

# Securing Structured Overlays against Identity Attacks

Krishna P.N. Puttaswamy, Haitao Zheng, and Ben Y. Zhao

**Abstract**—Structured overlay networks can greatly simplify data storage and management for a variety of distributed applications. Despite their attractive features, these overlays remain vulnerable to the Identity attack, where malicious nodes assume control of application components by intercepting and hijacking *key-based routing* requests. Attackers can assume arbitrary application roles such as storage node for a given file, or return falsified contents of an online shopper's shopping cart. In this paper, we define a generalized form of the Identity attack, and propose a lightweight detection and tracking system that protects applications by redirecting traffic away from attackers. We describe how this attack can be amplified by a Sybil or Eclipse attack, and analyze the costs of performing such an attack. Finally, we present measurements of a deployed overlay that show our techniques to be significantly more lightweight than prior techniques, and highly effective at detecting and avoiding both single node and colluding attacks under a variety of conditions.

**Index Terms**—Security, routing protocols, distributed systems, overlay networks.

## 1 INTRODUCTION

As the demand for Internet and Web-based services continues to grow, so does the scale of the computing infrastructure they are deployed on. Recent literature shows that structured peer-to-peer (P2P) overlays such as Chord [30], Pastry [24], and Tapestry [37] can greatly simplify data storage and management for a variety of large-scale distributed applications [19], [26], [39]. Finally, the usefulness of these infrastructures has been validated by recent studies of real-world deployments of structured overlay applications, including Amazon's distributed key-value storage system Dynamo [10], and the popular BitTorrent client Azureus [13].

Despite the success of new application deployments, these application infrastructures remain vulnerable to several critical malicious attacks. One such attack, the Identity attack [15], would allow a malicious peer in the network to hijack application-level requests and assume the responsibility of any application component. For example, Amazon uses Dynamo servers to store information about user shopping carts, which is read and converted to html for user consumption. While Dynamo servers are accessed only by internal servers, a compromised host can potentially hijack read operations on a user's shopping cart, and return to requesting Web servers a modified shopping cart.

At their core, structured overlays scale to large networks because each node stores routing state that scales *sublinearly* with the network size. Using this limited state, peers cooperatively map *keys* to physical network nodes using a

multihop lookup mechanism called *key-based routing* (KBR) [8]. KBR maps a given key to a specific live node called its *root*. Overlay applications then use KBR to "choose" nodes for specific application components, e.g., the root node of a file's content hash is the file's storage server [7], and the root of a multicast session key becomes the root of the multicast tree [26]. By using KBR to choose application components, these applications expose themselves to hijacking attempts by malicious peers, similar in principle to BGP hijacking attacks [2].

In our preliminary work [15], we first described a limited version of the Identity attack and a basic framework for its detection in the Tapestry [37] and Pastry [24] protocols. In this work, we greatly expand our study in several key dimensions. First, we present a generalized form of the Identity attack defense, and demonstrate its applicability to 16 of the most popular structured overlay protocols. Second, we propose a novel tracking mechanism that allows the network to accumulate statistics on prior attacks through self-verifying *evidence*. By routing evidence to current and potential victims, we effectively protect application traffic by redirecting KBR requests away from attackers. Not only does this render attackers harmless, but it also does not unfairly penalize the few nodes who are falsely accused due to network instability. Third, we describe how Identity attacks can be amplified using the Eclipse attack [28], and perform analysis to quantify its required cost in terms of node identifiers requested. Finally, we measure the effectiveness of our solution on a deployed structured overlay network. In addition to experiments on the overlay, we implement a version of the Cooperative File System (CFS) [7], and show how our techniques effectively reduce forged data blocks while requiring order of magnitude lower overheads than an alternative approach using redundant routing entries.

This paper is organized as follows: Section 2 summarizes background and related work. Section 3 summarizes the Identity attack, and describes our comprehensive framework

• The authors are with the Department of Computer Science, University of California at Santa Barbara, Santa Barbara, CA 93106-5110.  
E-mail: {krishnap, htzheng, ravenben}@cs.ucsb.edu.

Manuscript received 18 Feb. 2008; revised 19 Aug. 2008; accepted 14 Oct. 2008; published online 30 Oct. 2008.

Recommended for acceptance by J.C.S. Lui.

For information on obtaining reprints of this article, please send e-mail to: [tpds@computer.org](mailto:tpds@computer.org), and reference IEEECS Log Number TPDS-2008-02-0064. Digital Object Identifier no. 10.1109/TPDS.2008.241.




Fig. 1. *Prefix routing.* Top: Node 1023 sends a message to key 3222 in a structured overlay using prefix routing. Bottom right: Routing table for node 1023. The node itself fills one entry in each column. Entries that no nodes match are crossed out. Names are represented in base 4.

for lightweight detection and countermeasures. Next, in Section 4, we examine Identity attacks enhanced by collusion via the Eclipse attack, and analyze the costs of collusion. We present detailed evaluation of our defense framework in Section 5, and measurements of application-level impact via a study of CFS in Section 6.

## 2 BACKGROUND AND RELATED WORK

We begin with a background discussion of structured overlays and their use of KBR. We then summarize work on security for structured overlays and related topics.

*Structured overlays and KBR.* A structured overlay is an application-level network connecting any number of nodes, each representing an instance of an overlay participant. Nodes are assigned nodeIds uniformly at random from a large identifier space. To enhance overlay security, we assume that nodes register with a centralized certificate authority (CA) for a public/private key pair, and the CA binds the node's random nodeId with its public key using a public key certificate. Application-specific objects are assigned unique identifiers called keys from the same space.

The overlay dynamically maps each key to a unique live node, called its *root node*. While a key's root can change with network membership, at any given time in a consistent network, a single node is responsible for each key. The root is usually defined as the peer with nodeId closest to the key. Our detection mechanisms verify this overlay *invariant* to detect identity attacks. To deliver a message based on a key to its root node (KBR [8]), each node forwards the message using a locally maintained routing table of overlay links. P2P applications use this to deterministically choose peers to perform specific functions.

The large set of existing structured overlay protocols differ in the specifics of their routing algorithms. In order to support a network of size  $N$ , most protocols require per-node routing state that scales as  $O(\log N)$  and provide worst case  $O(\log N)$  overlay hops between any two nodes. Routing proceeds by forwarding the message incrementally closer in the namespace to the desired key. Fig. 1 shows an example of overlay routing in Chimera [1], a structured overlay using prefix routing similar to Tapestry [37], and Pastry [24]. Finally, while each system defines a function that maps keys to nodes, the exact function may vary slightly. For example, keys can be mapped to the live node with the closest nodeId as in Pastry, or the closest nodeId clockwise from the key as in Chord [30].

*Attacks on P2P systems.* Previous work describes two attacks on structured overlays, the Sybil attack [12] and the Eclipse attack [4], [28]. Attackers with significant resources perform a Sybil attack by generating arbitrarily large number of identities in the overlay network. In the Eclipse attack, attackers collude to increase their influence on a target by introducing each other as performance optimizing alternative entries into the victim's routing table. Both attacks focus on using multiple identities to gain influence and control in the overlay routing layer.

