

The LOCKSS Peer-to-Peer Digital Preservation System

PETROS MANIATIS

Intel Research

MEMA ROUSSOPOULOS

Harvard University

TJ GIULI

Stanford University

DAVID S. H. ROSENTHAL

Stanford University Libraries

and

MARY BAKER

HP Labs

The LOCKSS project has developed and deployed in a world-wide test a peer-to-peer system for preserving access to journals and other archival information published on the Web. It consists of a large number of independent, low-cost, persistent Web caches that cooperate to detect and repair damage to their content by voting in “opinion polls.” Based on this experience, we present a design for and simulations of a novel protocol for voting in systems of this kind. It incorporates rate limitation and intrusion detection to ensure that even some very powerful adversaries attacking over many years have only a small probability of causing irrecoverable damage before being detected.

Categories and Subject Descriptors: H.3.7 [Information Storage and Retrieval]: Digital Libraries; D.4.5 [Operating Systems]: Reliability

This work is supported by the National Science Foundation (Grant No. 0205667), by the Andrew W. Mellon Foundation, by Sun Microsystems Laboratories, by the Stanford Networking Research Center, by DARPA (contract No. N66001-00-C-8015), by MURI (award No. F49620-00-1-0330), and by Sonera Corporation. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of these funding agencies.

This article is the extended version of an earlier conference article [Maniatis et al. 2003].

Authors’ addresses: P. Maniatis, Intel Research, 2150 Shattuck Avenue Ste. 1300, Berkeley, CA 94704; email: petros.maniatis@intel.com; M. Roussopoulos, Harvard University, 33 Oxford Street, Cambridge, MA 02138; email: mema@eecs.harvard.edu; TJ Giuli, Computer Science Department, Stanford University, Gates Building 4A-416, 353 Serra Mall, Stanford, CA 94305-9040; email: giuli@cs.stanford.edu; D. S. H. Rosenthal, LOCKSS, 1454 Page Mill Rd., Palo Alto, CA 94304; email: dshr@stanford.edu; M. Baker, HP Labs, 1501 Page Mill Road, Mail Stop 1183, Palo Alto, CA 94304; email: mary.baker@hp.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 0734-2071/05/0200-0002 \$5.00

General Terms: Design, Economics, Reliability

Additional Key Words and Phrases: Rate limiting, replicated storage, digital preservation

1. INTRODUCTION

Academic publishing is migrating to the Web [Mogge 1999; Tenopir 2004], forcing the libraries that pay for journals to transition from purchasing copies of the material to renting access to the publisher's copy [Keller et al. 2003]. Unfortunately, rental provides no guarantee of long-term access. Librarians consider it one of their responsibilities to provide future readers with access to important materials. After millennia of experience with physical documents, they have techniques for doing so: acquire lots of copies of the document, distribute them around the world, and lend or copy them when necessary to provide access.

In the LOCKSS¹ program (Lots Of Copies Keep Stuff Safe), we model the physical document system and apply it to Web-published academic journals, providing tools for libraries to take custody of the material to which they subscribe, and to cooperate with other libraries to preserve it and provide access. The LOCKSS approach deploys a large number of independent, low-cost, persistent Web caches that cooperate to detect and repair damage by voting in "opinion polls" on their cached documents. The initial version of the system [Rosenthal and Reich 2000] has been under test since 1999 at about 50 libraries worldwide, and entered production use at many more libraries in 2004. Unfortunately, the protocol now in use does not scale adequately, and analysis of the first design for a revised protocol [Michalakis et al. 2003] showed it to be insufficiently resistant to attack.

In this work, we present a design for and simulations of a new peer-to-peer opinion poll protocol that addresses these scaling and attack resistance issues. We plan to migrate it to the deployed system shortly. The new protocol is based on our experience with the deployed LOCKSS system and the special characteristics of such a long-term large-scale application. Distributed digital preservation, with its time horizon of many decades and lack of central control, presents both unusual requirements, such as the need to avoid long-term secrets like encryption keys, and unusual opportunities, such as the option to make some system operations inherently very time-consuming without sacrificing usability.

Digital preservation systems must resist both random failures and deliberate attack for a long time. Their ultimate success can be judged only in the distant future. Techniques for evaluating their design must necessarily be approximate and probabilistic; they share this problem with encryption systems. We attempt to evaluate our design in the same way that encryption systems are evaluated, by estimating the computational effort an adversary would need to achieve a given probability of the desired result. In an encryption system, one such desired result is to recover the plaintext. In our case, it is to have the system deliver a corrupt copy of a document. These estimates can be converted to monetary costs

¹LOCKSS is a trademark of Stanford University.

using technology cost curves, and thus compared to the value of the plaintext or document at risk.

We introduce our design principles (Section 2) and the deployed test system (Section 3), then describe our new protocol and the reasons for the design decisions we made (Section 4). We present and analyze some attacks its adversary can mount (Section 5), introduce our experimental setup (Section 6), and then evaluate our system under three types of such attack, aimed at undetected direct corruption of a document, at discrediting the system by causing alarms, and at slowing the system down (Section 7). We conclude with related and future work (Sections 8 and 9), as well as a summary of our contributions (Section 10). The Appendix discusses some economic considerations.

Briefly, our simulations show a system that resists for decades an adversary capable of unlimited sustained effort, by preventing him from applying effectively more effort to the system than his victims do. Even assuming that an implementation flaw hands an adversary instantaneous control of one-third of the peers, his sustained effort can increase the probability of a reader seeing a damaged copy by no more than a further 3%, in the worst case. The system has a high probability of detecting an attack that can cause permanent damage before that damage becomes irreversible, while producing very few false positives due to random faults.

We believe this protocol and its underlying principles are novel and will prove useful in the design of other long-term large-scale applications operating in hostile environments.

2. DESIGN PRINCIPLES

Digital preservation systems have some unusual features. First, such systems must be very cheap to build and maintain, which precludes high-performance hardware such as RAID [Patterson et al. 1988] or complicated administration. Second, they need not operate quickly. Their purpose is to prevent rather than expedite change to data. Third, without central control and in the face of possible interference from attackers or catastrophic failures of storage media such as fire or theft, they must function properly for decades. These features, combined with our experience building and maintaining other large-scale distributed systems, lead to the very conservative design principles we use:

- Cheap storage is unreliable.* We assume that in our timescale no cheap and easy-to-maintain storage is reliable [ConservationOnline 2003; Fullmer 2002; Quirke 2002]. Note that write-once media are at least as unreliable as hard disks [Horlings 2003; Manns and Shahani 2003], precluding alternative designs dependent on storing documents or their hashes on CD-R (in our current deployment the CD-R containing the peer software is the cause of the vast majority of errors).
- Use no long-term secrets.* Or, to quote Diffie [2003], “The secret to strong security: less reliance on secrets.” Long-term secrets, such as private keys, are too vulnerable for our application. These secrets require storage that is effectively impossible to replicate, audit, repair, or regenerate. Over time they are

likely to leak; recovering from such leakage is extraordinarily difficult [Davis 1996; Venema 1996].

- Use inertia.* The goal of a digital preservation system is to prevent change. Some change is inevitable, and the system must repair it, but there is never a need for rapid change. A system that fails abruptly, without warning its operators in time for them to take corrective action and prevent total failure [Staniford et al. 2002], is not suitable for long-term preservation. Rate limiting has proved useful in other areas [Williamson 2002]; we can exploit similar techniques because we have no need for speed.
- Avoid third-party reputation.* Third-party reputation information is subject to a variety of problems, especially in the absence of strong peer identities. It is vulnerable to slander and subversion of previously reliable peers. If evidence of past good behavior is accumulated, an attacker can “cash in” a history of good behavior in low-value interactions by defecting in a single high-value interaction [Wingfield 2002].
- Reduce predictability.* Attackers predict the behavior of their victim to choose tactics. Making peer behavior depend on random combinations of external inputs and internal state reduces the accuracy of these predictions, thereby reducing their value to the attacker.
- Integrate intrusion detection intrinsically.* Conventional intrusion detection systems are extrinsic to the application being protected. They have to operate with less than full information about it, and may themselves become a target. Systems with bimodal behavior [Birman et al. 1999] can provide intrinsic intrusion detection by surrounding good states with a “moat” of forbidden states that are almost never reachable in the absence of an attack, and that generate an alarm when reached.
We believe this mechanism to be fundamentally more robust than layering an intrusion detection system on top of an application; it does, however, share with conventional intrusion detection systems the notion that repelling attacks on the system must be a cooperative effort between the software and the humans responsible for it.
- Assume a strong adversary.* The LOCKSS system preserves e-journals that have intrinsic value and contain information that powerful interests might want changed or suppressed. Today, a credible adversary is an Internet worm whose payload attacks the system using tens of thousands of hosts. We must plan for future, more powerful attacks.

