

Department of Computer & Information Science

Departmental Papers (CIS)

University of Pennsylvania

Year 2008

Provenance-aware Secure Networks

Wenchao Zhu*

Eric Cronin†

Boon Thau Loo‡

*University of Pennsylvania

†University of Pennsylvania

‡University of Pennsylvania, boonloo@cis.upenn.edu

Copyright 2008 IEEE. Reprinted from *Proceedings of the 24th IEEE International Conference on Data Engineering (ICDE 2008)*, pages 188-193.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons.

http://repository.upenn.edu/cis_papers/387

Provenance-aware Secure Networks

Wenchao Zhou Eric Cronin Boon Thau Loo
University of Pennsylvania

Abstract

Network accountability and forensic analysis have become increasingly important, as a means of performing network diagnostics, identifying malicious nodes, enforcing trust management policies, and imposing diverse billing over the Internet. This has led to a series of work to provide better network support for accountability, and efficient mechanisms to trace packets and information flows through the Internet. In this paper, we make the following contributions. First, we show that network accountability and forensic analysis can be posed generally as data provenance computations and queries over distributed streams. In particular, one can utilize declarative networks with appropriate security and provenance extensions to provide a unified declarative framework for specifying, analyzing and auditing networks. Second, we propose a taxonomy of data provenance along multiple axes, and show that they map naturally to different use cases in networks. Third, we suggest techniques to efficiently compute and store network provenance, and provide an initial performance evaluation on the P2 declarative networking system with modifications to support authenticated communication and provenance.

1 Introduction

The Internet was not designed with accountability as its primary goal. However, network accountability and forensic analysis have become increasingly important in recent years, as a means of performing network diagnostics, identifying malicious and misbehaving users, enforcing trust management policies, and imposing diverse billing over the Internet. This has led to a series of proposals (e.g. [22, 3, 23, 13, 26, 4, 11, 14]) on improving network support for accountability, and efficient mechanisms to trace packets and information flows through the Internet. While there have not been a lack of proposals, several of them narrowly tackle a specific security challenge, or target a limited set of network applications.

Provenance (also called *lineage*) has been studied in many different contexts. In the context of database systems, they have primarily used in databases to help “explain” to users why a tuple exists [7]. In this paper, we show that network accountability and forensic analysis can be posed generally as data provenance computations and queries over distributed streams. We argue that declarative networks [15, 17, 16] enhanced with the ability to maintain provenance of computations will enable a general extensible framework for specifying, analyzing, and auditing networks. Declarative networks utilize a database query language for specifying and implement-

ing networks, and its dataflow framework captures information flow naturally as distributed streams computations. We further demonstrate that with the appropriate security extensions [1] to the query language used in declarative networks, we can further allow provenance computations and queries to be authenticated in untrusted environments.

Contributions and Organization: In Section 2, we provide a background on declarative networks, its query languages, and recent security extensions obtained by unifying its core language with logic-based access control languages. Next, in Section 3, we survey various use cases of network provenance ranging from real-diagnostics, forensics, accountability, and trust management. In Section 4, we then provide a taxonomy of different types of data provenance (local vs distributed, on-line vs offline, authenticated, etc), several of which are derived from existing database literature, and show that they map naturally into existing use cases. We outline some possible optimizations (Section 5), perform initial performance evaluations based on extensions to the P2 declarative networking system (Section 6), and then conclude in Section 7.

2 Declarative Networks

As background, we briefly introduce declarative networking and its query language, including security extensions. The high level goal of *declarative networks* [17, 16, 15] is to build extensible network architectures that achieve a good balance of flexibility, performance and safety. Declarative networks are specified using *Network Datalog (NDlog)*, which is a distributed recursive query language used for querying network graphs. *NDlog* queries are executed using a distributed query processor to implement the network protocols, and continuously maintained as distributed views over existing network and host state. Declarative queries such as *NDlog* are a natural and compact way to implement a variety of routing protocols and overlay networks. For example, traditional routing protocols can be expressed in a few lines of code [17], and the Chord [25] distributed hash table in 47 lines of code [16]. When compiled and executed, these declarative networks perform efficiently relative to imperative implementations.