Some prior work limits the Eclipse attack by putting additional constraints such as low in-degree count or geographic proximity on how nodes choose their neighbors in the overlay [4], [20], [28]. While these limit attackers from attracting more than their share of normal traffic, they cannot prevent them from harming their portion of the overlay traffic. In contrast, our work seeks to identify and actively eliminate the influence of attackers on all traffic. Condie et al. proposed using periodic routing table resets (induced churn) to limit the impact of Eclipse attacks. Sit and Morris [29] also alluded to a variant of the identity attack in their initial study.

Recent works have also proposed a number of mechanisms to address the Sybil attack, including those based on computational puzzles [23] or social networks [35], [36]. These approaches can help limit the impact of Eclipse attacks, and are complementary to our defense mechanisms.

The structured overlay security work by Castro et al. [4] is highly relevant to this work. The authors propose a number of techniques to secure routing through the use of redundancy. A number of significant differences distinguish our work. First, mechanisms proposed in prior work require proper tuning of system parameters, and come at a significant cost in network bandwidth overhead. More specifically, the constrained routing table approach doubles the amount of neighbor maintenance traffic, and significantly reduces the overlay's ability to perform locality-aware optimizations such as proximity neighbor selection [17]. In contrast, we focus primarily on *identifying* the active attackers through a lightweight detection and evasion system (less than 1,500 lines of C, including our implementation of a simple CA). Our mechanisms do not limit routing optimizations, and add minimal per-node state. We provide a detailed comparison of the two approaches in Section 6.

Reputation systems have been explored as application-level security mechanisms for P2P systems. They can be used to quantify reliability of resources [9], [33] or individual peers [21], and have been applied to P2P network communities [31]. Finally, the PeerReview project [18] proposed secure message logs to improve accountability in distributed systems.

## 3 DEFENDING THE IDENTITY ATTACK

Structured P2P applications use KBR [8] as a way to assign application components such as traffic indirection points, storage servers, or measurement sensors to live nodes in the network. For example, distributed storage systems using the Distributed Hash Table (DHT) interface [7], [19] choose the storage server for each file or block as the root node of the file's content hash key. Thus, Put and Get




Fig. 2. *The Identity attack.* 1023 sends a message toward key 3222. Before the message reaches the root (3223), an attacker intercepts it and responds as the root.

operations are KBR requests that terminate at the storage node. Data streaming applications such as Scribe [26] hash session names to generate keys, and use KBR to choose the root of the multicast tree. While application peers can communicate directly via IP, they must first use KBR requests to locate each other.

An attacker can hijack and claim KBR messages as their own, by exploiting the fact that each node only sees a small subset of the overlay members. We call this an *identity attack*. Any malicious peer on the path of a KBR message can respond to the source node and claim to be the request's destination. Undetected, the attacker claims control over a particular key and its associated application data. Intuitively, this abuse of limited per-node state is similar to prefix-hijacking attacks in BGP routing [2]. Multiple attackers can collude and perform stronger attacks such as isolating a node from the network, effectively performing a manual partitioning of the overlay.

In this section, we quickly summarize the basic identity attack described in our preliminary work [15], then define a generalized lightweight detection framework for structured overlays, and introduce techniques to mark attackers and redirect traffic away from their influence. We will delay discussion of advanced attacks involving multiple identities to Section 4.

### 3.1 The Single-Node Identity Attack

In the single-peer attack [15], attacker *A* intercepts a KBR request from source peer *S* for key *K*, and responds to *S* that it is *K*'s root node by virtue of being closer to *K* than any other peer in the network. Without local knowledge of peers closer to *K* than *A*, *S* then interacts with *A* as required by the application. Thus, *A* has effectively hijacked the overlay connection between *S* and *K*'s root node at setup time. Fig. 2 illustrates the attack.

By claiming to be *K*'s root node, the attacker can intercept application requests and return data of its own choosing. The extent of control granted to the attacker varies across different applications. In distributed storage systems that use the DHT interface [7], [19], the attacker can prevent blocks from being written and provide forged data to peers by hijacking *put* and *get* operations, respectively. For directory services relying on decentralized object location and routing (DOLR) [22], an attacker can prevent the publication of resources, and redirect requests to malicious nodes hosting forged data. In multicast and anycast systems that rely on KBR for peer rendezvous [5], [26], clients can be redirected to join sessions run by malicious hosts.

Note that while applications can verify content integrity using mechanisms such as block checksums, the distribution of these mechanisms often rely on KBR and are themselves vulnerable to these attacks. Detecting the identity attack is difficult without an out-of-band peer-*rendezvous* and communication mechanism.

### 3.2 Detecting Identity Attacks

We now describe a generalized, lightweight detection mechanism suitable for most if not all of the structured overlays in current literature. We begin by stating several reasonable assumptions about the overlay network:

- We assume that all nodes are assigned unique nodeIds by a CA, along with a public/private key pair [12]. The two are embedded in a public key certificate. NodeId assignment can only be done by a centralized CA, which uses a pseudorandom nodeId generation algorithm such as a secure one-way hash like SHA-1. While this limits the impact of Sybil attacks, it is not a requirement for our mechanisms to work well.
- Local clocks at overlay nodes are loosely synchronized using the Network Time Protocol [11].
- Nodes digitally sign all of their responses to KBR messages with their public key.

Nodes detect identity attacks through the generation and timely dissemination of self-verifying "existence proofs." Overlay nodes periodically sign and distribute these proofs on behalf of well-defined regions of the namespace they reside in. For each region, a small number of randomly selected "proof managers" store these proofs and provide them on request.

Existence proofs are digitally signed certificates that prove at some time *t*, at least one live node existed whose nodeId resides inside a particular region of the overlay namespace. A node periodically constructs existence proofs for each namespace region it belongs to, and sends them to the set of *proof managers* for that region. An existence proof includes the signer's nodeId and a timestamp signed by the sender, and is valid for some time period after its issue. The use of the local timestamps prevents malicious nodes from replaying existence proofs and falsely accusing well-behaved nodes.

Prefix-based routing protocols such as Tapestry [37] and Pastry [24] use prefixes of different lengths to identify specific namespace regions. In a namespace where nodes have IDs of *L* digits, we use a prefix of length  $l < L$  to define a *prefix group* corresponding to the set of all nodes whose IDs start with that prefix. Shorter prefixes will have prefix groups with more members. If nodes use *L*-digit IDs, each node potentially belongs to *L* different prefix groups. This corresponds to different sized partitions of the namespace whose members are nodes with IDs in the range.

After receiving a signed reply to a KBR request for key *K*, a node *S* examines its local routing table to find the longest prefix column for which it has all entries filled. This threshold *T* is an approximate measure of the number of nodes in the network. A legitimate response should match the desired key *K* with at least *T* prefix digits. If not, *S* becomes suspicious, and attempts to verify the existence of




Fig. 3. *Identity attack detection.* Nodes 3220 and 3223 provide signed existence proofs for prefix 322 to three proof managers. Node 1023 seeks proof of a 322\* node.

nodes matching longer prefixes of key  $K$ .  $S$  computes addresses of proof managers by applying a secure one-way hash (SHA-1) to the region identifier and several small natural numbers (1, 2, 3). Proof managers are the “root nodes” of the resulting keys.  $S$  queries the relevant proof managers matching  $K$ 's first  $T$  prefix digits for any existence proofs. If successful, the proof provides indisputable evidence of an attempted identity attack. Fig. 3 shows node 1023 seeking proofs of nodes with prefix 322 by contacting three proof managers.