The LOCKSS design is very conservative, appropriately so for a preservation system. By applying these design principles, our goal is to build a system in which, with high probability, even a powerful adversary fails to cause irrecoverable damage without detection.

3. THE EXISTING LOCKSS SYSTEM

The LOCKSS system models the techniques that librarians use for physical documents to solve the similar problem of preserving access to e-journals; it makes it appear to a library’s patrons that pages remain available at their

original URLs even if they are not available there from the publisher. We thus preserve access to the material via common techniques such as links, bookmarks, and search engines. To do this, participating libraries run persistent Web caches that

- collect* by crawling the journal Web sites to preload themselves with newly published material,
- distribute* by acting as a limited proxy cache for the library’s local readers, supplying the *publisher’s* copy if it is available and the local copy otherwise, and
- preserve* by cooperating with other caches that hold the same material to detect and repair damage.

Caches cooperate by participating in “opinion polls” in a peer-to-peer network. In each, a sample of peers votes on a specified part of the content.

Polls provide peers with confidence in content *authenticity* and *integrity*. Journal publishers do not currently sign the material they distribute, nor do they provide a manifest describing the files forming a given paper, issue, or volume, and the crawling process is unreliable. Furthermore, no completely reliable long-term storage medium is available. Catastrophic failures such as fire, theft, and hacker break-in can wipe out or alter any storage medium without possibility of recovery. Evidence that many peers independently obtained and agree with each other on the material seems to be the best available guarantee that content is authentic and correctly preserved.

Peers vote on large *archival units* (AUs), normally a year’s run of a journal. Because each peer holds a different set of AUs, the protocol treats each AU independently. If a peer loses a poll on an AU, it calls a sequence of increasingly specific partial polls within the AU to locate the damage. Other peers cooperate with the damaged peer if they remember it agreeing with them in the past about the AU, by offering it a good copy, in the same way they would for local readers.

This mechanism defends against two important problems endemic to peer-to-peer systems: *free-loading* and *theft*. First, the only benefit a peer obtains from the system is a repair, and to obtain it the peer must have participated in the past, which precludes free-loading. Second, a peer only supplies material to a peer that proved in the past that it had that material, so the system does not *increase* the risk of theft. In this way LOCKSS peers provide a distributed, highly replicated, self-healing store of data that does not materially increase the risks that publishers already run. This is important; under the Digital Millennium Copyright Act (DMCA),² publishers must give permission for libraries to preserve their material.

Library budgets are perennially inadequate [ARL—Association of Research Libraries 2001]. To be effective, any digital preservation system must be affordable in the long term. Minimizing the cost of participating in the LOCKSS system is essential to its success, so individual peers are built from low-cost,

²Available from the Electronic Frontier Foundation. Go to Web site <http://www.eff.org/IP/DCMA/>.

unreliable technology. A generic PC with three 180-GB disks currently costs under \$1000 and would preserve about 210 years of the largest journal we have found (the *Journal of Biological Chemistry*) for a worst-case hardware cost of less than \$5 per journal/year. This is equivalent to less than 1¢ per SOSOP proceedings.

Using these generic PCs we can build a system with acceptable performance. If peers check each AU every 3 months and split their time equally between calling polls and voting in polls called by others, each peer has 45 days in which to call one poll for each of its AUs. If there are 210 AUs, each poll should last about 5 h. With our new protocol, this size of AU costs the caller about 1040 s for each peer it invites to vote (Section 6.1). Each poll could thus involve about 17 peers, more than in the current tests.

Peers require little administration [Rosenthal 2003], relying on cooperation with other caches to detect and repair failures. There is no need for offline backups on removable media. Creating these backups, and using them when readers request access to data, would involve excessive staff costs and latencies beyond a reader's attention span [U. S. Government 2004].

4. THE NEW OPINION POLL PROTOCOL

In this section we outline, describe, and justify our new LOCKSS opinion poll protocol. We give an overview of the protocol and introduce relevant apparatus and notation before specifying the protocol in more detail in Section 4.1. In Section 4.2, we distill the main techniques we employ in our design and explain how and why the protocol applies them.

To simplify the analysis of the new protocol, we abstract the relevant features of the existing system. We consider a population of peers preserving a copy of a single AU, obtained but no longer available from a publisher. We ignore the divide-and-conquer search for damage in a real, multifile journal. Each peer uses one of a number of independent implementations of the LOCKSS protocol to limit common-mode failures. Each peer's AU is subject to the same low rate of undetected random damage.

While a *peer* is any node that participates with benign or malicious intent in the LOCKSS protocol, we make the following distinctions between different types of peers in the rest of this article:

- A *malign* peer is part of a conspiracy of peers attempting to subvert the system.
- A *loyal* peer is a nonmalign peer, that is, one that follows the LOCKSS protocol at all times.
- A *damaged* peer is a loyal peer with a damaged AU.
- A *healthy* peer is a loyal peer with the correct AU.

The overall goal of our design is that there be a high probability that loyal peers are in the healthy state despite failures and attacks, and a low probability that even a powerful adversary can damage a significant proportion of the loyal peers without detection.

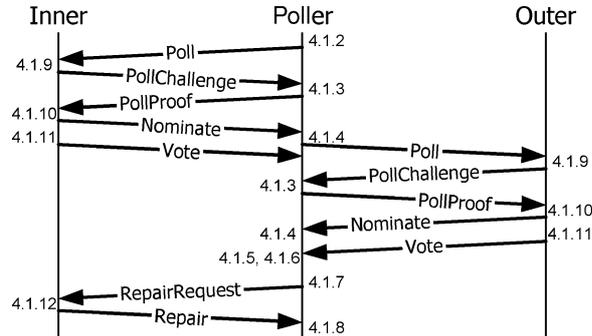


Fig. 1. The protocol messages exchanged during a poll between the poller and voters. Time flows from top to bottom. Next to each phase in the execution of the protocol, we give the section number that provides the pertinent description.

A LOCKSS peer calls opinion polls on the contents of an AU it holds at a rate much greater than any anticipated rate of random damage. It invites into its poll a small subset of the peers it has recently encountered, hoping they will offer votes on their version of the AU. Unless an invited peer is busy, it computes a fresh digest of its own version of the AU, which it returns in a vote. If the caller of the poll receives votes that overwhelmingly agree with its own version of the AU (a *landslide win*), it is satisfied and waits until it has to call a poll again. If it receives votes that overwhelmingly disagree with its own version of the AU (a *landslide loss*), it repairs its AU by fetching the copy of a voter who disagreed, and reevaluates the votes, hoping now to obtain a landslide win for its repaired AU. If the result of the poll justifies neither a landslide win nor a landslide loss (an *inconclusive* poll), then the caller raises an alarm to attract human attention to the situation.

The protocol supports two roles for participating peers (see Section 4.1). First, the *poller* calls polls on its own AUs and is the sole beneficiary of the poll result. Second, the *poll participant* or *voter* is a peer who is invited into the poll by the poller and who votes if it has the necessary resources. A voter need not find out the result of a poll in which it votes. Voters for a given poll are divided into two groups: the *inner circle* and the *outer circle*. Inner circle voters are chosen by the poller from those peers it has already discovered. The poller decides the outcome of the poll solely on inner circle votes. Outer circle voters are chosen by the poller from peers nominated by inner circle voters. The poller uses outer circle votes to perform discovery, that is, to locate peers that it can invite into future inner circles for its polls.