2.1 Network Datalog Language

NDlog is based on Datalog [19]: a Datalog program consists of a set of declarative *rules*. Each rule has the form $p :- q_1, q_2, \dots, q_n$, which can be read informally as “ q_1 and q_2 and \dots and q_n implies p ”. Here, p is the *head* of the rule, and q_1, q_2, \dots, q_n is a list

of *literals* that constitutes the *body* of the rule. Literals are either *predicates* with *attributes* (which are bound to variables or constants by the query), or boolean expressions that involve function symbols (including arithmetic) applied to attributes. Predicates in datalog are typically relations, although in some cases they may represent functions.

Datalog rules can refer to one another in a cyclic fashion to express recursion. The order in which the rules are presented in a program is semantically immaterial; likewise, the order predicates appear in a rule is not semantically meaningful. Commas are interpreted as logical conjunctions (*AND*). The names of predicates, function symbols, and constants begin with a lowercase letter, while variable names begin with an uppercase letter. We illustrate *NDlog* using a simple example of two rules that computes all pairs of reachable nodes:

```
r1 reachable(@S,D) :- link(@S,D).
r2 reachable(@S,D) :- link(@S,Z),
                       reachable(@Z,D).
```

The rules `r1` and `r2` specify a distributed transitive closure computation, where rule `r1` computes all pairs of nodes reachable within a single hop from all input links (denoted by the `link`), and rule `r2` expresses that “if there is a link from `S` to `Z`, and `Z` can reach `D`, then `S` can reach `D`.” By modifying this simple example, we can construct more complex routing protocols, such as the distance vector and path vector routing protocols.

NDlog supports a *location specifier* in each predicate, expressed with `@` symbol followed by an attribute. This attribute is used to denote the source location of each corresponding tuple. For example, all `reachable` and `link` tuples are stored based on the `@S` address field. The output of interest is the set of all `reachable(@S,D)` tuples, representing reachable pairs of nodes.

When executed, the above *NDlog* query is essentially a distributed stream computation, where stream of `link` and `reachable` tuples are joined at different nodes to compute routing tables. In a recent work [2], we show that *sliding windows* commonly used in stream processing can be used to process *soft-state* [20] data in declarative networks, where the time-based window size essentially corresponds to the soft-state lifetime of all routes.

2.2 Secure Network Datalog

Secure Network Datalog (*SeNDlog*) [1] is a unified declarative language for networks and security policies, which combines language features from *NDlog*, and *Binder*, a logic-based language for access control in distributed systems. *SeNDlog* utilizes *Binder*’s notion of *context* that represents a component (or security principal) in a distributed environment and a distinguished operator “*says*”. We illustrate *SeNDlog* via the same reachable example as before, with the additional use of the “*says*” operator:

```
At S:
s1 reachable(S,D) :- link(S,D).
s2 linkD(D,S)@D :- link(S,D).
s3 reachable(Z,Y)@Z :- Z says linkD(S,Z),
                       W says reachable(S,Y).
```

The rules `s1-s3` are within the context of the principal `S`. An additional *localization rewrite* [15] ensures that all rule bodies are localized within a context (i.e. have the same location specifier). Assuming an untrusted network, this allows rules to execute only based on trusted local data, or authenticated data from remote sources. The “*says*” construct is an abstraction for the details of authentication. In one specific implementation, communication happens via signed certificates, where derived tuples signed using the private key of the exporting context can be imported into another context and checked using the corresponding public key. E.g. node `S` will import the `reachable(S,Y)` fact from its neighbor `W`, and verify that it is indeed from `W` via the signature stored with the fact. Node `S` then derives the `reachable(Z,Y)` fact which is signed and exported to node `Z`. Note that the implementation of “*says*” may depend on the system and its context. In a hostile world, “*says*” may require digital signatures, while in a more benign world, “*says*” may simply append a cleartext principal header to a message—and this will of course be cheaper. The policy writer could additionally provide hints along with rules, indicating that some “*says*” are more important than others, e.g. by supporting multiple “*says*” operators with different security levels.

3 Provenance in Practice

In this section, we survey a (non-exhaustive) list of existing work in the networking literature that motivates the use of network provenance. We classify the use-cases as *real-time diagnostics*, *forensics*, *accountability*, and *trust management*.