*Optimizing detection for scalability and robustness.* We summarize several techniques that improve the scalability and robustness of the detection system. We refer the reader to [15] for additional details.

*Limiting prefix groups*—Certifying every possible prefix group in the network would result in a large number of existence proofs. However, given a rough estimate of the network size, each node needs to only certify a small constant number of namespace ranges. Each node can estimate the number of active nodes  $n$  by examining the density of its local routing table, and need to certify at most  $\log_2 n$  prefix groups or namespace ranges. Application of prior measurement and analysis results [38] shows that certificate proofs of *three* prefix groups are sufficient to provide coverage of attacked nodes for different network sizes.

*Replicating proof managers*—Several factors can limit the success of proof managers. Node churn limits their availability, attackers between a client and proof manager can drop verification requests, and proof managers themselves can be malicious and deny any knowledge of requested

existence proofs. The latter means a proof manager can collude with an attacker to help it escape detection. We expect all these attack variants in our attack model. We make detection more robust by using multiple proof managers for each namespace region. This “replication” increases the probability that one or more nonmalicious managers will be online despite network churn. Replicating the proof managers also addresses the problem that our verification mechanism is dependent on overlay routing. Multiple independent proof managers increase the number of verification requests that avoid interception by malicious peers, providing more reliable routing without any complex techniques.

The number of replicas should be chosen to ensure the availability of at least one uncompromised manager, while avoiding the unnecessary overhead of sending proofs to too many proof managers. We quantify the benefits of proof manager replication through detailed experiments in Section 5. Our results show that a small replication factor like 3 is sufficient for attacks compromising up to 20 percent of all nodes, and highly replication factors like 6 or 7 prove effective against massive compromise attacks.

*Caching existence proofs*—An overlay node can cache existence proofs it observes in the network. Since proofs are self-certifying, any node can observe and cache them as they are sent on their way to proof managers. Locally cached unexpired proofs can be referenced as an alternative to sending verification requests.

*Defining namespace regions.* Our defense against identity attacks relies on certifying that nodes exist in different regions of varying sizes. This mechanism generalizes to any structured overlay with a notion of continuous namespaces. Only the mechanism for specifying a region needs to be customized for each protocol. For example, prefix-routing protocols like Tapestry [37] and Pastry [24] define regions as all nodes whose nodeId share a matching prefix, e.g., *prefix*{123}. The length of the prefix defines the region size. In contrast, protocols like Chord [30] that use pure numerical closeness for routing can define the same region using a center identifier and a region size, e.g., *range*{1235,9}. To illustrate the generality of our mechanism, we show in Table 1 how it applies to 16 protocols described in literature. We describe how each protocol

TABLE 1  
The Namespace Region Certification Technique Is Applicable to Most Known Structured Overlays

Overlay Protocol(s)	How Regions are Named	Region Naming Example
Tapestry [37], Kademlia [IPTPS'02], Pastry [24], Bamboo [Usenix'04], LAND [SODA'04], Z-Ring [ICNP'05]	Numeric range (Prefix-based Routing)	[123*] (prefix length $L = 3$ )
Chord [30], Symphony [Usenix'03], Viceroy [PODC'02], Accordion [NSDI'05]	Numeric range	[1235, 9] (center, size)
Koorde [IPTPS'03]	Numeric (de Bruijn Routing)	[1235, 9] (center, size)
SkipNet [USITS'03]	2 Ranges: numeric & alphabet	[abc*] and [123*]
CAN [SIGCOMM'01]	$D$ -dimensional: Numeric	[1235, 9], [5675, 9] ...
Ulysses [ICNP'03]	Numeric range and Level	[1235, 9], $L$
Kelips [IPTPS'03], Tulip [IPTPS'05]	Gossip-based: no ranges	Name of affinity group

This shows how different sized regions are specified in 16 popular protocols. Gossip protocols can only certify regions of a fixed size.




Fig. 4. *Tracking the attacker.* Node 1023 detects an attack by node 3203 and sends an alert toward the attacker. 3203's predecessor 3011 receives the alert, increments 3203's local blacklist counter, and switches its traffic to an alternate route.

defines ranges in its namespace, and give an example to show how to define regions of varying sizes. Note that for protocols like Tulip and Kelips, namespace regions can only be divided into fixed affinity groups.

### 3.3 Tracking and Avoiding Attackers

Merely detecting an attack is insufficient. Nodes must reliably perform KBR despite the presence of malicious peers. In this section, we describe novel mechanisms that use self-certifying evidence of the attack to track down and mark attackers, allowing overlay peers to locate and avoid attacker nodes in favor of more reliable alternative routes.

*Self-verifying evidence of an attack.* For overlay nodes to take action after an attack by  $A$  on key  $K$ , nodes must observe unforgeable evidence of  $A$ 's malicious behavior. This evidence comes in the form of two components, the original reply to the KBR message signed by  $A$ , and an existence proof signed by a node who is a better root node for  $K$ . Since both the reply and proof contain signed timestamps, a third party can observe whether the KBR reply was sent while the existence proof was still valid. Assuming peers are loosely synchronized via the NTP protocol, peer timestamps should be synchronized within 200 ms [11]. Since we expect existence proofs to be valid on the order of minutes (30 seconds to 2 minutes), errors in time synchronization should not prevent validation of the evidence from an attack.

*Tracking attackers via blacklists.* After detecting an attack, the message source  $S$  can forward evidence of the attack to interested third parties in the network. If node  $A$  attacks a KBR message for key  $K$ ,  $S$  prepares an *alert* message including  $K$ ,  $A$ 's signed reply, and the existence proof from a proof manager, and sends the alert as a special message to attacker  $A$ . Each node forwards it toward  $A$ , and checks to see if its next-hop is node  $A$ . If so, the node is a predecessor of  $A$ , verifies the evidence is valid, and adds  $A$  to a local *blacklist*. Each node on the blacklist has an associated counter that is incremented each time a new alert is presented showing that node performed an attack. Fig. 4 shows this tracking mechanism in detail.

Blacklists are not foolproof. First, nodes in a dynamic network can observe a small number of false positives in attack detection due to network churn. Over time, nodes blacklist values slowly decay using an exponentially weighted moving average, thus preventing long-term accumulation of false positives. Second, multiple attackers can collude where one attacker forwards traffic to other

attackers, and acts as a shield by dropping alert messages. We are actively investigating an effective approach to identify and avoid these collusion groups.

*Evading attackers via malice-aware routing.* Once attackers have been identified with blacklists, nodes can actively avoid them when routing KBR requests. The blacklist counter for a routing entry acts as a simplistic reputation value that indicates a likelihood of malicious behavior. Routing policies can consider blacklist values in conjunction with a node's link quality and network latency when choosing between multiple routes. This proactive avoidance approach to routing increases application reliability without unfairly punishing nodes involved in false positives. Finally, malice-aware routing policies allow nodes to dissociate from attackers over time, thus reducing attackers' network in-degree and future attacks.