LOCKSS peers communicate in two types of exchanges (see Section 4.1). First, a poller uses unicast datagrams to invite voters into a poll. Second, a poller may contact its voters to request a repair for its AU using a bulk transfer protocol. In both cases, communication is encrypted via symmetric session keys, derived using Diffie-Hellman key exchanges [Diffie and Hellman 1976] between the poller and each of its voters. After the poll, session keys are discarded. Figure 1 shows a typical message exchange between a poller and its inner and outer circles.

The LOCKSS opinion poll protocol requires both pollers and voters to expend provable computational effort [Dwork and Naor 1992] in amounts related to underlying system operations (hashing of an AU), as a means of limiting Sybil attacks [Douceur 2002]. We describe in Section 4.2.1 how these amounts are determined, and how proofs of effort are constructed, verified, and used; we give extensive details in Appendix A.3. In the protocol description below we simply refer to the generation and verification of effort proofs.

In the remainder of this article, we use the following notation for system parameters:

- A* Maximum number of discredited challenges allowed in a poll (Section 4.1.2)
- C* Proportion of the reference list refreshed using friends at every poll (*friend bias* in Section 4.1.8)
- D* The maximum number of votes allowed to be in the minority of a poll (Section 4.1.6)
- E* Maximum age of unused reference list entries (Section 4.1.8)
- I* Number of outer circle nominations per inner circle voter (Section 4.1.10)
- N* Number of inner-circle peers invited into a poll (Section 4.1.2)
- Q* Number of valid inner-circle votes required to conclude a poll successfully (*quorum*) (Section 4.1.2)
- R* Mean interval between two successive polls called by a peer on an AU (Section 4.1.1)

We use the following convenience variables:

- L* Number of loyal voters in the inner circle (Section 5.5.1)
- M* Number of malign voters in the inner circle (Section 5.5.1)
- V* Number of inner-circle peers whose vote is received and verified to be valid (Section 4.1.6)

4.1 Detailed Description

In this section, we present in detail how the opinion poll protocol works. In Section 4.2, we explain the reasoning behind our major design decisions.

Each peer maintains two peer lists for every AU it holds: the *reference list*, which contains information about other LOCKSS peers it has recently encountered; and the *friends list*, which contains information about LOCKSS peers with whose operators or organizations the peer has an out-of-band relationship. A peer maintains for every AU it holds a *poll counter* that records the number of polls the peer has called on that AU since first acquiring it.

Reference list entries have the form [*peer IP address, time inserted*]. They are added to or removed from the list by the protocol. The value of the *time inserted* field is set to the value of the poll counter at the time the entry is inserted into the reference list. Friends list entries contain only a peer IP address. They are added to or removed from the list by the peer's operator, as her affiliations with other institutions change over time.

A peer who is not in the process of calling a poll for an AU also maintains a *refresh timer* for that AU. When the timer expires, the peer calls a new poll for the AU (Section 4.1.2).

In what follows, we describe in detail the different protocol steps, including bootstrapping (Section 4.1.1), the poller's point of view (Sections 4.1.2 to 4.1.8), the voter's point of view (Sections 4.1.9 to 4.1.12), and the intrusion detection alarms (Section 4.1.13).

4.1.1 Bootstrapping. When a peer first enters a LOCKSS network for a given AU, or when it reinitializes after a failure, it copies all entries from its current friends list into its reference list, and sets its refresh timer with a random expiration time with mean value R . In our simulations, we choose this random value uniformly from an interval centered at R .

4.1.2 Poll Initiation. To call a new poll on an AU, a LOCKSS peer chooses a fresh, random poll identifier and N random peers from its reference list, which it inserts into its *inner circle list*. For each inner circle peer, the poller chooses a fresh, random Diffie-Hellman public key, and sends that peer a Poll message of the form [Poll ID, DH Public Key]. Then the poller waits for PollChallenge messages from the invited inner circle peers (see Section 4.1.9) and sets a *challenge timer* to stop waiting.

The poller removes from its inner circle list those peers who respond with a negative PollChallenge message, those who do not respond by the time the challenge timer expires, and those from whom the poller receives multiple PollChallenge messages with conflicting contents.

Peers removed because of conflicting PollChallenge messages are said to be *discredited*. A discredited peer may be self-inconsistent because of a local fault; alternatively, it may be the victim of a spoofer located near it or near the poller. Either way, the poller cannot tell the intended from the faulty or malicious PollChallenge messages, so it removes discredited peers from the poll. If the poller discredits more than A peers in a poll, it suspects a local spoofer and raises a spoofer alarm (see Section 4.1.13).

For all inner circle peers who send a valid, affirmative challenge, the poller computes the provable effort for its poll invitation (Section 4.1.3). If the poller ends up with fewer inner circle peers than Q , the minimum required, it invites additional peers into the poll via more Poll messages, or aborts the poll if it has no more peers in its reference list.

4.1.3 Poll Effort. For every received affirmative PollChallenge message, the poller produces some computational effort that is provable via a *poll effort proof* (see Section 4.2.1). The effort and its proof are cryptographically derived from the poll identifier and the potential voter's challenge. The poller returns this poll effort proof to the sender of the associated PollChallenge message within a PollProof message of the form [Poll ID, poll effort proof, poller's challenge], encrypted using the session key. The poller's challenge is a fresh nonce chosen by the poller to ensure that the votes returned by voters are freshly produced.

The poller also sends PollProof messages to voters who responded to the initial invitation with a negative PollChallenge. The poller need not expend

computational effort for negative voter's challenges; it can use a random value as the poll effort proof (see Section 4.2.5).

After sending all PollProof messages, the poller waits for Nominate messages (Section 4.1.10) and sets the *nomination timer* to stop waiting. When all Nominate messages arrive or the timer expires, the poller forms its outer circle.

4.1.4 Outer Circle Invitation. The poller discovers new peers that maintain the same AU by forming an outer circle based on the Nominate messages returned from its inner circle voters (Section 4.1.10). Discovery is important when the reference list is short (close to N), but less necessary when the reference list is long. Therefore, the poller picks an outer circle size that, when added to its current reference list, would achieve a target reference list size (in our simulations $3 \times N$).

To form its outer circle, the poller removes from every nomination list peers already contained in its reference list, and then it chooses an *equal* number of peers from every nomination list at random for its *outer circle list*; as a result, every inner circle nominator affects the outer circle equally. The poller invites outer circle peers into the poll in a manner identical to inviting the inner circle, as outlined in Sections 4.1.2 and 4.1.3. Nominate messages from outer circle voters are ignored.

The poller starts collecting Vote messages once it has sent its inner circle PollProof messages (Section 4.1.3). As soon as it finishes the construction of poll effort proofs for the outer circle, it sets a *vote timer* to stop collecting Vote messages. When all expected Vote messages have arrived, or the vote timer has expired, the poller verifies the votes.

4.1.5 Vote Verification. Vote verification deems votes to be one of *invalid*, *valid but disagreeing* with the poller's AU, or *valid and agreeing* with the poller's AU. Votes are constructed in rounds (Section 4.1.11) and are thus verified in rounds.

In each such round, the poller verifies the proof of computational effort included in the Vote message for the corresponding voting round, ensuring that the effort proof is cryptographically dependent on the poller's challenge (sent to the voter within the PollProof message). If the proof is incorrect, the poller deems the vote invalid and verification stops. Otherwise, if the vote has yet to be deemed disagreeing, the poller hashes the proof with the corresponding portion of its own copy of the AU; if the result does not match the hash in the vote, the vote is declared disagreeing. The poller skips hashing the AU if it has already deemed the vote disagreeing and uses the values in the Vote message to proceed with validity verification, instead.