Real-time Diagnostics: Provenance is useful for real-time diagnostics and debugging [24, 8, 21] of distributed systems. In a declarative monitoring system, one can add additional queries that monitor a network for runtime anomalies, e.g. lack of convergence, network traffic spike suggesting possible intrusion. To illustrate, a continuous query specified in *SeNDlog* can be used to compute the number of changes changes to a routing table entry over past `T` seconds, and generate an alarm event when the number of changes exceed a threshold as an indication of possible divergence. Upon receiving the alarm, the system may generate a distributed recursive query over the network provenance to detecting the source of malicious activities.

Forensics: In addition to real-time data, historical data is often required to correlate traffic patterns of attackers. A common area of research has been in providing “*traceback*”[22] of traffic, either by the receiver or by an involved third party, to determine where packets are

originated from without trusting the unauthenticated IP headers. One can store annotations either in the packet (i.e. piggyback each tuple with its complete “path” or “provenance”), or maintain state at each router, to allow for subsequent traceback via a distributed query during forensic analysis. To reduce the storage and communication overhead, *ForNet* [23] and *Time Machine*[13] have proposed techniques that trade-off accuracy for performance, by using summarization (via bloom filters) and sampling techniques to compress the provenance.

Accountability: Forensics analysis is essentially a form of *call-detail* used in voice telephone networks, where historical information on the caller, callee, length of call, and call status both in real-time and in many cases historically through the examination of call detail records. One important use of the call-detail information is to enforce *accountability*, or proper usage in networks. For example, PlanetFlow [11] is a network auditing service provided on PlanetLab [18], to provide accountability for all traffic generated by PlanetLab services, to ensure that all users are in accordance with PlanetLab policies.

Trust Management: In our final use case, network provenance is useful for enforcing distributed trust management [5] policies in networked information systems. Using an example from Internet routing, the path-vector protocol used in BGP carries the entire path during route advertisement, in order to allow for ASes to enforce their respective policies. More generally, provenance in our system enables any networked information node to trace the origins of its data, and hence enforce trust policies to accept or reject incoming updates based on the source origins. The Orchestra [12] p2p data-integration engine uses provenance in this manner, to accept or reject updates from neighboring nodes by examining the provenance of updates and the trust relationships among nodes. Taking this idea one step further, one can maintain a *quantifiable* notion of trust, e.g. accepting an update *only if* over K principals assert the update.

4 Taxonomy of Data Provenance

In this section, we present a taxonomy of data provenance, most of which are derived from existing database literature. We then map that to the use cases presented in Section 3. To illustrate, we make use of an example network which consists of three nodes a, b, c and three unidirectional links $\text{link}(a, b), \text{link}(a, c), \text{link}(b, c)$.

We give the derivation tree for $\text{reachable}(@a, c)$ in Figure 1 as a result of executing the *NDlog* query in Section 2.1. This derivation tree essentially represents the lineage or provenance of the tuple, and one can use this tree to figure out the initial input base tuples (at the leaves of the tree). The ovals in the diagram represent the rules ($r1, r2$, or union) to combine their results that are used for the derivation of $\text{reachable}(@a, c)$.

Declarative networks are essentially computations over distributed streams, with time-based sliding win-

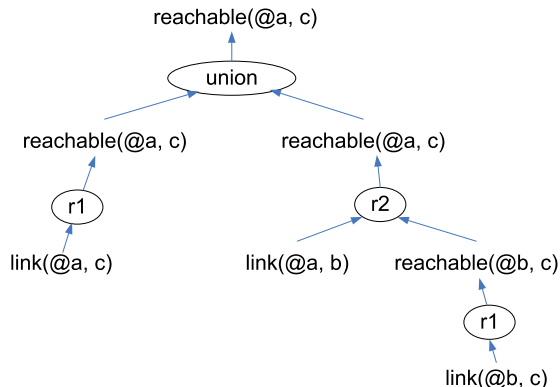


Figure 1: *NDlog* derivation tree for $\text{reachable}(a,c)$.

dows for soft-state derived tuples. In order to incorporate provenance into distributed streams, we make the following changes to traditional provenance. First, we annotate each derivation with its location (denoted by the location specifier “@”). Second, since tuples are soft-state with lifetimes, we also add creation timestamps and time-to-live to the nodes in the tree.