When considering blacklist values in conjunction with other route selection metrics, routing policies must place a weight on their consideration of blacklist values. Heavier weights will result in higher sensitivity to blacklists, possibly resulting in the selection of lower performance routes. Weights set too low can result in low sensitivity, possibly selecting attacker routes based on their improved performance. The tuning of this parameter can be application-specific, but a simple policy of equal weights for performance and blacklist sensitivity would likely achieve the desired balance.

*Summary of attack models.* We have described a number of different attack mechanisms in the previous sections. We summarize these attacks into two different attack models. The attack models are as follows:

- *Verification denials: Type 1.* We first consider an attack model where malicious peers perform the Identity attack by hijacking all KBR messages including application messages (e.g., *put/get* for a DHT). Proof managers can also be malicious. Malicious proof managers will deny access to any existence proofs it stores, thereby helping other attackers escape detection. Note that we assume that malicious nodes route other messages normally, including verification protocol messages (e.g., *existence proofs, verification request/responses*). We refer to this model as Attack Type 1.
- *Denials and existence proof drops: Type 2.* We also consider a more severe attack model. In addition to all attacks present in Attack Type 1 (Identity attack and malicious proof managers), malicious peers also collude by dropping existence proofs being routed en route to their respective proof managers. We refer to this model as Attack Type 2.

Since these are the possible attack variants based on the Identity attack, we will use them to evaluate our detection mechanisms in Section 5.2. Later, we will also evaluate our detection mechanisms against a combination of Identity and Eclipse attacks in Section 6.

## 4 COLLUSION-ENHANCED IDENTITY ATTACKS

P2P systems that provide zero-cost identities suffer from the Sybil attack [12], where a single user can obtain large




Fig. 5. *The Eclipse attack.* Multiple peers collude against target node  $T$  to infiltrate its routing table. Dark nodes denote colluders. Entries where  $T$  fills its own routing table are marked with a cross.

numbers of virtual identities. These virtual nodes can collude to infiltrate the routing table of a single target and perform a collusion-enhanced Identity attack (or the Eclipse attack). In this section, we describe this attack in more detail, analyze its cost in virtual identities, and outline defense mechanisms.

*Eclipse attacks and user collusion.* In the Eclipse attack [4], [6], [28], malicious peers exploit overlay optimization algorithms such as Proximity Neighbor Selection [17] to introduce colluding peers into a target's routing table. An example is shown in Fig. 5. If successful, a group of peers control the majority of outgoing traffic from the target, and can then be used as a platform to launch powerful attacks. First, colluding peers can perform a colluding Identity attack by hijacking all KBR requests from a peer, effectively *isolating* it from the network, surrounding it with a self-consistent virtual network where all exchanges are controlled by a colluder. Second, colluding peers can drop all outgoing verification requests, making it extremely difficult to detect ongoing attacks.

*Estimating the cost of Eclipse attacks.* To perform a successful Eclipse attack, an attacker must have control over specific nodeIds that satisfy routing constraints in the target node's routing table. Because nodeIds are assigned at random, an attacker must obtain a large number of nodeIds in order to obtain the requisite nodeIds that infiltrate a particular target. Thus, we quantify the "cost" of launching an Eclipse attack as *the average number of nodeIds must an attacker request to attack a target T*. Since the network is likely to be sparse, a node's routing table will contain valid entries (other than the local node) only for the first few levels. Therefore, the cost depends on what levels of the routing table would the attacker have full control over.

For our analysis, the relevant parameters are  $b$ , the base of the nodeId,  $N$ , the size of the namespace,  $L = \log_b N$ , the total number of levels in the routing table, and  $m$  ( $1 \leq m \leq L$ ), the number of levels in the routing table that the attacker would have full control over.

Mathematically, we can reduce this problem into a variant of *The General Coupon Collector's Problem* [34]. In each level of the routing table, the local node fills one entry, leaving  $b - 1$  entries to be filled (recall from Fig. 1). Each randomly obtained nodeId will have probability of  $b^{-l}$  to fill one of the  $b - 1$  entries at level  $l$  ( $1 \leq l \leq L$ ). Hence, the attacker needs to obtain nodeIds to fill  $m(b - 1)$  routing entries ( $m$  levels and  $b - 1$  entries per level). Finally, we introduce an artificial "catch-all" entry that represents all the remaining routing entries that the attacker does not




Fig. 6. *Mapping per-column-filling to the General Coupon Collector's problem.*

need to have control over (shown in Fig. 6). The probability of any nodeId falls into this entry is  $1 - \sum_{l=1}^m (b - 1)b^{-l}$ .

By mapping each of the  $m(b - 1)$  routing entries and the catch-all entry into an object, we form the corresponding General Coupon Collector's Problem as: *Let  $K = m(b - 1) + 1$  objects to be picked repeatedly and randomly; with probability  $p_i$  that the  $i$ th object is picked on a given try, and  $\sum_{i=1}^K p_i = 1$ ; the expected number of tries  $T_E$  required after which all  $K$  objects have been picked at least one is [32]*

$$T_E = \sum_{1 \leq i_1 \leq K} \frac{1}{p_{i_1}} - \sum_{1 \leq i_1 < i_2 \leq K} \frac{1}{p_{i_1} + p_{i_2}} + \dots + (-1)^{K-1} \frac{1}{p_1 + p_2 + \dots + p_K}. \quad (1)$$

Since the probability of filling the catch-all entry is significantly higher than that of the  $m(b - 1)$  routing entries,  $T_E$  can accurately represent the cost of Eclipse attack, i.e., the number of nodeIds randomly acquired in order to fill the first  $m$  levels of the routing table.

We apply our analytical result to prefix-routing networks using digit bases of  $b = 4, 8,$  and  $16$ . Fig. 7 shows the cost of filling the first  $m$  levels of the routing tables. Note that the number of nodeIds required increases exponentially with each additional level. Using base 16, the cost of filling the first two levels of the routing table is only 100-1,000 nodeIds, but jumps sharply to 100,000 at the 4th level and 1,000,000 at the 5th level. The results clearly show that the cost of the attack increases exponentially with each level, and quickly becomes prohibitively expensive past a few levels. We now examine how to leverage this result to evade the Eclipse attack.

*Evading the Eclipse attack.* Our goal is to limit the impact, and to allow successful detection of Eclipse attacks. We adopt two complementary approaches in Chimera. First, nodes periodically induce artificial "churn," resulting in periodic refresh of routing table state [6]. This limits the overall proportion of malicious peers in a node's routing




Fig. 7. *Expected number of randomly obtained nodeIds required to fill  $L$  levels of the target's routing table.*

table. Second, nodes use one-hop indirection for verification requests, forwarding them through one or more random members of its leafset peers. These leafset entries represent entries at the highest levels of the routing table. Therefore, they share long prefixes with the target and are difficult to compromise using randomly assigned nodeIds. In addition, leafset nodes are likely to be geographically distributed across the physical network, making them more resistant to Eclipse attacks exploiting physical proximity optimizations. They serve a similar function as *constrained routing entries* in secure Pastry [4]. We evaluate the effectiveness of these mechanisms in Section 6.