If all proofs of effort are correct, the poller deems the vote valid. If all AU hashes match, the poller deems the vote agreeing, and disagreeing otherwise. Invalid votes result in the removal of the offending voter from the poll (inner or outer circle), and from the poller's reference list, since they indicate fault or malice. When the poller has verified all received Vote messages, it tabulates the results.

4.1.6 Vote Tabulation. The poller tabulates the valid votes from the inner circle to determine whether its AU replica is correct. If the number V of valid inner circle votes is greater than or equal to the quorum Q , then the voter acts as follows:

- Agreeing votes are no more than D .* The poll is a landslide loss. The poller considers its current AU copy damaged and repairs it (Section 4.1.7).
- Agreeing votes are at least $V - D$.* The poller considers its current copy the prevailing one (landslide win). This is the only way in which an opinion poll *concludes successfully*. The poller updates its reference list (Section 4.1.8) and schedules another poll at a time whose distance into the future is uniformly random with mean R .
- Agreeing votes are more than D but fewer than $V - D$.* The poller considers the poll inconclusive and raises an alarm (Section 4.1.13).

If the poller has been unable to accumulate Q valid votes from its inner circle, then it does not make a decision on its AU; it updates its reference list (Section 4.1.8), and immediately calls another poll. If it has failed to obtain Q votes in a poll on this AU for a long time, the poller raises an inter-poll interval alarm (Section 4.1.13).

4.1.7 Repair. If the poller decides that its AU is damaged, it picks at random one of the disagreeing inner circle voters and sends it an encrypted Repair-Request message containing the poll identifier. If it receives a Repair message (see Section 4.1.12), the poller re-verifies any disagreeing votes against the new AU (Section 4.1.5) and retabulates the results (Section 4.1.6). If it does not receive a Repair message, it picks another disagreeing inner circle voter and tries again.

The poller discards repairs that disagree with the vote of the supplier of the repair and removes the supplier from the reference list. Any inconsistency between a vote and the AU on which that vote was purportedly computed indicates a fault or malice at the repair supplier.

Note that the poller need only make up to D repair attempts. If during repairs the poller has agreed with more than D but fewer than $V - D$ voters in total, it knows that reaching a landslide win through subsequent repairs is impossible. It deems the poll inconclusive, raising the corresponding alarm (Section 4.1.13).

4.1.8 Reference List Update. Once a poll has concluded successfully, whether initially or after a repair, the poller updates its reference list by the following four steps. First, it removes those peers on whose votes it based its decision. Specifically, it removes all disagreeing inner circle voters and enough randomly chosen agreeing inner circle voters to make the total number of removals Q (see Section 4.2.2). Second, it resets the *time inserted* field of the remaining agreeing inner circle voters in the reference list with the current poll counter. Third, it inserts all outer circle peers whose votes were valid and (eventually) agreeing. Fourth, it inserts randomly chosen entries copied from its friends list until there are a fraction C of them in its reference list (see Section 4.2.4). Finally, it removes all peers that have not voted in the last

E polls it has called, that is, those entries whose *time inserted* is at least E polls less than the current poll counter.

A poll may fail to attract Q or more valid votes from inner circle voters. If so, the poller ignores all disagreeing votes, but refreshes or inserts into the reference list the agreeing votes from both circles.

4.1.9 Vote Solicitation. This and subsequent sections describe the opinion poll protocol from the point of view of an invitee.

When a LOCKSS peer receives a Poll message from a poller, it chooses a fresh, random *voter's challenge* value, a fresh, random Diffie-Hellman public key, and computes a symmetric session key from it and from the poller's public key included in the Poll message. If the peer is not currently the poller or a voter of another poll, it decides to vote in this new poll. It sends back a PollChallenge message of the form [*Poll ID*, *DH Public Key*, *voter's challenge*, *YES*]. Otherwise, it declines to vote and responds with a PollChallenge message of the form [*Poll ID*, *DH Public Key*, *voter's challenge*, *NO*]. In either case, the voter's challenge and the YES/NO bit are encrypted with the session key.

Finally, the peer sets an *effort timer* and waits for a PollProof message from the poller (see Section 4.1.3). If the message never arrives, the peer discards all poll state. Otherwise, the peer verifies the PollProof message.

4.1.10 Poll Effort Verification. A voter verifies the poll effort proof it receives in a PollProof message using the poll identifier and the challenge it sent to the poller (Section 4.1.9). If the verification succeeds, the voter chooses I other peers at random from its own reference list, and nominates them for inclusion into the poller's outer circle via a Nominate message of the form [*Poll ID*, *Nominations*] encrypted with the session key. Then the voter constructs its vote.

4.1.11 Vote Construction. A vote consists of a hash of the AU interleaved with provable computational effort.

Vote computation is divided into rounds, each returning a proof of computational effort and a hash of this proof with a portion of the AU. In each round, the computational effort and the AU portion that is hashed both double in size (see Section 4.2.1). The first round takes as input, and is cryptographically dependent upon, the poller's challenge, poll identifier, and identity of the voter. Subsequent rounds take as input, and are cryptographically dependent upon, the output of the previous round. The voter sends the proofs of computational effort and AU hashes from all rounds in a single encrypted Vote message to the poller.

Note that we use the term *cryptographically dependent* informally to convey that a polynomially bound adversary has only negligible probability of deriving the output without the input. For example, the output of the first round, above, includes one-way computations on the poller's challenge, the poll identifier, etc.; and an adversary has negligible probability of computing or guessing the output of that first round without first knowing the poller's challenge or any of the other inputs to the round.

A peer who refused to participate in the poll sends back to the poller an encrypted Vote message with bogus contents.

4.1.12 *Repair Solicitation.* After the vote, the poller may request a voter to supply a repair via a RepairRequest message (Section 4.1.7). If that voter has conducted a poll on the same AU in the past, in which the poller supplied a valid agreeing vote, then the voter responds to the request with a Repair message. The Repair message contains the poll identifier and its own copy of the AU, encrypted with the symmetric session key. Otherwise, the voter discards the request.

4.1.13 *Alarms.* LOCKSS peers raise alarms when they suspect that an attack is under way. The alarm requests human involvement in suppressing the attack, and is thus expensive.

An *inconclusive poll alarm* suggests that the library should contact others, examine the differences between their copies and determine a cause. Any compromised nodes found during this process are repaired. If the institutions hosting the peers voting for bad copies cannot be identified or do not cooperate, their peers are blacklisted.

A *local spoofing alarm* suggests that the network surrounding the peer should be audited and any compromised nodes removed. The cost of this alarm can be reduced by placing peers on their own subnets.

An *interpoll interval alarm* suggests that a poller has failed to conclude a poll after several consecutive attempts (currently 10). An attrition attack may be underway (Section 5.2), or the peer may no longer be fast enough to keep up with the system; human attention is needed in either case. Logs with large numbers of poll requests from previously unknown peers might lead to potential attackers who should be blacklisted.

4.2 Protocol Analysis

To defend the LOCKSS system from attack, we make it costly and time-consuming for an adversary to sway an opinion poll in his favor or to waste loyal peers' resources. This means the protocol must

- prevent the adversary from gaining a foothold in a poller's reference list (prevent him from populating it with malign peers),
- make it expensive for the adversary to waste another peer's resources, and
- make it likely that the adversary's attack is detected before it progresses far enough to cause irrecoverable damage.

We use provable recent effort, rate limiting, friend bias in the reference list, and obfuscation of protocol state to make it expensive and slow for an adversary to gain a significant foothold in a peer's reference list or to waste other peers' resources. We raise alarms when we detect signs of attack.

4.2.1 *Effort Sizing.* One application of our principle of inertia (Section 2) is that large changes to the system require large efforts. In a protocol where some valid messages cost nothing to produce, but cause the expenditure of great effort—for example, a cheap request causing its recipient to hash a large amount of data—this principle is unfulfilled. To satisfy our inertia requirement in the LOCKSS system, we adjust the amount of effort involved in message exchanges for voting, discovery, and poll initiation, by embedding extra, otherwise unnecessary effort.