4.1 Local vs Distributed Provenance

The derivation tree shown in Figure 1 can be stored either locally or in a distributed fashion. In *local provenance*, the tree is stored at node a , which is the final storage location of $\text{reachable}(@a, b)$. In order to have a locally complete provenance, each tuple that is derived needs to piggy-back its entire provenance when shipped from one node to another.

On the other hand, one can utilize *distributed provenance*, which only stores pointers to the previous node to reconstruct its provenance on demand. Hence, node a only needs to store the fact that it is derived from $\text{link}(@a, b)$ which is available locally, and $\text{reachable}(@b, c)$ which is stored at node b . The analogy here is IP traceback, where one can either store the entire traversed path within each packet (similar to local provenance), or only maintain enough state at each router to traceback the route on demand.

There are evidently tradeoffs between local and distributed provenance. In local provenance, computation is more expensive for each tuple, but provenance querying is cheap. Also, since each node has the provenance available locally, it can also better enforce trust policies. On the other hand, distributed provenance requires no extra communication overhead, but incurs expensive cost of querying the provenance.

4.2 Online vs Offline Provenance

Along another axis, we can further classify provenance as either *online* or *offline*. Online provenance is maintained for network state that is currently valid (i.e. not expired), and *offline* provenance is kept even when the derivations have expired. The purpose of online prove-

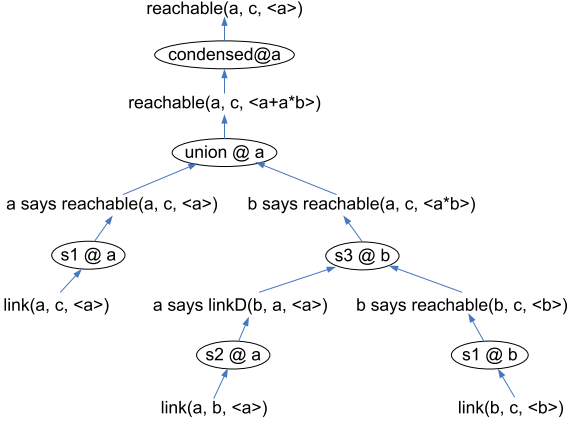


Figure 2: *SeNDlog* derivation tree for $\text{reachable}(a,c)$ with annotations for condensed provenance.

nance is for runtime reaction to network anomalies. For example, when a node is detected to be suspicious, one can query the online provenance to delete all routing entries associated with the malicious node. However, online provenance by itself has limited usage given that most networked data are maintained as soft-state with TTLs. In this case, we can additionally maintain offline provenance for data that has long expired. Offline provenance is also useful for real-time diagnostics, and can additionally be used to support forensics and enforce accountability. Offline provenance can result in high storage overhead. We will revisit this issue in Section 5.

4.3 Authenticated Provenance

Up to this point, we have assumed that all nodes who compute the provenance are trusted. In practice, authentication is required to ensure the validity of provenance computed by other nodes (e.g. to prevent spoofing of messages from malicious attackers).

Figure 2 shows an alternative derivation tree based on the *SeNDlog* query presented in Section 2.2. We note the following differences. First, since all rule bodies are localized within the context of a security principal, we can omit the location specifiers for each tuple. However, we annotate each operator (denoted by the oval) with the location (or context) where the rule is executed. Second, each node in the tree is asserted by a principal using “says”. In an untrusted environment, this means that individual nodes in the provenance tree need to have digital signatures to validate the authenticity of the computed provenance.

4.4 Condensed Provenance

When computing local provenance, the overhead of shipping the entire provenance with each tuple may be expensive. With authenticated provenance, the overhead is increased due to the digital signatures. We note that in several instances, local provenance is desired (e.g. for deciding whether to accept a tuple based on its origins).

To reduce the overhead of computing and sending provenance, we present an existing technique to *condense* the size of local provenance, yet retain sufficient information for enforcing trust based on source origins. This technique is inspired by *provenance semirings* [9] in Orchestra [12] system, where tuples are annotated with provenance expressions that are based on the unique keys of base input tuples. These provenance expressions can themselves be encoded in boolean expressions stored in *Binary Decision Diagrams* (BDD) [6], and further compressed as presented in [2].

