constraints in the deployment environment. Traditional optimization approaches use imperative languages such as C++ or Java and often result in cumbersome and error-prone programs that are difficult to maintain and customize. Moreover, due to scalability and management constraints imposed across administrative domains, it is often necessary to optimize in a *distributed* setting in which multiple *local* solvers must coordinate with one another. Each local optimizer handles a portion of the whole problem, and they together achieve a global objective.

To meet the above challenges, the fourth research thrust centers around *Cologne* [33] (*COnstraint LOGic EngiNE*), a distributed declarative constraint optimization platform that enables constraint optimization problems (COP) to be declaratively specified and incrementally executed in distributed systems. Cologne integrates a *declarative networking engine* [37] with an off-the-shelf constraint solver. We highlight two representative use cases to which we have applied our platform.

First, we have developed the *Policy-based Unified Multi-radio Architecture* (PUMA), a declarative constraint solving platform for optimizing wireless mesh networks. In PUMA, network operators can flexibly vary the choice of routing via adaptable *hybrid* routing protocols [28, 29]. The hybrid technique combines several existing protocols (e.g., proactive, reactive, and epidemic) with specific criteria for determining when particular protocols are to be used. The hybrid compositional capabilities are particularly useful for routing in heterogeneous network settings in which application requirements and network conditions continuously change over time. In addition, PUMA enables policies for *wireless channel selection* [26, 27] to be declaratively specified and optimized; such policies may reduce network interference and maximize throughput while not violating deployment constraints. PUMA has been successfully deployed on the ORBIT wireless testbed.

Second, in our *Cloud Orchestration Policy Engine* (COPE) [30], we use our optimization framework to declaratively control the provisioning, configuration, management and decommissioning of cloud resource orchestration. COPE enables the automatic realization of customer service level agreements while simultaneously conforming to operational objectives of the cloud providers. COPE is developed in collaboration with AT&T Labs Research, in a joint research effort at developing TROPIC [31], a transactional cloud resource orchestration platform [31]. We have experimentally deployed TROPIC and COPE on AT&T's ShadowNet testbed. A follow-up joint project with AT&T further resulted in a cross data-center deployment [10].

Beyond the above use cases, we have explored the other domains to apply the constraint optimizations work, namely in the context of bandwidth allocation in data centers [68], green data centers [23], real-time network functions virtualization [24], real-time big data analytics [47, 69–81], and transport protocols [3, 4].

Generalizing the ideas from Cologne, in the context of adaptive query processing in stream processing systems, we developed a declarative cost-based optimizer [21] that recomputes the optimal execution plan incrementally given new cost information. Our cost-based optimization is written in Datalog, and by leveraging incremental recursive view maintenance techniques [34], we incrementally compute the optimal query plan by recomputing only portions of the search space.

6 Research Mentoring and Student Development

An important aspect of my research program is the development of doctoral students into independent researchers, as well as promoting research interests at the undergraduate and masters level. My research group currently comprises of five doctoral students and several masters students and undergraduates. I have graduated eight doctoral students over the past nine years, and my 9th PhD student will graduate by the end of summer of 2016. Three of my graduated PhD students have received dissertation awards: Micah Sherr (Ph.D. 2009) received the *Rubinoff award* for the most outstanding dissertation in Computer and Information Science; Wenchao Zhou (Ph.D. 2012) received honorable mention for the *ACM SIGMOD dissertation award* for outstanding dissertation in the field of data management; and finally, Zhuoyao Zhang (Ph.D. 2014, female) received honorable mention for the *SPEC dissertation award* for outstanding dissertation in experimental system analysis.