For this purpose we need a mechanism satisfying at least three requirements. First, it must have an adjustable cost, since different amounts of additional effort are needed at different protocol steps. Second, it must produce effort measurable in the same units as the cost it adjusts (hashing in our case). Third, the cost of generating the effort must be greater than the cost of verifying it, which makes abuse expensive.

We use a mechanism for provable effort based on a class of *memory-bound functions* [Abadi et al. 2003] (MBF) proposed by Dwork et al. [2003] to prevent email spam. These cryptographic functions have a computation phase, yielding a short *proof of effort*, and a verification phase, which checks the validity of that proof. A parameter of the system sets the asymmetry factor by which the computation phase is more time consuming than the adjustable cost of the verification phase. We provide details on the MBF mechanism in Appendix A.1 and on how we set its parameters for our protocol in Appendix A.3.

MBFs are attractive for our purposes because the inherent cost of the hashing necessary for voting is also memory-bound, and because the difference in performance between available memory systems is much less than the difference in other characteristics such as CPU speed [Douceur 2002]. This observation has persisted across generations of technology. Nevertheless, if another mechanism for imposing cost becomes more attractive, the protocol could easily be revised to use it; it is the concept of imposing costs on peers that is important rather than the particular mechanism we use.

In voting, the cost of constructing a vote must be greater than the cost of processing the vote. We interleave AU hashing with effort proof generation and transmit the resulting proofs in the Vote message (Section 4.1.11). This ensures that bogus votes causing the poller to hash its AU in vain are more expensive to create than the effort they waste. The extra effort is interleaved with the hashing *in rounds* to prevent a cheap, bogus vote from wasting a lot of verification effort. The rounds ensure that generating a vote that is valid up to round $i - 1$ but then invalid costs more than its verification up to the i th round (see Appendix A.3).

In discovery, we require peers found via others' nominations to participate first in the outer circle of a poll and generate a valid agreeing but ineffectual vote before they are invited into the reference list. They must thus prove substantial effort (and wait on average more than time R) before they are able to affect the result of a poll. This makes it expensive and time-consuming for a malign peer to get an opportunity to vote maliciously and effectively.

Finally, in poll initiation, a poller must expend more effort than the cumulative effort it imposes on the voters in its poll. Otherwise, a malign peer would initiate spurious polls at no cost, causing loyal peers to waste their resources. We require the poller to prove more effort in the PollEffort message (Section 4.1.3) than the voter needs to verify that effort and then construct its vote.

4.2.2 Timeliness Of Effort. Our principle of avoiding third-party reputation leads us to use several techniques to ensure that only proofs of recent effort can affect the system. They prevent an adversary from exploiting evidence of good behavior accumulated over time. The requester of a poll effort

proof supplies a challenge on which the proof must depend; the first round of a subsequent vote depends on that proof, and each subsequent round depends on the round preceding it. Neither proofs nor votes can be precomputed.

Peers must supply a vote, and thus a proof of effort, to be admitted to the reference list, except for friends (see Section 4.2.4). If several polls take place without the admitted peer taking part, perhaps because it died, the poller removes it. If the admitted peer is invited and does take part in a poll, it must supply a vote and thus further proof of effort. After the poll the poller “forgets” the peer if it disagreed, or if it agreed but was chosen among the Q peers removed (Section 4.1.8). To limit the effect of this removal on loyal peers’ reference list size, we treat the poll as if it had a bare quorum and remove only the corresponding number of agreeing peers; additional agreeing peers do not affect the result and are thus treated as part of the outer circle for this purpose.

By these means we ensure that any peer identity, whether loyal or malign, must continually be sustained by at least a minimum rate of expenditure of effort if it is not to disappear from the system. Although the lack of long-term secrets makes it cheap for an adversary to create an identity, sustaining that identity for long enough to affect the system is expensive. Unless the adversary’s resources are truly unlimited, there are better uses for them than maintaining identities that do not contribute to his goals.

4.2.3 Rate Limiting. Another application of our principle of inertia (Section 2) is that the system should not change rapidly no matter how much effort is applied to it. We use rate-limiting techniques to implement this.

Loyal peers call polls autonomously and infrequently, but often enough to prevent random undetected damage from affecting readers significantly. This sets the effort required of the voters, and means that an adversary can damage a loyal peer only when that peer calls a poll. *The rate at which an attack can make progress is limited by the smaller of the adversary’s efforts and the efforts of her victims.* The adversary cannot affect this limit on the rate at which she can damage loyal peers.

4.2.4 Friend Bias. A protocol attacker (Section 5.2) needs to populate a loyal peer’s reference list with malign peers as a precondition for damaging its AU. We reduce the predictability of the mechanism by which the reference list is updated using *friend bias*.

It is important for a peer to avoid depending on a fixed set of peers for maintenance of its AU, because those peers may become faulty or subversion targets. It is equally important not to depend entirely on peers nominated by others whose motives may be unclear or unwholesome.

The poller hampers attempts to fill its reference list with malign conspirators by maintaining a fixed bias towards friends in that list, even as it adds outer circle agreeing voters. Absent an attack, the proportion of malign peers in both the outer circle and the friends list matches the population as a whole. An adversary’s attempt to subvert the random sampling process by nominating only malign peers raises the proportion of malign peers in the outer circle but not in the friends list. Friend bias reduces the effect of the attack because,

on average, the friends list is less malign than the outer circle, *even when the initial friends list contains subverted peers*.

The conference version of this article [Maniatis et al. 2003] used the term *churn* for this process. This caused confusion with the common use of the term to mean change in a population by additions or departures of members (e.g., Rhea et al. [2004]). Along with the change in terminology, we have updated the definition of friend bias for clarity. Whereas churn specified the number of friends added to the reference list at the end of every poll as a fraction of the reference list size, friend bias is the (target) proportion of friends in the reference list. For comparison, a churn of 2% results in an effective friend bias of approximately 8% and a churn of 10% results in an effective friend bias of approximately 40%. With the redefinition of this concept, we allow finer-grained control on how a peer biases its reference list toward its friends.

4.2.5 Obfuscation of Protocol State. Our design principles (Section 2) include assuming a powerful adversary, capable of observing traffic at many points in the network. We obfuscate protocol state in two ways to deny her information about a poll other than that obtained from the malign voters.

First, we encrypt all but the first protocol message exchanged by a poller and each potential voter, using a fresh symmetric key for each poll and voter. Second, we make all loyal peers invited into a poll, even those who decline to vote, go through the motions of the protocol, behind the cover of encryption. This prevents an adversary from using traffic analysis to infer state such as the number of loyal peers who actually vote in a specific poll. Note that in our modeled adversary and simulations we conservatively assume that the adversary *can* infer such information.

4.2.6 Alarms. In accordance with our design principle that intrusion detection be inherent in the system, the protocol raises an alarm when a peer determines that a poll is inconclusive, suspects local spoofing, or has been unable to complete a poll for a long time. Raising an alarm is thus expensive; a significant rate of false alarms would render the system useless (Section 7.1).

The expectation is that alarms result in enough loss to the adversary, for example by causing operators to remove damage, malign peers and compromised nodes, that a rational adversary will be highly motivated to avoid them, unless raising alarms is her primary goal.

5. ADVERSARY ANALYSIS

A peer-to-peer system running on a public network must expect to be attacked, even if the attackers have nothing tangible to gain. We present the capabilities we assume of adversaries (Section 5.1), explore the space of attacks they can mount in terms of their goals (Section 5.2), and the protocol roles that they can play (Section 5.3). We then dig deeper into the specific attack techniques available to adversaries (Section 5.4) and present in detail the particular adversary strategies we study in this paper (Section 5.5).