To provide the intuition behind the condensation process, we revisit the derivation tree in Figure 2. Each tuple has an additional field denoted by $\langle \dots \rangle$ that stores the condensed provenance, where $+$ represents union, and $*$ represents a join operation. An expression such as $\langle a+a*b \rangle$ for $\text{reachable}(a, c)$ can be compressed simply into $\langle a \rangle$. Intuitively, whether the principal b is trusted or not is inconsequential given a . As long as principal a is trusted by the node that receives the $\text{reachable}(a, c)$ tuple, this tuple will be accepted, regardless of whether principal b is trusted or not.

4.5 Quantifiable Provenance

The provenance semirings formulation of provenance also permits quantifiable notions of trust that can leverage the different levels of “says” described in Section 2.2. For example, consider the non-condensed expression $\langle a+a*b \rangle$, where principal a has security level 2, and b has security level 1. One can conclude that $\text{reachable}(a, c)$ derivation has a trust level of $\max(2, \min(2, 1))=2$, assuming that the higher security level is more trusted. Other quantifiable notions of trust are also possible, e.g. the *count* [10] of the number of ways each derivation is achievable, or *vote*, representing the number of principals that agree on a derivation concurrently.

4.6 Summary

We summarize the types of provenance that are applicable to each usage scenario. In real-time diagnostics, *online* provenance of existing data is required. The provenance can be *local* or *distributed*, and can further be authenticated. On the other hand, forensics and accountability require *offline* provenance, and in practice, would be used in conjunction with *online* provenance. Trust management is best enforced locally at each node, and one can reduce communication overhead by using *condensed* provenance to store only the source principals necessary to enforce trust, or *quantifiable* provenance if trust is based on security levels

5 Optimizations

A key challenge in maintaining network provenance is in lowering the storage, communication, and distributed querying overheads. In the previous section, we have

seen how condensed provenance encoded via BDDs can result in a compact representation of provenance that can be evaluated locally for trust management. In addition, we outline three possible optimizations that we would like to further explore as future work:

Proactive vs reactive provenance: In *proactive provenance*, all the provenance of new tuples are eagerly maintained and propagated throughout the network. In a more *reactive* mode of operation, one can maintain *lazy provenance*, whose computation is triggered only by specified network events. For example, in the earlier path computation example, we can start computing the provenance of `nextHop` only when route divergence is detected. Similarly, offline provenance for forensics can be aged out over time to reduce storage, unless explicitly marked to persist as a result of network anomaly.

Sampling: A straightforward optimization is to only record a portion of the provenance (both online and offline) via sampling techniques. For example, IP Traceback [22] (which generates a new message 1/20,000th of the time) and *ForNet* [23] (which uses Bloom filters) are examples of this approach. The sampling techniques can also be applied when querying distributed provenance. One example existing technique is the use of random moonwalks [26] to avoid querying all provenance.

Provenance granularity: In reconstructing network provenance, there are different granularities at which systems can operate. To reduce overhead, provenance can be aggregated and maintained at the AS granularity. While it may not be conducive to detect all attacks, AS granularity is likely sufficient for detecting aggregated events such as a large number of spoofed packet injections from a group of malicious nodes within the AS.

6 Preliminary Evaluation

In this section, we present a preliminary evaluation study on the overhead of authenticated communication and computation of network provenance. We modified the P2 declarative networking system [16] to support the *SeNDlog* query language, which is compiled into distributed dataflows that exchange messages that are signed with RSA signatures. We further modify various relational operators (particularly joins) in the P2 system to support provenance. In particular, we focus on evaluating the performance of *authenticated provenance* (Section 4.3) which is individually signed by the principal that asserted each fact, and we further apply the condensation (Section 4.4) to reduce communication and storage overhead.

We utilize the OpenSSL v0.9.8b, and Buddy BDD v2.4 libraries to support encryption and provenance. Our experiments are performed on a quad-core machine with Intel Xeon 2.33GHz CPUs and 4GB RAM running Fedora Core 6 with kernel version 2.6.20. In our experiments, we execute up to 100 P2 processes representing different nodes on the machine.