Micah Sherr is now a tenured Associate Professor at Georgetown University. Wenchao Zhou is a tenure-track Assistant Professor at Georgetown University, while Anduo Wang (Ph.D 2013, female) recently joined Temple University as a tenure-track Assistant Professor, after a post-doctoral stint at UIUC. Several of my Ph.D. students are now in the industry. Yun Mao (Ph.D. 2008) joined AT&T Labs Research and is now leading a team at Facebook. Changbin Liu (Ph.D. 2012) similarly joined AT&T Labs Research and recently co-founded a big data storage company called Termaxia with me. Mengmeng Liu (Ph.D. 2016), Dong Lin (Ph.D. 2015) and Zhuoyao Zhang (Ph.D. 2014) are at Walmart Labs, LinkedIn, and Google respectively.

I actively promote the involvement of female students and undergraduates in my research group. In total, there are fifteen female research participants involved in various aspects of my research. Among my Ph.D. students past and present, five of them are female (Behnaz Arzani, Miao Cheng, Mengmeng Liu, Anduo Wang, and Zhuoyao Zhang). Mengmeng Liu and Zhuoyao Zhang received best papers for ICDE'09 [34] and ICAC'12 respectively. Two female undergraduates (Sandy Sun and Trisha Kothari) have published papers [14, 88] as undergraduates and are recent winners of Microsoft undergraduate fellowships. Three female Master students (Qiong Fei, Yiqing Ren, and Karen Tao) played key roles in major ACM publications and demonstrations [48, 58, 87, 88, 91]. Sangeetha Joshi graduated from my research group and is a Ph.D. student at UIUC.

In terms of undergraduate research participants, I have supervised 60 undergraduates on various senior design projects, of which, several are directly affiliated with NetDB@Penn. Penn undergraduate Bill Marczak [43, 44] received honorable mention for the CRA outstanding award. Students graduating from my group are now pursuing Ph.D. studies at top universities such as Harvard, MIT, Princeton, University of Washington Seattle, and UC Berkeley, and have published papers [14, 43, 44, 53, 88] as undergraduates. Several undergraduates and masters students regularly participate in conference demonstrations.

7 Research Impact, Tech Transfer, and Future Directions

My research group's widespread impact is shown not only in its publications across several subfields in computer science, but also in its impact on industry and the research community at large. My research publications have been cited 5577 times (Google scholar Mar 4 2016) in total, and 2517 (since 2011).

We strive towards open-source releases of our research prototypes. These includes the A3 [53] system, Veracity [50], RapidNet [45], and NetTrails [88]. The NetEgg project will also be turned into an open-source release soon.

We have participated in multi-PI research initiatives that use declarative networking. These include:

- DARPA's *XD3* project, which I am leading as PI, that uses declarative component-based techniques to mitigate distributed denial of service (DDoS) attacks.
- DARPA's *SAFER* initiative on anonymous communications, in collaboration with Georgetown University.
- DARPA's *Wireless Networks After Next* (WNaN) program on policy-based adaptive wireless communications, in collaboration with Raytheon BBN Technologies and several other institutions.
- *NEBULA*, an NSF-funded multi-university effort led by Penn on a Future Internet Architecture (FIA) for secure cloud computing.
- *Expeditions on Computer Augmented Program Engineering* (ExCAPE), NSF Expeditions on program synthesis led by Penn in collaboration with several institutions.
- The *Presidio* MURI on collaborative policies and assured information sharing, in collaboration with Stanford University and several other institutions.

Overall, NSF, DoD, and the industry has recognized the significance and inter-disciplinary nature of my work, by awarding 26 research gifts/grants (13 as PI) in the past 9 years, including the NSF CAREER award and the AFOSR Young Investigator award. My NSF funding is cross-cutting from different programs in networking, security, databases, programming languages, and industry innovations/partnerships. I have also received competitive awards from Amazon, AT&T, and Google. In total, I have raised $24M over nine years in research funding in collaboration with my colleagues at Penn and at other institutions.

Given the inter-disciplinary nature of my research, I have also been invited to program committees of top conferences across different communities. These include systems (OSDI, USENIX ATC, SOCC), security (PETS), databases (SIGMOD, VLDB, ICDE, CIDR), networking (NSDI, CoNEXT, INFOCOM, SOSR) and programming languages (ICLP). NSF has also recognized my research breadth and invited me to serve on panels across multiple different programs.