## 5 SYSTEM EVALUATION

We use detailed measurements of a deployed prototype to evaluate the effectiveness of our mechanisms for detecting and avoiding single-node and collusion-enhanced Identity attacks. We examine the effectiveness of attack detection under multiple attack environments and churn models, and also quantify the overhead of our mechanisms.

### 5.1 Measurement Methodology

We implemented our protection framework on Chimera, a lightweight structured P2P overlay implemented as a C library [1]. Since Chimera uses prefix routing similar to Tapestry [37] and Pastry [24], we implemented a version of our detection framework customized for prefix routing. We deployed a network of 1,500 Chimera peers on a 32 machine server cluster connected via switched Gigabit Ethernet. NodeIds are 160 bits long, generated from SHA-1 hashes of public keys, and represented as 40 hexadecimal digits. Unless otherwise specified, each prefix group stores existence proofs at three proof managers, nodes sign and distribute proofs every 15 seconds, and each proof expires after 30 seconds.

We implemented a simple centralized Certification Authority (CA) that runs at a well-known address. Each new node entering the network obtains a public-private key pair and certificate from the CA. We assume that peers communicate with the CA through a secure channel, and the CA can verify each peer's identity.

Each experiment includes results from at least three runs. For each experimental run, we choose a random subset of the network nodes to be attackers. All attackers exhibit the same behavior and remain malicious throughout the experiment. Malicious source peers still forward their KBR requests correctly. Malicious destination peers also behave correctly, since they do not need to hijack traffic destined for them. We use Chimera's built-in reliable transmission mechanisms to ensure no messages are lost and all KBR messages are responded to by its destination or a hijacker.

Note that the percentage of messages hijacked can be significantly higher than the percentage of malicious nodes in the network. A message is attacked if any of the nodes in its path are malicious. If  $p$  is the fraction of attackers in the network, the probability a message is attacked on a path with  $h$  hops is  $1 - (1 - p)^h$ . Our experiments confirm this expectation. Finally, in the absence of prior studies on level of malice in P2P networks, we assume a default 20 percent rate of malicious peers unless otherwise specified.




Fig. 8. Attack detection with 20 percent malicious nodes, three proof managers.

### 5.2 Measurement Results

In our tests, each peer in the network sends a predefined number of KBR messages to random destinations. Malicious peers attack messages they receive according to the experiment's attack model. After all messages are sent, replies received and processed by our framework, we use detailed per-message logs to compute the total number of attacks and detections.

#### 5.2.1 Basic Detection and Analysis

Our first experiment measures the effectiveness of our detection mechanisms against Identity attack with Verification Denials (Type 1), assuming 20 percent of all network peers are malicious attackers. The results plotted in Fig. 8 show that basic detection rate is very high (average of 95 percent or higher). Through careful analysis of the logs, we can attribute all detection failures to one of three scenarios. While we discuss our observations on Chimera, similar scenarios exist for other protocols such as Chord.

The few detection failures we observe can be attributed to three scenarios. First, because our certification mechanism has fixed namespace regions centered around a fixed point, keys near the edge of each region are more vulnerable. For example, key 5999's root node might be 6001, but an attacker at 5900 matches the namespace region for  $prefix\{5\}$  despite being further away in the namespace. We refer to this as the *Edge of Region Effect*. Second, since existence proofs each cover a namespace region, their precision is limited to the granularity of the smallest region. If the real root and the attacker reside in the same region, they cannot be distinguished from each other. We refer to this as the *Limits of Precision*. Finally, since existence proofs are self-certifying, a valid response from a single proof manager is sufficient to detect an attack. However, attackers can compromise all proof managers and complete deny access to the proofs (*Unavailable Proofs*). This accounts for  $\approx 1$  percent of all attacks going undetected. This is exactly as we expected, given that we use three proof managers per group in a network with 20 percent attackers:  $0.2^3 \approx 1$  percent. Our experiments show that these scenarios are responsible for all undetected attacks under attack model 1. Fig. 8 shows the undetected attacked accounted for by the Edge of Region effect.

#### 5.2.2 Attacks and Proof Manager Replication

Fig. 8 shows the detection rate for Attack Type 2, which allows all malicious nodes to collude by dropping all existence proofs en route to their proof managers. There is a visible drop in detection rate compared to Type 1, since the




Fig. 9. Effect of proof manager replication on detection robustness in a 1,000-node network.

chance of an existence proof routing through and being dropped by a malicious node is significant. Detection rates remain high (90 percent) because of the redundancy provided by multiple nodes providing existence proofs for each shared namespace region. Since this is the most powerful single node attack model, we use this as the default attack model for the remainder of our single-node attack measurements.

*Proof manager replication.* To understand the impact of replicating proof managers, we vary the proof manager “replication factor” (RF) and plot the detection rate with increasing malice in Fig. 9. Results show that increasing the replication factor from 2 to 4 provides the greatest gain, with diminishing returns at RF 6 and 8. They also show that that higher replication (RF 6 or 8) can provide robust detection rates (70 percent) even when a large majority (70 percent) of the network is malicious.

### 5.2.3 Detection under Network Churn

Realistic evaluation of overlay networks must include tests under network dynamics (network churn). We evaluate our implementation using a range of artificial churn models that follow exponential peer lifetime distributions. This exponential distribution produces even more extreme network churn compared to recent churn measurements on Skype [16] and Gnutella [27]. Our experiments maintain a near-constant network size by introducing new peers into the network as others leave. Our experiments using those churn models produced very high detection rates, which we omit here for brevity.

We observe that the main impact of churn is messages lost en route or being processed by a node leaving the network. By default, messages in Chimera are sent via UDP and acknowledged per overlay hop. If no ack is received, a message is retransmitted up to three times. Chimera also measures overlay links for loss, and performs route switching based on loss rates. Fig. 10 shows the same experiment (1,000 nodes,




Fig. 10. Detection with and without one-hop acknowledgments, 1,000 nodes, 20 percent malice.




Fig. 11. Attack detection under churn.

20 percent malicious nodes, Attack Type 2) with and without one-hop retransmissions. Clearly, local retransmissions made transmission of proofs and requests reliable, thus dramatically improving detection under churn, e.g., from 26 percent to 80 percent for networks with node life time distributions around 300 seconds.

Fig. 11 shows that detection rates drop significantly for high churn rates (lifetime = 300 seconds). With higher churn, new proof managers might not have received existence proofs and those with proofs might have left the network, making proofs unavailable and lowering detection rates. We also plot the impact of increasing number of proof managers under churn in Fig. 12. The results are consistent with our nonchurn experiment, and show less incremental improvement for more than four managers.

### 5.2.4 False Positives

Churn can produce temporary inconsistencies in peer’s routing state. In some cases, peers might be unaware of better root nodes for a given KBR request, resulting in an observed Identity attack. Note that our tracking and avoidance mechanisms do not adversely impact falsely accused nodes. While very few instances are actually observed in our experiments, false positives for our system can occur under two rare scenarios:

1. *Unexpired existence proofs.* A node sends an existence proof to its proof managers just before leaving. Its proofs are valid for some time after it has left (expiration time of 30 seconds in our tests). Requests for keys controlled by the departed node can cause false positives. Reducing the proof expiration period reduces false positives but increases system overhead. Table 2 shows a clear correlation between our false positive rate (out of all detected attacks) and shorter existence proof expiration periods.
2. *New roots.* A node between the KBR key and its root enters the network as a KBR message reaches its root. Existence proofs from the new node arrive




Fig. 12. Impact of proof manager replication under churn, Attack Type 2.