5.1 Adversary Capabilities

The adversary we study in this article controls a group of malign peers. We believe the following abilities match our “powerful adversary” design principle:

- Total information awareness.* Malign peers know each other. Any information known to one malign peer, including the identities of other malign peers involved in a poll, is immediately known to all.
- Perfect work balancing.* Any of the adversary’s nodes can perform work on his behalf and relay it instantaneously to the node presumed to have performed it.
- Perfect digital preservation.* The malign peers have magically incorruptible copies of both the good AU and as many bad ones as they require.
- Local eavesdropping.* The adversary is aware of the existence and contents of packets originating or terminating at a network on which he controls a physical node. He cannot observe the existence or contents of packets that originate and terminate anywhere else.
- Local spoofing.* From a routing realm in which he controls a physical node, the adversary can send IP packets whose ostensible source is any local address and destination is any Internet address; or he can send packets whose ostensible source is any Internet address and destination is a local address.
 However, the adversary cannot usefully send IP packets whose ostensible source and destination addresses are from routing realms within which he has no control. The protocol’s encryption handshake prevents peers from taking part in polls unless they can receive packets at their ostensible address (see Section 4.1.9). The adversary cannot benefit from spoofing IP addresses on whose traffic he cannot eavesdrop.
- Stealth.* A loyal peer cannot detect that another peer executing the LOCKSS protocol is malign.
- Unconstrained identities.* The adversary can increase the number of identities he can assume in the system, by purchasing or spoofing IP addresses.
- Exploitation of common peer vulnerabilities.* The adversary can instantaneously take over those peers from the LOCKSS population that run an implementation with the same exploitable vulnerability. He can then instantly change the state and/or the logic of the afflicted peers, causing them to become malign.
- Complete parameter knowledge.* Although loyal peers actively obfuscate protocol state (Section 4.2.5), we assume that the adversary knows the values of protocol parameters, including A , D , N and Q , as set by loyal peers.

We measure the adversary by the extent to which he can subvert loyal peers by any means, and by the total computational power he has available. Our analyses and simulations do not depend on identifying the cause of an individual peer becoming malign.

5.2 Adversary Attacks

Attacks against the LOCKSS system can be divided according to their functional approach into *platform* attacks and *protocol* attacks. Platform attacks seek to subvert the system hosting the LOCKSS implementation, either to tamper with its logic and data from the inside or to use a LOCKSS peer as a jumping point for further unrelated attacks. Our implementation makes considerable efforts to resist platform attacks [Rosenthal 2003]. We do not address platform attacks further in this article. We do, however, allow our simulated adversary to mount such attacks successfully, taking over a substantial portion of the peer population at the beginning of the simulations (see Section 6.1).

We can divide protocol attacks according to the role within the LOCKSS protocol whose operations the adversary manipulates for her purposes. In general, inner circle peers within a poll command greater power for mischief than outer circle peers, pollers can only affect their own polls but can waste loyal peers' resources, and spoofing and eavesdropping are most helpful when available near a loyal poller's network. We analyze attacks by role in Section 5.3.

Finally, we can divide attacks according to the goals of the adversary. Different adversary goals require different combinations of platform and protocol attacks. We list brief descriptions of some possible adversary goals below:

- Stealth modification.* The adversary wishes to replace the protected content with her version (the *bad* content). Her goal is to change, through protocol exchanges, as many replicas of the content held by loyal peers as possible without being detected, that is, before the system raises an alarm. Measures of her success are the proportion of loyal peers in the system hosting replicas of the bad content at the time of the first alarm, and the probability that a reader requesting the content from any peer in the system obtains the bad content.
- Nuisance.* The adversary wishes to raise frequent, spurious LOCKSS alarms to dilute the credibility of alarms and to waste (primarily human) resources at loyal peers. A measure of adversary success is the time it takes for the adversary to raise the first alarm (the lower, the better for the adversary).
- Attrition.* The adversary wishes to prevent loyal peers from repairing damage to their replicas caused by naturally occurring failures. Toward that goal, she wastes the computational resources of loyal peers so that they cannot successfully call polls to audit and repair their replicas. Measures of her success are the time between successful polls called by loyal peers or the busyness of loyal peers (in both cases the higher, the better for the adversary).
- Theft.* The adversary wishes to obtain published content without the consent of the publisher. For example, she wishes to obtain for-fee content without paying the fee. A measure of adversary success is the time to obtain a copy of the restricted content (the lower, the better for the adversary).
- Free-loading.* The adversary wishes to obtain services without supplying services to other peers in return. For example, she wishes to obtain repairs for her own replicas, without supplying repairs to those who request them from her. A measure of her success is the ratio of repairs supplied to repairs received (the lower, the better for the adversary).

In this article, we focus on adversaries with the stealth modification, nuisance, and attrition goals. We choose these three goals because they have the greatest potential, especially in concert, to disrupt the preservation work of a community of LOCKSS peers, since they can directly or indirectly hurt the preserved content *and* discredit the functions of the system that can detect the damage. In the rest of this section, we briefly describe what the LOCKSS system can do about theft and free-loading.

The LOCKSS system does not materially increase the risk of theft. Repairs are the only protocol exchanges in which content is transferred. A peer supplies a repair only if the requester has previously proved with an agreeing vote that it once had the same content (see Section 4.1.12). The protocol cannot be used to obtain a first instance of an AU. Most e-journals authorize ranges of IP addresses for institutional subscribers; existing LOCKSS peers use this mechanism to authorize their crawls for first instances of content. If more secure forms of content authorization (e.g., Shibboleth [Erdos and Cantor 2002]) become widely accepted, LOCKSS peers can use them.

The repair mechanism also limits the problem of free-loading. The primary good exchanged in the LOCKSS system, content repair, is only available to peers who prove through voting that they have had the content in the past. A peer who votes dutifully—so as to be able to obtain repairs if necessary—but does not itself supply repairs to others can hardly be described as *free-loading*, given the cost of the voting required. Tit-for-tat refusal to supply repairs to such antisocial peers might further deter this behavior.

In the next section, we present attack techniques available to the adversary. We combine these techniques in later sections, to produce adversary strategies toward the stealth modification, nuisance, and attrition goals.

5.3 Attack Analysis by Peer Role

We present a different categorization of the attacks that divides them according to the role played in the poll by the malign peer. For each role, we examine attacks that affect the result of the poll, and those that degrade the quality of the poller's reference list.

We also describe additional attacks that deny service, and some spoofing attacks.

5.3.1 Poller. The worst a malign poller can do is to deny service.

- Affect poll result.* A malign poller can affect the result of a poll it initiates, but to no effect on the rest of the system, since it is already malign.
- Degrade reference list.* A malign poller can degrade its own reference list, but to no effect on the rest of the system, since it is already malign.
- Deny service.* The attrition adversary acts as a malign poller and invites large numbers of loyal peers into polls to waste their resources and prevent them from coming to consensus. We set the poll effort cost to make this attack very expensive (Section 4.2.1), and raise an interpoll interval alarm if we detect it.

5.3.2 *Inner Circle.* A malign peer invited into the inner circle of a loyal peer's poll can take part in all three kinds of attacks.

- Affect poll result.* If enough malign peers are invited into a poll they can affect the result without detection (Section 5.5.1). The nuisance adversary attacks a poll if it meets the criteria discussed in Section 5.5.2.
- Degrade reference list.* A malign inner circle invitee has the opportunity to degrade the quality of the poller's reference list, by recommending other malign peers into the poll's outer circle. Malign peers recommended for the outer circle will only get into the loyal poller's reference list if they expend effort and vote with the consensus.

However, the malign peers have an advantage over the loyal peers. They know the identities of the other malign peers. Loyal peers recommend some malign and some loyal peers, but malign peers recommend only malign peers. Doing so costs the adversary effort over time because the malign invitees have to exert effort to get in and stay in the loyal peer's reference list. Loyal peers also have to exert effort to stay in other loyal peers' lists. Malign peers do not have to exert effort to stay in their co-conspirator's lists.