We have also collaborated with several industry partners, including AT&T Research [10,30–33], Intel [82], LogicBlox [43, 44], Microsoft Research [89], Raytheon BBN Technologies [27–29, 33], SRI International [58, 60, 61], and HP Labs [69–81]. Our work [43, 44] has been integrated into LogicBlox's commercial data management system.

Finally, our work has also resulted in successful technology transfer. For example, Scalanytics [13, 15] has been commercialized as Gencore (rebranded as Netsil Inc.), a NSF-funded and venture-backed San Francisco based cloud microservices analytics company that I co-founded. A more recent research project, NetEgg [66, 67] participated in the NSF Innovation Corps (ICorps) program, geared towards identifying commercialization and licensing opportunities for selected university research. With former Ph.D. student Changbin Liu, I recently co-founded Philly-based Termaxia, a big data storage company.

Future directions. In the past decade, I have successfully carried out cross-cutting research *within* the computer science field, for example, bridging research in formal methods, programming languages, security, and databases. In the process of doing so, I have collaborated widely and published in a wide range of venues in computer science. Moving forward, I am reshaping my research agenda over the next decade, to aim to reach out to scientific communities *outside* of computer science. For example, I recently started a new research initiative (with two pending NSF grants), leading an inter-disciplinary team of researchers from Wharton's management, marketing, operations research, and statistics department, and Penn's Annenberg's School of Communication. We are developing a next-generation big data platform geared towards storing and analyzing large volumes of temporal social network data collected over time. We have developed an initial proof-of-concept system [12] specifically targeted at analyzing the impact of social network engagements on startup crowdfunding success. Generalizing beyond our initial work, the eventual platform will support several use cases by social scientists across advertising, business, marketing, and mass communications, The project integrates several core research ideas developed by NetDB@Penn over the past decade (namely declarative query languages, distributed provenance, large-scale data management and scalable cloud and data center design), with the intent of solving the data collection, cleaning, and analysis bottlenecks faced by social scientists.

References

[1] Abadi, M., and Loo, B. T. Towards a Language and System for Secure Networking. In *NetDB* (2007).

[2] Arzani, B., Ciraci, S., Loo, B. T., Schuster, A., and Outhred, G. Taking the blame game out of data centers operations with netpoirot. In *ACM SIGCOMM Conference on Data Communication* (2016).

[3] Arzani, B., Gurney, A., Cheng, S., Guerin, R., and Loo, B. T. Deconstructing mptcp performance. In *22nd IEEE International Conference on Network Protocols (ICNP)* (2014).

[4] Arzani, B., Gurney, A., Cheng, S., Guerin, R., and Loo, B. T. Impact of path selection and scheduling policies on mptcp performance. In *4th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS)* (2014).

[5] Chen, A., Wu, Y., Haeberlen, A., Zhou, W., and Loo, B. T. Differential provenance: Better network diagnostics with reference events. In *14th ACM Workshop on Hot Topics in Networks (HotNets)* (2015).

[6] Chen, A., Wu, Y., Haeberlen, A., Zhou, W., and Loo, B. T. The good, the bad, and the differences: Better network diagnostics with differential provenance. In *ACM SIGCOMM Conference on Data Communication* (2016).

[7] Chen, C., Jia, L., Loo, B. T., and Zhou, W. Reduction-based security analysis of internet routing protocols. In *2nd International Workshop on Rigorous Protocol Engineering (WRiPE)* (2012).

[8] Chen, C., Jia, L., Xu, H., Luo, C., Zhou, W., and Loo, B. T. A program logic for verifying secure routing protocols. In *34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems (FORTE)* (2014).

[9] Chen, C., Jia, L., Xu, H., Luo, C., Zhou, W., and Loo, B. T. A program logic for verifying secure routing protocols. *Logical Methods in Computer Science Journal* (2015).