**TABLE 2**  
Percentage of False Positives and Proof Expiration Times

Proof Expiration Time (s)	30	60	90	120
% of False Positives	6.0	11.9	12.8	14.6




Fig. 13. Bandwidth measurement of certification and verification messages. Attack Type 1 on a 1,000-node network.

**TABLE 3**  
Comparison of Certification Overhead with CRT Overhead, Network has 1,000 Nodes

Cert. Overhead (bytes per second per node)	102.86
CRT Overhead (bytes per second per node)	196.11

before any verification requests, leading to a false positive.

**5.2.5 System Bandwidth Overhead**

Fig. 13 summarizes our detailed bandwidth measurement results from a 1,000-node network with Attack Type 1.<sup>1</sup> Certification overhead per second per node remains constant as rate of malicious attackers increases. The low number of verification requests per actual attack means our system is issuing few spurious verifications. The decrease in verification attempts per attack can be attributed to the additional attacks actually decreasing the number of spurious verifications attempted.

We also compare the overhead of our approach to the previously proposed constrained routing tables solution [4]. We implemented constrained routing tables on Chimera and maintained entries in the routing table with the same heartbeat period as Chimera routing entries: 20 seconds. Measurements from this implementation for a 1,000-node Chimera network are shown in Table 3. Intuitively, our approach results in much lower overhead because we incur a *single* per-prefix group proof to be sent to proof managers periodically, which is relatively low compared to existing routing entry liveness probes. In contrast, the constrained routing table approach requires at least one additional periodic probe message for each entry in the routing table, thereby at least doubling the routing entry maintenance traffic. Also note that while certification overhead of our proposal stays constant across different network sizes, the constrained routing table maintenance overhead will increase for larger networks as the number of routing entries increases. We also compare the overheads of these systems in an application scenario in Section 6.

1. We performed bandwidth measurements for all attack types, and found that Type 1 incurs the highest overhead since fewer messages are dropped and more replies are made to verification requests.




Fig. 14. Effectiveness of blacklist counting. Alert messages track down the attacker’s predecessor and increment a blacklist count for the attacker. Attack Type 2, 1,000-node network.




Fig. 15. In-degree changes of the malicious nodes. Attack Type 2, 1,000-node network, KBR messages sent at a rate of 1 msg/second.

**5.2.6 Blacklists and Avoiding Malicious Nodes**

We measure the effectiveness of alerts to track down attackers by their predecessors. We run a Chimera network of 1,000 nodes, and randomly mark 20 percent of nodes as malicious attackers. Each node sends 50 KBR messages to randomly chosen key destinations; malicious nodes attack all messages; and peers send alerts for each detected attack.

In Fig. 14, the X-axis sorts each predecessor-attacker combination by the number of attacks performed across the link, and the Y-axis plots the corresponding blacklist value. We introduce a random variance of 0.15 on Y-values to distinguish multiple points. As the results show, blacklist values are closely correlated to the actual number of attacks, meaning alert messages are locating the attacker’s predecessor. Predecessors whose traffic the attacker hijacks will accrue high blacklist values for the attacker.

To examine the impact of our countermeasures, we implemented a routing policy in Chimera to prioritize attacker-avoidance over other metrics such as reliability or link latency when performing route selection. By default, Chimera maintains three routes per entry in the routing table.

We take snapshots of the “in-degree” of each attacker at 10-second intervals in an experiment that totals 10,000 seconds, where in-degree is the number of predecessors for which the node is the preferred route for some path. We bucketed the nodes based on their in-degree into buckets of size 30 starting from 0, and plot the average in-degree of the nodes in each bucket for every snapshot as a dot in Fig. 15. We see from the figure that as malicious nodes attack, they are detected and blacklisted, causing their in-degree counts to drop quickly. Malicious attackers with a high in-degree at the beginning of the experiment quickly drop as many of their predecessors receive alerts of their attacks. The results confirm that the malice-aware routing policy is highly effective in marking and redirecting traffic away from attackers.




Fig. 16. CFS block lookups comparing our approach against constrained routing tables. Attack Type 2, 1,000-node network.

## 6 COOPERATIVE FILE SYSTEM MEASUREMENTS

Ultimately, our detection and avoidance framework should effectively protect applications from the Identity attack. To quantify the impact of our system on a real application, we implemented a version of the CFS [7] on Chimera, which supports the PUT, GET block operations, and block replication. To measure just the effect on lookup, we ensured that only GET messages are hijacked by malicious nodes. Upon detecting a GET hijack, nodes retry the operation. To contrast our approach against existing work, we implemented the secure constrained routing approach proposed by Castro et al. [4], and compare the two systems side-by-side on effectiveness and overheads.

Since we have already shown the effectiveness of our mechanisms against two variants of the Identity attack, we want to explore in these experiments the impact of augmenting a Type 2 Identity attack with a continuous Eclipse attack. In this model, attackers are continuously requesting new nodeIds and infiltrating the routing table of the target node.

### 6.1 Effectiveness of Malice-Aware Routing

To improve the lookup success of CFS under attack, we implemented the malice-aware routing policy described earlier in CFS. Nodes favor routes with lower blacklist counters, and use backup routes or a leafset entry to evade attackers unless the blacklisted entry is itself the root of the blockId.

We run a 1,000-node Chimera network using a Type 2 Identity attack model. Each node performs a PUT and GET on 25 unique blocks, resulting in read and write operations on a total of 25,000 blocks. Fig. 16 shows that malice-aware routing mechanism significantly improves the lookup success of CFS compared to CFS with normal routing. Results also show that malice-aware routing performs better than constrained routing with redundancy factors of 4 and 8. These values represent the number of parallel outgoing requests following an attack, and are set high to maximize the effectiveness of constrained routing.

The next step is to understand the overheads incurred by the defense mechanisms. We measured the overhead of our proposal as the number of all request, proof, and verification messages exchanged. To measure the overhead of constrained routing, we measured the number of routing failure tests triggered, the number of test successes, and the number of replica messages sent in the block lookup tests. Combining these values and certificate and hash sizes and equations from [4], we compute the total overhead. We




Fig. 17. Comparing overhead of CFS block lookups in our system and constrained routing. Attack Type 2, 1,000-node network.

optimized our implementation of the constrained routing approach to use the best parameters that minimize errors in the routing failure test. Where appropriate, we used identical parameter values from our implementation (e.g., leafset size = 8, IdSize = 40 bytes). All other parameters used in the implementation and calculations are taken directly from [4]. Results in Fig. 17 show that our proposal incurs more than an order of magnitude less overhead compared to constrained routing.

### 6.2 Effectiveness of Defense against Collusion-Based Attacks

We now evaluate our solution for evading collusion-enhanced Eclipse attacks from the perspective of a single target node. As discussed in Section 4, we use a combination of periodic induced churn [6] and one-hop redirection through other nodes.