- Deny service.* A malign inner circle peer can do three things to deny service, none of which is effective:
 - It can cheaply generate a bogus vote in a poll the malign peers are certain to lose. The poller can eliminate the vote with little effort and can then remove the malign peer from its reference list. This is not a good strategy.
 - It can recommend loyal peers into the outer circle then discredit the poll by inviting the same loyal peers into a bogus poll with the same ID. The effort of creating the bogus poll is larger than the cost to these loyal peers of detecting that it is bogus. The adversary would benefit more from using the effort to degrade the poller's reference list further.
 - It can refuse to send a vote, raising the probability that the poll fails to get a quorum, at the risk of losing its slot in the poller's reference list if the poll does achieve a quorum. Unless the adversary is resource-constrained, the malign peer should vote and recommend other malign peers.

5.3.3 *Outer Circle.* A malign peer invited into the outer circle of a poll is offered a chance to occupy a place in the loyal poller's reference list. This is a valuable opportunity to attack a future vulnerable poll.

- Affect poll result.* A malign outer circle invitee cannot change the result of the poll because the poller only uses outer circle votes for discovery, not for auditing.
- Degrade reference list.* To degrade the reference list, the malign outer circle invitee must vote with the consensus. It must then continue to act as a healthy peer in future polls until the malign peers decide to attack from within the inner circle.
- Deny service.* A malign outer circle invitee can also remain mute, causing a shortage of outer circle voters, thereby impeding the expansion of the poller's

reference list. However, this only ensures that the adversary's presence in that poller's reference list decreases.

5.4 Attack Techniques

In this section, we identify possible attack vectors against the LOCKSS opinion poll protocol, and we describe how they can be exploited.

5.4.1 Adversary Foothold in a Reference List. The composition of a loyal peer's reference list is of great importance. A loyal peer chooses inner circle voters for a poll that it initiates from its reference list at random (see Section 4.1.2); as a result, the proportion of malign peers in the inner circle approximates their proportion in the poller's reference list. An adversary wishing to control the outcome of a poll initiated by a loyal peer can increase her chances by increasing the proportion of malign peers in that peer's reference list.

To gain an initial foothold in loyal peers' reference lists, the adversary must take over peers that used to be loyal. She does this, for example, by exploiting common implementation vulnerabilities, or by coercing peer operators to act as she wishes. She can then increase that foothold, but to do so she must wait until a loyal peer invites a malign peer into the inner circle of a poll. When invited, the adversary causes the malign peer to nominate other malign peers unknown to the poller. Note that loyal peers also nominate malign peers inadvertently. This means that, assuming the average foothold in loyal peers' reference lists is B , then after a poll that contains M malign inner circle voters and L loyal inner circle voters, a fraction $(M + L \times B)/(M + L)$ of the outer circle nominations are for malign peers, which is greater than the current foothold ratio B regardless of the value of B .

Malign peers in loyal peers' reference lists must behave as loyal peers until an attack requires otherwise. Each such peer thus consumes adversary resources as it must both vote in and call polls to avoid detection and maintain its position in the list (Section 4.2.2).

We measure the adversary's success at maintaining a foothold in loyal peers reference lists with the average *foothold ratio* over the population of loyal peers. For a given reference list, the foothold ratio is the fraction of the list occupied by malign peers.

The LOCKSS protocol has two lines of defense against reference list takeover. First, loyal peers only change their reference lists after a poll that they call; the adversary must wait until they do so before she can increase her foothold (see Section 4.2.3). Second, friend bias allows the operator of a loyal peer to trade off the risks of depending too much on a static set of friendly peers against those of depending too much on peers it discovers via the outer circle nomination process (Section 4.2.4).

5.4.2 Delayed Commitment. The adversary need not decide in advance how each of her malign peers will react to particular poll invitations. Instead she can determine how a particular malign peer behaves after having collected all available information about a poll or its poller, as per her Total Information Awareness and Stealth capabilities (Section 5.1).

Loyal peers can defend against this adaptive adversary by requesting commitments on future protocol steps as early as possible. During repairs, the requester checks that the repair it receives is consistent with the vote of the repair supplier. During voting, the poller could request from potential voters an early commitment on the hash of a few bytes of the AU, chosen randomly. Later, the poller could verify that each vote is consistent with the AU version to which that voter committed. This would reduce the adversary's ability to attack only polls she is sure of winning and increase the probability of detection. Our simulations currently use the former defense but not the latter.

5.4.3 Peer Profiling. Using the Local Eavesdropping capability, the adversary can observe a loyal peer's traffic and attempt to infer useful information such as likely members of the peer's reference list, likely voters in a poll, and whether or not the peer has agreed to vote in a poll. The adversary can use this information to make better decisions, as with delayed commitment, or to mount flooding attacks against invited loyal peers whose sessions she can hijack (Section 5.4.4).

Loyal peers defend against eavesdroppers by actively obfuscating their protocol state (Section 4.2.5).

5.4.4 Session Hijacking. The LOCKSS system lacks stable identities because it cannot support them without long-term secrets. An attacker can thus impersonate loyal peers and hijack their sessions, using her Local Spoofing capability (Section 5.1). By spoofing local source addresses in messages it sends to remote peers, or remote source addresses in messages it sends to local peers, a malign peer within spoofing range of a poller can affect the poll result by either hijacking sessions between the poller and loyal invitees, or discrediting loyal invitees of the adversary's choosing.

The malign peer can hijack a session by responding to the poller's Poll message with a spoofed PollChallenge message establishing a session key. If the poller also receives the genuine PollChallenge message, the two conflict, the invitee is discredited, and the hijack fails. If, however, the loyal invitee fails to respond with a PollChallenge or the adversary manages to suppress it, perhaps by flooding the loyal invitee's link, the hijack can succeed and the malign peer can vote in the loyal invitee's place. Once established, a session is protected by a session key and can no longer be hijacked.

Alternatively, the malign peer local to the poller can selectively discredit loyal invitees by also responding with a spoofed, conflicting PollChallenge. Votes sent from invitees whose PollChallenge message has thus been discredited are not tallied, as the poller lacks a means to distinguish between the spoof and the original. By suppressing loyal votes in this way the adversary can increase her foothold ratio (Section 5.4.1), or waste human resources by raising spoofing alarms (Section 4.2.6).

The operators of a loyal peer can defend against hijacking by checking its network vicinity for subverted peers or routers, by providing it with a dedicated subnet, and by monitoring for packets spoofing the router's MAC address. Loyal peers could retransmit their potentially suppressed PollChallenge messages at random intervals throughout the poll. If any of these retransmissions get to the

poller, the hijacked session is discredited. This would force the adversary to suppress traffic from the hijacked peer for many hours, increasing the probability of detection.

5.5 Attack Strategies

In this section, we take a closer look at three kinds of adversaries: first, an adversary whose goal is to modify an AU stealthily across loyal peers so as to change the scientific record persistently (Section 5.5.1); second, an adversary whose goal is to discredit the LOCKSS alarm mechanism, thereby rendering it incapable of reacting to attacks (Section 5.5.2); and third, an adversary whose goal is to slow down LOCKSS polls long enough for random faults to cause irrecoverable damage in the content (Section 5.5.3). For all three types of adversaries, we describe an attack strategy. In Section 6, we measure through simulation how the LOCKSS system fares against such attacks.

5.5.1 Stealth Modification Strategy. The stealth adversary has to balance two goals: changing the consensus on the target AU and remaining undetected. To achieve these goals, he must repeatedly find a healthy poller, convince it that it has a damaged AU replica without causing it to raise any alarms, and then, if asked, conveniently offer it a repair with the bad version of the AU. This strategy relies primarily on building a foothold in loyal peers' reference lists (Section 5.4.1) and on delayed commitment (Section 5.4.2).