[10] Chen, C., Liu, C., Liu, P., Loo, B. T., and Ding, L. A scalable multi-datacenter layer-2 network architecture. In *Symposium on SDN Research (SOSR)* (2015).

[11] Chen, C., Loh, L. K., Jia, L., Zhou, W., and Loo, B. T. Automated verification of safety properties in declarative networking programs. In *17th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP)* (2015).

[12] Cheng, M., Sriramulu, A., Muralidhar, S., Loo, B. T., Huang, L., and Loh, P.-L. Collection, Exploration and Analysis of Crowdfunding Social Networks. . In *3rd International Workshop on Exploratory Search in Databases and the Web (ExploreDB), co-located with SIGMOD* (2016).

[13] Gill, H., Lin, D., Han, X., Nguyen, C., Gill, T., and Loo, B. T. Scalanytics: A declarative multi-core platform for scalable composable traffic analytics. In *22nd International ACM Symposium on High Performance and Distributed Computing (HPDC)* (2013).

[14] Gill, H., Lin, D., Kothari, T., and Loo, B. T. Declarative multicore programming of software-based stateful packet processing. In *Declarative Aspects and Applications of Multicore Programming (DAMP), co-located with POPL* (2012).

[15] Gill, H., Lin, D., Nguyen, C., Gill, T., and Loo, B. T. Declarative platform for high-performance network traffic analytics. *Cluster Computing journal* (2014).

[16] Gill, H., Lin, D., Sarna, L., Mead, R., Lee, K. C., and Loo, B. T. Sp4: Scalable programmable packet processing platform. In *ACM SIGCOMM Conference on Data Communication (demonstration)* (2012).

[17] Green, T. J., Huang, S. S., Loo, B. T., and Zhou, W. Datalog and recursive query processing. *Foundations and Trends in Databases* (2013).

[18] Gurney, A. J., Han, X., Li, Y., and Loo, B. T. Route shepherd: Stability hints for the control plane. In *ACM SIGCOMM Conference on Data Communication (demonstration)* (2012).

[19] Gurney, A. J. T., Haeberlen, A., Zhou, W., Sherr, M., and Loo, B. T. Having your Cake and Eating it too: Routing Security with Privacy Protections. In *10th ACM Workshop on Hot Topics in Networks (HotNets-X)* (2011).

[20] Gurney, A. J. T., Jia, L., Wang, A., and Loo, B. T. Partial specifications of routing configurations. In *1st International Workshop on Rigorous Protocol Engineering (WRiPE), co-located with ICNP* (2011).

[21] Ives, M. L. Z., and Loo, B. T. Enabling Incremental Query Re-Optimization. In *Proceedings of ACM SIGMOD International Conference on Management of Data* (2016).

[22] Jia, L., Chen, C., Jyothi, S. A., Zhou, W., Mapara, S., and Loo, B. T. Towards a secure and verifiable future internet. In *Off the Beaten Track: Underrepresented Problems for Programming Language Researchers, co-located with POPL* (2012).

[23] Li, Y., Chiu, D., Liu, C., Phan, L. T., Gill, T., Aggarwal, S., Zhang, Z., Loo, B. T., Maier, D., and McManus, B. Towards dynamic pricing-based collaborative optimizations for green data centers. In *Second International Workshop on Data Management in the Cloud (DMC)* (2013).

[24] Li, Y., Phan, L. T., and Loo, B. T. Network functions virtualization with soft real-time guarantees. In *IEEE International Conference on Computer Communications (INFOCOM)* (2016).

[25] Lin, D., Sherr, M., and Loo, B. T. Scalable and anonymous group communication with mtor. *Proceedings on Privacy Enhancing Technologies (PoPETS)* (2016).

[26] Liu, C., Correa, R., Gill, H., Gill, T., Li, X., Muthukumar, S., Saeed, T., Loo, B. T., and Basu, P. PUMA: Policy-based Unified Multi-radio Architecture for Agile Mesh Networking. In *4th International Conference on Communication Systems and Networks (COMSNETS) (paper and demonstration)* (2012).