For the query workload, we utilize the *Best-Path* recursive query that computes the shortest paths between all pairs of nodes. This query is obtained from the *NDlog* all-pairs reachability query presented in Section 2, with additional predicates to compute the actual path, cost of the path, and two extra rules for computing the best paths. As input, we insert link tables for N nodes with average outdegree of three, and vary the size of N from 10 to 100. To isolate the individual overhead of authenticated communication and provenance, we execute three versions of the *Best-Path* query: *NDlog* version without authentication and provenance, *SeNDLog* with authentication but without provenance, and *SeNDLogProv* with both authentication and provenance. Our metrics of evaluation are as follows:

Query completion time (s): Time taken for a query to finish execution. As our example programs are recursive, this means the time elapsed before the system reaches a distributed fixpoint, where all nodes finish computing their best paths.

Bandwidth usage (MB): The total combined bandwidth usage across all nodes required for executing the distributed query.

In our experiments, we measure the computation and bandwidth overheads of encryption and provenance by comparing *NDLog*, *SeNDLog* and *SeNDLogProv*. Figure 3 and 4 shows the query completion time and bandwidth utilization respectively, averaged over 10 experimental runs. We summarize our results as follows:

SeNDlog overhead: The use of authenticated communication in *SeNDLog* incurs in the average 53% delay in query completion time and additional 36% bandwidth utilization compared to *NDlog*. As N increases, the additional overhead decreases. For example, when N is 100, the overhead is 44% and 17% respectively. Given that we are running multiple P2 processes on a single node and generating a signature for each tuple, this represents an upper bound on the encryption overhead.

Condensed provenance overhead: The query completion time of *SeNDLogProv* increases by 41% compared to *SeNDLog* due to the overhead of computing and shipping provenance. In addition, *SeNDLogProv* requires 54% more bandwidth than *SeNDLog*. Similar to the *SeNDlog* overhead above, we observe that provenance overhead decreases as the number of nodes increases. For example, when N is 100, *SeNDLogProv* only incurs additional 6% and 10% costs in computation and bandwidth overhead respectively. Our results demonstrate that the BDD-encoded condensed provenance is efficient for recording derivation of tuples, at reasonably low overhead especially for larger networks.

7 Conclusion

In this paper, we argue that network accountability and forensic analysis can be posed as data provenance com-

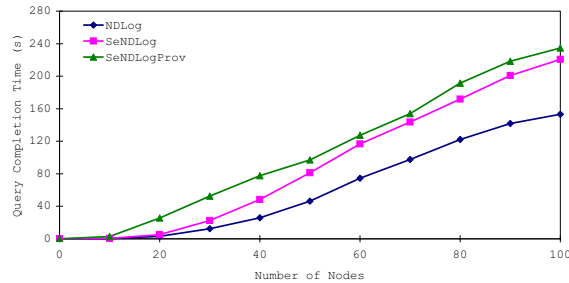


Figure 3: Query completion time (s) for *Best-Path* query

putations and queries over distributed streams. In particular, one can utilize provenance-aware secure networks with appropriate security extensions to provide a flexible declarative framework for specifying, analyzing and auditing networks. To prove our case, we propose a taxonomy of data provenance along multiple axes, and show that they map naturally to several use cases ranging from network forensics and diagnostics to trust management. We suggest techniques to efficiently compute and store network provenance, and provide an initial performance evaluation using the P2 declarative networking system.

Our future work is proceeding along several fronts. First, while we focus on forensics and accountability over the Internet, we intend to explore the general applicability of these techniques to overlay networks and sensor networks. Second, we are in the process of evaluating a variety of secure networks specified and implemented by using *SeNDlog* (e.g. secure Chord routing, DNSSEC), and studying the usage of network provenance for a variety of networks. This will enable us to investigate cross-layer analysis opportunities that arise as a result of having a single integrated system that unifies network and security specifications.