We model the Eclipse attack by initializing the network with 20 percent randomly chosen malicious nodes. We assume a node's routing table starts with the same rate, and introduce more malicious peers into the routing table at a constant rate, updated every 20 seconds. The target resets its routing table every 100 seconds, reducing the percentage of attackers back down to the same level as the rest of the network (20 percent). All verification requests are routed using one-hop indirection through four nodes chosen from our leafset or routing entries. We run the experiment for 1,200 seconds, and measure the portion of messages attacked, and the portion of all attacks detected. These results are plotted against infiltration rate (per 20 seconds) in Fig. 18.

We plot two separate results in Fig. 18. First, we look at the scenario where the attackers have sufficient resources to request enough nodeIds to attack all routing levels of the target node. Therefore, we introduce malicious peers uniformly at random across all routing levels (and leafset)




Fig. 18. Effectiveness of Eclipse attack defenses with induced churn and one-hop indirection. Attack Type 2, 1,000-node network.




Fig. 19. Effectiveness of various Eclipse attack defenses in CFS. Attack Type 2, 1,000-node network.

of the target node. We use indirection across the leafset (LS) nodes, and plot the results as “indirection using LS.” While indirection results in high detection at low infiltration rates, as infiltration rates increase, leafset nodes are also affected, and indirection is less successful. In the second scenario, we assume the attacker has limited resources, and only has enough nodeIds to attack the lower half of the routing table. In this case, we can use indirection across both leafset nodes and higher level routing entries. The result is detection rates similar to those of single node attacks. Note that we are effectively using leafset entries and higher level routing entries as “hard-to-compromise” nodes, similar to constrained routing tables.

The second scenario described above is more practical and follows directly from our analytical results. Our results show that our proposal works well in this practical case, and is generally able to detect around 90 percent of all attacks.

We implemented various proposals from Castro et al. and Condie et al. to evade collusion-based Eclipse attacks in CFS. We ran the CFS block lookup experiments under the Identity and Eclipse attack model, and plotted the results in Fig. 19. Our proposal performs the best, and works well even under extremely high rates of node compromise. More than 80 percent of lookups using our proposal succeed even when all lower level neighbor entries are malicious, while alternative solutions provide lookup success rates around 50 percent.

## 7 DISCUSSION AND CONCLUSIONS

In summary, we have proposed a general defense for the Identity attack on structured overlays. Using existence proofs, blacklists, and malice-aware routing, our system effectively detects, marks, and redirects traffic away from attackers. We showed that our techniques are easily applied to and highly effective on real applications such as CFS. Measurements on CFS show that they perform at least as well as other proposed approaches in detection and recovery, but require an order or magnitude lower costs in bandwidth overhead. In addition, our analysis shows that performing a basic Eclipse attack and corrupting the lower levels of a victim’s routing table is practical, but can be successfully defended using our mechanisms. The vulnerabilities of KBR we described here are common to all structured overlay protocols. As more critical applications are deployed on structured overlays, application designers must be aware of these attacks, and integrate defense mechanisms into their protocols and applications.

Finally, we recognize that the reliance of our defense mechanisms on a centralized CA might be restrictive for certain applications. For example, the fully decentralized DHT implemented in Azerus does not currently support a central CA. We believe, however, that this assumption would apply to any wide-area applications that has a central component such as an introduction server or load balancer. Examples include applications such as distributed storage services [22], [25], Voice-over-IP networks [16], and content delivery networks [14]. As an alternative to centralized CAs, we are also evaluating the feasibility of using identity-based encryption schemes [3] to generate unique nodeIds from nonforgeable identity information.

## ACKNOWLEDGMENTS

The authors would like to thank Lakshmi Ganesh (Cornell) for her help during the initial stages of the project. They also gratefully acknowledge the support from DARPA through the Control Plane Program (BAA04-11) and the US National Science Foundation (NSF) through CAREER Award 0546216.

## REFERENCES

- [1] R. Alebouyeh, M.S. Allen, K.P.N. Puttaswamy, and B.Y. Zhao, *Chimera Software Distribution*, <http://current.cs.ucsb.edu/projects/chimera>, 2008.
- [2] H. Ballani, P. Francis, and X. Zhang, “A Study of Prex Hijacking and Interception in the Internet,” *Proc. ACM SIGCOMM '07*, Sept. 2007.
- [3] D. Boneh and M. Franklin, “Identity Based Encryption from the Weil Pairing,” *SIAM J. Computing*, vol. 32, no. 3, pp. 586-615, 2003.
- [4] M. Castro et al., “Security for Structured Peer-to-Peer Overlay Networks,” *Proc. Fifth Symp. Operating Systems Design and Implementation (OSDI '02)*, Dec. 2002.
- [5] M. Castro et al., “Splitstream: High-Bandwidth Multicast in a Cooperative Environment,” *Proc. 19th ACM Symp. Operating Systems Principles (SOSP '03)*, Oct. 2003.
- [6] T. Condie et al., “Induced Churn as Shelter from Routing-Table Poisoning,” *Proc. 13th Ann. Network and Distributed System Security Symp. (NDSS '06)*, Feb. 2006.
- [7] F. Dabek, M.F. Kaashoek, D. Karger, R. Morris, and I. Stoica, “Wide-Area Cooperative Storage with CFS,” *Proc. 18th ACM Symp. Operating Systems Principles (SOSP '01)*, Oct. 2001.
- [8] F. Dabek, B. Zhao, P. Druschel, J. Kubiatowicz, and I. Stoica, “Towards a Common API for Structured P2P Overlays,” *Proc. Second Int'l Workshop Peer-to-Peer Systems (IPTPS '03)*, Feb. 2003.
- [9] E. Damiani et al., “A Reputation-Based Approach for Choosing Reliable Resources in Peer-to-Peer Networks,” *Proc. Ninth ACM Conf. Computer and Comm. Security (CCS '02)*, Nov. 2002.
- [10] G. DeCandia et al., “Dynamo: Amazon’s Highly Available Key-Value Store,” *Proc. 21st ACM Symp. Operating Systems Principles (SOSP '07)*, Oct. 2007.
- [11] D. Deeths and G. Brunette, “Using NTP to Control and Synchronize System Clocks,” Technical Report 816-1475-10, Sun Microsystems Inc., July 2001.
- [12] J.R. Douceur, “The Sybil Attack,” *Proc. First Int'l Workshop Peer-to-Peer Systems (IPTPS '02)*, Mar. 2002.
- [13] J. Falkner, M. Piatek, J.P. John, A. Krishnamurthy, and T. Anderson, “Profiling a Million User DHT,” *Proc. Internet Measurement Conf. (IMC '07)*, Oct. 2007.
- [14] M.J. Freedman, E. Freudenthal, and D. Mazires, “Democratizing Content Publication with Coral,” *Proc. First Symp. Networked Systems Design and Implementation (NSDI '04)*, Dec. 2004.
- [15] L. Ganesh and B.Y. Zhao, “Identity Theft Protection in Structured Overlays,” *Proc. First Workshop Secure Network Protocols (NPsec '05)*, June 2005.
- [16] S. Guha, N. Daswani, and R. Jain, “An Experimental Study of the Skype Peer-to-Peer Voip System,” *Proc. Fifth Int'l Workshop Peer-to-Peer Systems (IPTPS)*, 2006.