[27] Liu, C., Correa, R., Gill, H., Gill, T., Li, X., Muthukumar, S., Saeed, T., Loo, B. T., and Basu, P. Puma: Policy-based unified multi-radio architecture for agile mesh networking. *IEEE/ACM Transactions on Networking (ToN)* (2013).

[28] Liu, C., Correa, R., Li, X., Basu, P., Loo, B. T., and Mao, Y. Declarative Policy-based Adaptive MANET Routing. In *17th IEEE International Conference on Network Protocols (ICNP)* (2009).

[29] Liu, C., Correa, R., Li, X., Basu, P., Loo, B. T., and Mao, Y. Declarative Policy-based Adaptive Mobile Ad Hoc Networking. *IEEE/ACM Transactions on Networking (ToN)* (2011).

[30] Liu, C., Loo, B. T., and Mao, Y. Declarative Automated Cloud Resource Orchestration. In *ACM Symposium on Cloud Computing (SOCC)* (2011).

[31] Liu, C., Mao, Y., Chen, X., Fernadez, M. F., Loo, B. T., and der Merwe, J. E. V. TROPIC: Transactional Resource Orchestration Platform In the Cloud. In *USENIX Annual Technical Conference (USENIX ATC)* (2012).

[32] Liu, C., Mao, Y., Chen, X., Fernandez, M. F., Loo, B. T., and der Merwe, K. V. Towards Transactional Cloud Resource Orchestration, 2011.

[33] Liu, C., Ren, L., Loo, B. T., Mao, Y., and Basu, P. Cologne: A Declarative Distributed Constraint Optimization Platform. In *Proceedings of VLDB Conference* (2012).

[34] Liu, M., Taylor, N., Zhou, W., Ives, Z., and Loo, B. T. Recursive Computation of Regions and Connectivity in Networks. In *Proceedings of IEEE Conference on Data Engineering (ICDE)* (2009).

[35] Liu, M., Zhou, W., Taylor, N., Ives, Z., and Loo, B. T. Maintaining Recursive Views of Regions and Connectivity in Networks. *IEEE Transaction on Knowledge and Data Engineering (TKDE)* (2009). Special issue on best papers of ICDE 2009.

[36] Loo, B. T., Condie, T., Garofalakis, M., Gay, D. E., Hellerstein, J. M., Maniatis, P., Ramakrishnan, R., Roscoe, T., and Stoica, I. Declarative Networking: Language, Execution and Optimization. In *Proceedings of ACM SIGMOD International Conference on Management of Data* (2006).

[37] Loo, B. T., Condie, T., Garofalakis, M., Gay, D. E., Hellerstein, J. M., Maniatis, P., Ramakrishnan, R., Roscoe, T., and Stoica., I. Declarative Networking. In *Communications of the ACM (CACM)* (2009).

[38] Loo, B. T., Condie, T., Hellerstein, J. M., Maniatis, P., Roscoe, T., and Stoica, I. Implementing Declarative Overlays. In *Proceedings of ACM Symposium on Operating Systems Principles* (2005).

[39] Loo, B. T., Gill, H., Liu, C., Mao, Y., Marczak, W. R., Sherr, M., Wang, A., and Zhou, W. Recent Advances in Declarative Networking. In *Fourteenth International Symposium on Practical Aspects of Declarative Languages (PADL), colocated with POPL*. Invited paper and keynote speech.

[40] Loo, B. T., Hellerstein, J. M., Stoica, I., and Ramakrishnan, R. Declarative Routing: Extensible Routing with Declarative Queries. In *Proceedings of ACM SIGCOMM Conference on Data Communication* (2005).

[41] Loo, B. T., and Zhou, W. *Declarative Networking*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2012.

[42] Mao, Y., Loo, B. T., Ives, Z., and Smith, J. M. MOSAIC: Unified Platform for Dynamic Overlay Selection and Composition. In *CoNEXT* (2008).