References

- [1] M. Abadi and B. T. Loo. Towards a Language and System for Secure Networking. In *NetDB*, 2007.
- [2] Anonymous. Paper is under submission.
- [3] K. Argyraki, P. Maniatis, D. Cheriton, and S. Shenker. Providing packet obituaries. In *Proc. of 2006 ACM SIGCOMM Workshop on Mining Network Data (MineNet '06)*. ACM Press, Sept. 2006.
- [4] A. Bender, N. Spring, D. Levin, and B. Bhattacharjee. Accountability as a service. In *USENIX Steps to Reducing Unwanted Traffic on the Internet*, 2007.
- [5] M. Blaze, J. Feigenbaum, and A. D. Keromytis. The role of trust management in distributed systems security. In *Secure Internet Programming*, pages 185–210, 1999.
- [6] R. E. Bryant. Symbolic boolean manipulation with ordered binary decision diagrams. *ACM Computing Surveys*, 24(3), 1992.
- [7] P. Buneman, S. Khanna, and W. C. Tan. Why and where: A characterization of data provenance. In *ICDT*, 2001.
- [8] D. Geels, G. Altek, P. Maniatis, T. Roscoe, and I. Stoica. Friday: Global Comprehension for Distributed Replay. In *USENIX Symposium on Networked Systems Design and Implementation*, 2007.




Figure 4: Bandwidth utilization (MB) for *Best-Path* query.

- [9] T. J. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *ACM Symposium on Principles of Database Systems*, 2007.
- [10] A. Gupta, I. S. Mumick, and V. S. Subrahmanian. Maintaining Views Incrementally. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, 1993.
- [11] M. Huang, A. Bavier, and L. Peterson. PlanetFlow: Maintaining accountability for network services. *Operating Systems Review*, 40(1):89–94, Jan. 2006.
- [12] Z. Ives, N. Khandelwal, A. Kapur, and M. Cakir. ORCHESTRA: Rapid, collaborative sharing of dynamic data. In *CIDR*, January 2005.
- [13] S. Kornexl, V. Paxson, H. Dreger, A. Feldmann, and R. Sommer. Building a time machine for efficient recording and retrieval of high-volume network traffic. In *Internet Measurement Conference (IMC)*, 2005.
- [14] P. Laskowski and J. Chuang. Network monitors and contracting systems: Competition and innovation. In *Proceedings of ACM SIGCOMM Conference on Data Communication*, 2007.
- [15] B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative Networking: Language, Execution and Optimization. In *ACM SIGMOD*, June 2006.
- [16] B. T. Loo, T. Condie, J. M. Hellerstein, P. Maniatis, T. Roscoe, and I. Stoica. Implementing Declarative Overlays. In *ACM SOSP*, 2005.
- [17] B. T. Loo, J. M. Hellerstein, I. Stoica, and R. Ramakrishnan. Declarative Routing: Extensible Routing with Declarative Queries. In *ACM SIGMOD*, 2005.
- [18] PlanetLab. Global testbed. 2006. <http://www.planet-lab.org/>.
- [19] R. Ramakrishnan and J. D. Ullman. A Survey of Research on Deductive Database Systems. *Journal of Logic Programming*, 23(2):125–149, 1993.
- [20] S. Raman and S. McCanne. A model, analysis, and protocol framework for soft state-based communication. In *SIGCOMM*, pages 15–25, 1999.
- [21] P. Reynolds, C. Killian, J. L. Wiener, J. C. Mogul, M. A. Shah, and A. Vahdat. Pip: Detecting the Unexpected in Distributed Systems. In *USENIX Symposium on Networked Systems Design and Implementation*, 2006.
- [22] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Practical network support for IP traceback. In *Proceedings of ACM SIGCOMM Conference on Data Communication*, 2000.
- [23] K. Shanmugasundaram, N. Memon, A. Savant, and H. Bronnimann. ForNet: A distributed forensics network. In *Proc. of 2nd International Workshop on Mathematical Methods, Models, and Architectures for Computer Network Security*, 2003.
- [24] A. Singh, P. Maniatis, T. Roscoe, and P. Druschel. Distributed Monitoring and Forensics in Overlay Networks. In *Eurosys*, 2006.
- [25] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. In *Proceedings of ACM SIGCOMM Conference on Data Communication*, 2001.
- [26] Y. Xie, V. Sekar, M. K. Reiter, and H. Zhang. Forensic analysis for epidemic attacks in federated networks. In *Proc. of the 2001 IEEE Symposium on Security and Privacy*, pages 43–53. IEEE Computer Society, May 2001.