- [17] K. Gummadi et al., "The Impact of DHT Routing Geometry on Resilience and Proximity," *Proc. ACM SIGCOMM '03*, Sept. 2003.
- [18] A. Haeberlen, P. Kouznetsov, and P. Druschel, "PeerReview: Practical Accountability for Distributed Systems," *Proc. 21st ACM Symp. Operating Systems Principles (SOSP '07)*, Oct. 2007.
- [19] A. Haeberlen, A. Mislove, and P. Druschel, "Glacier: Highly Durable, Decentralized Storage Despite Massive Correlated Failures," *Proc. Second Symp. Networked Systems Design and Implementation (NSDI '05)*, May 2005.
- [20] K. Hildrum and J. Kubiatowicz, "Asymptotically Efficient Approaches to Fault-Tolerance in Peer-to-Peer Networks," *Proc. 17th Int'l Symp. Distributed Computing (DISC '03)*, Oct. 2003.
- [21] S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks," *Proc. 12th Int'l Conf. World Wide Web (WWW '03)*, May 2003.
- [22] S. Rhea et al., "Pond: The OceanStore Prototype," *Proc. Second USENIX Conf. File and Storage Technologies (FAST '03)*, Apr. 2003.
- [23] H. Rowaihy, W. Enck, P. McDaniel, and T.L. Porta, "Limiting Sybil Attacks in Structured P2P Networks," *Proc. IEEE INFOCOM, 2007*.
- [24] A. Rowstron and P. Druschel, "Pastry: Scalable, Distributed Object Location and Routing for Large-Scale Peer-to-Peer Systems," *Proc. 18th IFIP/ACM Int'l Conf. Distributed Systems Platforms (Middleware '01)*, Nov. 2001.
- [25] A. Rowstron and P. Druschel, "Storage Management and Caching in PAST, a Large-Scale, Persistent Peer-to-Peer Storage Utility," *Proc. 18th ACM Symp. Operating Systems Principles (SOSP '01)*, Oct. 2001.
- [26] A. Rowstron et al., "SCRIBE: The Design of a Large-Scale Event Notification Infrastructure," *Proc. Third Int'l Workshop Networked Group Comm. (NGC '01)*, Nov. 2001.
- [27] S. Saroiu, P.K. Gummadi, and S. Gribble, "A Measurement Study of Peer-to-Peer File Sharing Systems," *Proc. Multimedia Computing and Networking (MMCN '02)*, Jan. 2002.
- [28] A. Singh, T.-W. Ngan, P. Druschel, and D. Wallach, "Eclipse Attacks on Overlay Networks: Threats and Defenses," *Proc. IEEE INFOCOM, Apr. 2006*.
- [29] E. Sit and R. Morris, "Security Considerations for Peer-to-Peer Distributed Hash Tables," *Proc. First Int'l Workshop Peer-to-Peer Systems (IPTPS '02)*, Mar. 2002.
- [30] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications," *Proc. ACM SIGCOMM, 2001*.
- [31] G. Swamyathan, B.Y. Zhao, and K.C. Almeroth, "Exploring the Feasibility of Proactive Reputations," *Proc. Fifth Int'l Workshop Peer-to-Peer Systems (IPTPS '06)*, Feb. 2006.
- [32] H. von Schelling, "Coupon Collecting for Unequal Probabilities," *Am. Math. Monthly*, vol. 61, pp. 306-311, 1954.
- [33] K. Walsh and E.G. Sirer, "Evaluation of a Deployed, Distributed Object Reputation System for Peer-to-Peer Filesharing," *Proc. Third Symp. Networked Systems Design and Implementation (NSDI '06)*, May 2006.
- [34] E.W. Weisstein, *Coupon Collector's Problem*, <http://mathworld.wolfram.com>, 2008.
- [35] H. Yu, P.B. Gibbons, M. Kaminsky, and F. Xiao, "A Near-Optimal Social Network Defense against Sybil Attacks," *Proc. IEEE Symp. Security and Privacy (S&P '08)*, May 2008.
- [36] H. Yu, M. Kaminsky, P.B. Gibbons, and A.D. Flaxman, "Sybilguard: Defending against Sybil Attacks via Social Networks," *IEEE/ACM Trans. Networking*, vol. 16, no. 3, pp. 576-589, June 2008.
- [37] B.Y. Zhao, L. Huang, J. Stribling, S.C. Rhea, A.D. Joseph, and J. Kubiatowicz, "Tapestry: A Global-Scale Overlay for Rapid Service Deployment," *IEEE J. Selected Areas in Comm.*, vol. 22, no. 1, Jan. 2004.
- [38] B.Y. Zhao, J.D. Kubiatowicz, and A.D. Joseph, "Tapestry: An Infrastructure for Fault-Tolerant Wide-Area Location and Routing," Technical Report CSD-01-1141, U.C. Berkeley, 2001.
- [39] L. Zhuang, F. Zhou, B.Y. Zhao, and A. Rowstron, "Cashmere: Resilient Anonymous Routing," *Proc. Second Symp. Networked Systems Design and Implementation (NSDI '05)*, May 2005.



**Krishna P.N. Puttaswamy** received the BS degree in computer science from the National Institute of Technology, Karnataka, India, in 2003. He is a PhD candidate in the Department of Computer Science, University of California at Santa Barbara, Santa Barbara. His research interests include the areas of security and privacy, and large-scale distributed systems.



**Haitao Zheng** received the BS degree (with highest honor) from Xian Jiaotong University in 1995 and the MSEE and PhD degrees in electrical and computer engineering from the University of Maryland, College Park, in 1998 and 1999, respectively. She joined the Wireless Research Lab, Bell-Labs, Lucent Technologies as a member of technical staff in August 1999, and moved to Microsoft Research Asia as a project lead and a researcher in March 2004.

Since September 2005, she has been an assistant professor in the Department of Computer Science, University of California at Santa Barbara, Santa Barbara. Her research interests include wireless systems and networking and multimedia computing. She was named as one of the 2005 MIT Technology Review Top 35 Innovators under the age of 35 for her work on cognitive radios. Her work was selected by MIT Technology Review as one of the 10 Emerging Technologies in 2006. She also received the 2002 Bell Laboratories President's Gold Award from Lucent Bell-Labs and the 1998-1999 George Harhalakis Outstanding Graduate Student Award from the Institute of System Research, University of Maryland.



**Ben Y. Zhao** received the BS degree from Yale University and the MS and PhD degrees in computer science from the University of California, Berkeley. He is a faculty member in the Department of Computer Science, University of California at Santa Barbara, Santa Barbara. His research interests include the areas of security and privacy, large-scale networking, and distributed systems. He is a recent recipient of the US National Science Foundation's CAREER award (2005), the Massachusetts Institute of Technology Review's TR-35 Award (Top 35 Young Innovators Under 35) in 2006, and ComputerWorld's Top 40 IT Innovators under 40.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).