[43] Marczak, W. R., Huang, S. S., Bravenboer, M., Sherr, M., Loo, B. T., and Aref, M. SecureBlox: Customizable Secure Distributed Data Processing. In *SIGMOD* (2010).

[44] Marczak, W. R., Zook, D., Zhou, W., Aref, M., and Loo, B. T. Declarative Reconfigurable Trust Management. In *Proceedings of Conference on Innovative Data Systems Research (CIDR)* (2009).

[45] Muthukumar, S. C., Li, X., Liu, C., Kopena, J. B., Oprea, M., and Loo, B. T. Declarative toolkit for rapid network protocol simulation and experimentation. In *SIGCOMM (demo)* (2009).

[46] Nigam, V., Jia, L., Loo, B. T., and Scedrov, A. Maintaining Distributed Logic Programs Incrementally. In *13th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP)* (2011).

[47] Phan, L. P., Zhang, Z., Zheng, Q., Loo, B. T., and Lee, I. An Empirical Analysis of Scheduling Techniques for Real-time Cloud-based Data Processing. In *4th IEEE International Workshop on Real-time Service-Oriented Architecture and Application (RTSOAA)* (2012).

[48] Ren, Y., Zhou, W., Wang, A., Jia, L., Gurney, A. J., Loo, B. T., and Rexford, J. FSR: Formal Analysis and Implementation Toolkit for Safe Inter-domain Routing. In *ACM SIGCOMM Conference on Data Communication (demonstration)* (2011).

[49] Sherr, M., Blaze, M., and Loo, B. T. Scalable Link-Based Relay Selection for Anonymous Routing. In *9th Privacy Enhancing Technologies Symposium* (August 2009).

[50] Sherr, M., Blaze, M., and Loo, B. T. Veracity: Practical Secure Network Coordinates via Vote-Based Agreements. In *USENIX Annual Technical Conference (USENIX ATC)* (June 2009).

[51] Sherr, M., Gill, H., Saeed, T. A., Mao, A., Marczak, W. R., Soundararajan, S., Zhou, W., Loo, B. T., and Blaze, M. The design and implementation of the a3 application-aware anonymity platform. *Computer Networks (COMNET), Elsevier Publishing* (2013).

[52] Sherr, M., Loo, B. T., and Blaze, M. Towards Application-Aware Anonymous Routing. In *2nd USENIX Workshop on Hot Topics in Security (HotSec)* (2007).

[53] Sherr, M., Mao, A., Marczak, W. R., Zhou, W., Loo, B. T., and Blaze, M. A3: An Extensible Platform for Application-Aware Anonymity. In *Network and Distributed System Security* (2010).

[54] Wang, A., Basu, P., Loo, B. T., and Sokolsky, O. Towards Declarative Network Verification. In *11th International Symposium on Practical Aspects of Declarative Languages (PADL)* (2009).

[55] Wang, A., Gurney, A., Han, X., Cao, J., Loo, B. T., Talcott, C., and Scedrov, A. A reduction-based approach towards scaling up formal analysis of internet configurations. In *33rd Annual IEEE International Conference on Computer Communications (INFOCOM)* (2014).

[56] Wang, A., Gurney, A. J., Han, X., Cao, J., Talcott, C., Loo, B. T., and Scedrov, A. Reduction-based analysis of bgp systems with bgpverif. In *ACM SIGCOMM Conference on Data Communication (demonstration)* (2012).

[57] Wang, A., Jia, L., Liu, C., Loo, B. T., Sokolsky, O., and Basu, P. Formally Verifiable Networking. In *SIGCOMM HotNets-VIII* (2009).

[58] Wang, A., Jia, L., Zhou, W., Ren, Y., Loo, B. T., Rexford, J., Nigam, V., Scedrov, A., and Talcott, C. FSR: Formal analysis and implementation toolkit for safe inter-domain routing. *IEEE/ACM Transactions on Networking* (2012).

[59] Wang, A., Loo, B. T., Liu, C., Sokolsky, O., and Basu, P. A Theorem Proving Approach towards Declarative Networking. In *22nd International Conference on Theorem Proving in Higher Order Logics (TPHOLs) emerging trends proceedings* (2009).

[60] Wang, A., Talcott, C., Gurney, A. J., Loo, B. T., and Scedrov, A. A Calculus of Policy-Based Routing Systems. In *31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC) (Brief Announcement)* (2012).

[61] Wang, A., Talcott, C., Gurney, A. J. T., Loo, B. T., and Scedrov, A. Reduction-based Formal Analysis of BGP Instances. In *18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)* (2012).

[62] Wang, A., Talcott, C., Jia, L., Loo, B. T., and Scedrov, A. Analyzing BGP Instances in Maude. In *International Conference on Formal Techniques for Networked and Distributed Systems (FMOODS/FORTE)* (2011).

[63] Wu, Y., Chen, A., Haeberlen, A., Zhou, W., and Loo, B. T. Automated network repair with meta provenance. In *14th ACM Workshop on Hot Topics in Networks (HotNets)* (2015).

[64] Wu, Y., Haeberlen, A., Zhou, W., and Loo, B. T. Answering why-not queries in software-defined networks with negative provenance. In *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)* (2013).

[65] Wu, Y., Zhao, M., Haeberlen, A., Zhou, W., and Loo, B. T. Diagnosing missing events in distributed systems with negative provenance. In *ACM SIGCOMM Conference on Data Communication* (2014).

[66] Yuan, Y., Alur, R., and Loo, B. T. Netegg: Programming network policies by examples. In *13th ACM Workshop on Hot Topics in Networks (HotNets-XIV)* (2014).

[67] Yuan, Y., Lin, D., Alur, R., and Loo, B. T. Scenario-based programming for sdn policies. In *11th Conference on emerging Networking EXperiments and Technologies (ACM CoNEXT)* (2015).

[68] Yuan, Y., Wang, A., Alur, R., and Loo, B. T. On the feasibility of automation for bandwidth allocation problems in data centers. In *Formal Methods in Computer-Aided Design (FMCAD)* (2013).

[69] Zhang, Z., Cherkasova, L., and Loo, B. T. Autotune: Optimizing execution concurrency and resource usage in mapreduce workflows. In *9th ACM International Conference on Autonomic Computing (ICAC) Management of Big Data Systems (MBDS) track* (2013).

[70] Zhang, Z., Cherkasova, L., and Loo, B. T. Benchmarking approach for designing a mapreduce performance model. In *4th ACM/SPEC International Conference on Performance Engineering (ICPE)* (2013).

[71] Zhang, Z., Cherkasova, L., and Loo, B. T. Getting more for less in optimized mapreduce workflows. In *IFIP/IEEE Integrated Network Management Symposium (IM)* (2013).

[72] Zhang, Z., Cherkasova, L., and Loo, B. T. Performance modeling of mapreduce jobs in heterogeneous environments. In *IEEE 6th International Conference on Cloud Computing (CLOUD)* (2013).

[73] Zhang, Z., Cherkasova, L., and Loo, B. T. Exploiting cloud heterogeneity for optimized cost/performance mapreduce processing. In *Proceedings of the Fourth International Workshop on Cloud Data and Platforms (CloudDP)* (2014).

[74] Zhang, Z., Cherkasova, L., and Loo, B. T. Optimizing cost and performance trade-offs for mapreduce job processing in the cloud. In *IFIP/IEEE Network Operations and Management Symposium (NOMS)* (2014).

[75] Zhang, Z., Cherkasova, L., and Loo, B. T. Parameterizable benchmarking framework for designing a mapreduce performance model. *Concurrency and Computation: Practice and Experience journal 26*, 12 (Aug. 2014), 2005–2026.

[76] Zhang, Z., Cherkasova, L., and Loo, B. T. Automating platform selection for mapreduce processing in the cloud. In *IEEE International Conference on Cloud and Autonomic Computing (ICCAC)* (2015).

[77] Zhang, Z., Cherkasova, L., and Loo, B. T. Exploiting cloud heterogeneity to optimize performance and cost of mapreduce processing. *SIGMETRICS Perform. Eval. Rev. 42*, 4 (June 2015), 38–50.

[78] Zhang, Z., Cherkasova, L., Verma, A., and Loo, B. T. Automated Profiling and Resource Management of Pig Programs for Meeting Service Level Objectives. In *9th ACM International Conference on Autonomic Computing (ICAC)* (2012).

[79] Zhang, Z., Cherkasova, L., Verma, A., and Loo, B. T. Meeting Service Level Objectives of Pig Programs. In *2nd International Workshop on Cloud Computing Platforms (CloudCP)* (2012).

[80] Zhang, Z., Cherkasova, L., Verma, A., and Loo, B. T. Optimizing Completion Time and Resource Provisioning of Pig Programs. In *Workshop on Cloud Computing Optimization (CCOPT)* (2012).

[81] Zhang, Z., Cherkasova, L., Verma, A., and Loo, B. T. Performance modeling and optimization of deadline-driven pig programs. *ACM Transactions on Autonomous and Adaptive Systems (TAAS) 8*, 3 (Sept. 2013), 14:1–14:28.

[82] Zhang, Z., Phan, L. T., Tan, G., Jain, S., Duong, H., Loo, B. T., and Lee, I. On the Feasibility of Dynamic Rescheduling on the Intel Distributed Computing Platform. In *ACM/IFIP/USENIX 11th International Middleware Conference (Industry Track)* (2010).

[83] Zhao, M., Zhou, W., Gurney, A. J., Haeberlen, A., Sherr, M., and Loo, B. T. Private and Verifiable Interdomain Routing Decisions. In *Proceedings of ACM SIGCOMM Conference on Data Communication* (2012).

[84] Zhao, M., Zhou, W., Gurney, A. J. T., Haeberlen, A., Sherr, M., and Loo, B. T. Private and verifiable interdomain routing decisions. *IEEE/ACM Transactions on Networking (ToN)* (2015).

[85] Zhou, W., Ding, L., Haeberlen, A., Ives, Z., and Loo, B. T. TAP: Time-aware Provenance for Distributed Systems. In *3rd USENIX Workshop on the Theory and Practice of Provenance (TaPP '11)* (2011).

[86] Zhou, W., Fei, Q., Narayan, A., Haeberlen, A., Loo, B. T., and Sherr, M. Secure Forensics without Trusted Components. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI) demo* (2011).

[87] Zhou, W., Fei, Q., Narayan, A., Haeberlen, A., Loo, B. T., and Sherr, M. Secure Network Provenance. In *Proceedings of ACM Symposium on Operating Systems Principles* (2011).

[88] Zhou, W., Fei, Q., Sun, S., Tao, T., Haeberlen, A., Ives, Z., Loo, B. T., and Sherr, M. NetTrails: A Declarative Platform for Provenance Maintenance and Querying in Distributed Systems. In *SIGMOD (demonstration)* (2011).

[89] Zhou, W., Mao, Y., Loo, B. T., and Abadi, M. Unified Declarative Platform for Secure Networked Information Systems. In *Proceedings of IEEE Conference on Data Engineering (ICDE)* (2009).

[90] Zhou, W., Mapara, S., Ren, Y., Li, Y., Haeberlen, A., Ives, Z., Loo, B. T., and Sherr, M. Distributed Time-aware Provenance. In *39th International Conference on Very Large Databases (VLDB)* (2013).

[91] Zhou, W., Sherr, M., Tao, T., Li, X., Loo, B. T., and Mao, Y. Efficient Querying and Maintenance of Network Provenance at Internet-Scale. In *Proc. SIGMOD* (2010).

[92] Zhou, W., Tao, T., Loo, B. T., and Mao, Y. Declarative secure distributed information systems. *Computer Languages, Systems and Structures (COMLAN), Elsevier Publishing 39*, 1 (Apr. 2013), 1–24.