The stealth adversary acts in two phases. First he *lurks*, seeking to build a foothold in loyal peers' reference lists but otherwise behaving as a loyal peer, voting and repairing with the correct version of the AU. Then he *attacks*, causing his malign peers to vote and repair using either the correct or the bad version of the AU, as needed. During the attack phase malign peers vote with the correct copy unless a poll is *vulnerable*, i.e., the overwhelming majority of the poll's inner circle is malign. In vulnerable polls malign peers vote with the bad copy, because by doing so they can change the loyal poller's AU without detection. Polls are vulnerable if the following three conditions hold:

$$M + L \geq Q, \quad (1)$$

$$M > L, \quad (2)$$

$$L \leq D. \quad (3)$$

Condition 1 ensures that the $V = M + L$ peers agreeing to vote satisfy the quorum Q . Condition 2 ensures that the M malign peers determine the result with an absolute majority of the votes. Condition 3 ensures that the L loyal peers are not enough to raise an inconclusive poll alarm at the poller. The adversary we model has Complete Parameter Knowledge (Section 5.1) and can evaluate this vulnerability criterion exactly, in accordance to our "strong adversary" design principle. In a practical system an adversary would have only estimates of L , Q , and D , and would thus run a higher risk of detection than in our simulations.

The protocol provides several defenses that are especially relevant against the stealth adversary. Individually none is very strong; in combination they are quite effective. First, in accord with our rate-limiting principle, the adversary cannot induce loyal peers to call vulnerable polls but has to wait until they occur.

Second, a damaged peer continues to call and vote in polls using its now bad copy of the AU. Unlike malign peers, it does not evaluate the vulnerability criterion or decide between the good and bad versions. If more than D damaged peers take part in a poll called by a healthy peer but the adversary deems the poll invulnerable, an inconclusive poll alarm is raised.

Third, a damaged peer continues to call polls and may invite enough healthy peers into its inner circle to repair the damage. For each loyal peer the stealth adversary damages, he must expend resources to maintain his foothold in the peer's reference list and vote whenever his malign peers are invited into polls, until the bad version of the AU prevails everywhere.

Finally, if the stealth adversary fools the poller of a vulnerable poll into requesting a repair, he must ensure that the request will go to one of the malign peers. The poller requests repairs only from peers in whose polls it has voted; others would refuse the request as they lack evidence that the requester once had a valid copy (Section 4.1.12). Thus the stealth adversary must expend effort to call polls as well as vote in polls called by the loyal peers.

5.5.2 Nuisance Strategy. The nuisance adversary has a simple goal: to raise alarms at loyal peers as fast as possible. The nuisance adversary can cause all three types of alarms (Section 4.2.6).

To cause an inconclusive poll alarm, the adversary can use delayed commitment. Every time some of her malign peers are invited into a poll, she evaluates the vulnerability criterion

$$M + L \geq Q, \quad (4)$$

$$L > D, \quad (5)$$

$$M > D. \quad (6)$$

Apart from reaching a quorum (Condition 4), the criterion means that neither are the loyal votes few enough to lose quietly (Condition 5), nor are the malign votes few enough to allow a landslide win by the loyal votes (Condition 6). When the adversary detects that this criterion is met, she instructs her malign peers to vote in the poll with a bogus AU, thereby causing the poll to raise an inconclusive alarm. If the criterion is not met, the adversary instructs her malign peers to vote with the correct AU. This is necessary, so that malign peers can remain in the reference list of the poller (see Section 4.1.6) for another try in the future; the adversary can foster the satisfaction of the vulnerability criterion by building a foothold in the reference list of the poller (Section 5.4.1).

The adversary can cause a spoofing alarm by sending conflicting PollChallenge messages to the poller. This strategy can yield spoofing alarms when the adversary has more than A malign peers in the invited inner circle of a poll. Then, without having a spoofer near the poller, the adversary can cause the poller to suspect a local spoofer and to raise frivolous alarms, which waste human resources at the poller.

The adversary can cause an interpoll interval alarm by foiling loyal peers' attempts to call a poll successfully for long enough. This strategy is similar to a strategy with the goal of attrition (Section 5.5.3).

The nuisance adversary need not have her malign peers call polls, because she does not seek to obtain or supply repairs. However, she needs to have her malign peers invited in polls. As a result, she must follow the observable aspects of the protocol (voting with valid votes when asked) unless she can attack a particular poll.

The LOCKSS system defends against the nuisance adversary primarily via the autonomy with which loyal peers decide when to call polls. As a result, a nuisance adversary must persist for a while, increasing her foothold in the reference lists of loyal peers, before she can attack a particular poll raising an inconclusive poll alarm.

5.5.3 Attrition Strategy. The attrition adversary’s goal is to occupy the time the loyal peers have available for voting, making it less likely that a poll called by a loyal peer gains a quorum of voters. Success is measured by the average time between successful, quorate polls at loyal peers. If the attrition adversary can increase this interval enough, random damage at the loyal peers can accumulate and degrade the system.

The attrition adversary’s strategy is to call polls as fast as possible, inviting only loyal peers. The adversary’s peers do not vote in polls called by other peers; there is no need to persuade loyal peers to fetch repairs from them. We do not yet use a “newcomer pays extra” strategy [Friedman and Resnick 2001], so the attrition adversary can currently use one-time throw-away identities to call the polls.

The attrition adversary’s impact on the system is currently limited only by the rate at which he can compute the proofs of effort demanded by the loyal voters he is trying to involve in his polls. We are investigating techniques that limit his impact more effectively (see Section 9).

6. SIMULATION

We have evaluated our new protocol’s resistance to random failures and malicious attacks using the simulation we present in this section. We first describe our simulation environment in Section 6.1. Then we explain how we simulate loyal peers (Section 6.2). Then in Section 6.3 we describe how we simulate the different adversary strategies from Section 5.5. Section 7 collects our simulation results.

6.1 Simulation Environment

In this section, we describe the simulation environment we use for our evaluation. This includes our simulator, the network model we employ, and our application-layer overlay topology initialization.

We use Narses, a Java-based discrete-event simulator [Giuli and Baker 2002] designed for scalability over large numbers of nodes, large amounts of traffic, and long periods of time. Narses offers facilities for a variety of flow-based network models allowing tradeoffs between speed and accuracy. The simulator can also model expensive computations, such as hashes and proofs of effort, allowing some realism in our simulation of protocols involving cryptographic primitives.

Since we simulate a LOCKSS network for up to 30 (simulated) years, we use a faster-to-simulate network model that considers propagation latency but not traffic congestion. We simulate the underlying network topology as a star at the center of which lies the “core.” Individual nodes are linked to the core via a link whose bandwidth is chosen at random among 1.5, 10, and 100 Mbps, and whose propagation latency is chosen uniformly at random from 1 to 30 ms. The core has infinite switching capacity; as a result, the effective bandwidth of a flow from node A to node B is the minimum bandwidth of the two links, and its propagation latency is the sum of the propagation latencies of the two links.

Every simulation run starts with an initial population of 1000 peers, each storing an AU that takes 120 s to hash, which is the equivalent of a 4-GB AU for our deployed machines. The dynamic contents of the reference lists of these peers determine the application-layer topology of the LOCKSS overlay. As the protocol requires, the reference list of each peer is initialized with the content of its friends list. We initialize each peer’s friends list with a clustering technique. Peers are randomly assigned to clusters of 30 peers. For each peer, we add 29 other peers to its friends list, 80% of which are chosen randomly from its own cluster and the rest chosen randomly from other clusters.

We simulate a provable effort mechanism similar to the MBF scheme devised by Dwork et al. [2003]. In keeping with the constraints placed by that scheme and with the requirements we set out in Section 4.2.1, we derive one possible set of provable effort sizes for the protocol (Appendix A.3). Given that hashing the AU costs S , the poll effort construction size (Section 4.1.3) is $(20/3)S$, the verification of a poll effort proof (Section 4.1.10) costs $(5/3)S$, the cost of computing a valid vote (Section 4.1.11) is $5S$, and the cost of verifying a vote (Section 4.1.5) is $2S$ for agreeing and between S and $2S$ for disagreeing votes (depending on where in the AU the disagreement occurs first).

If the cost of hashing the AU is 120 s, the poller spends 800 s per invitee generating the PollProof message and 240 s per invitee verifying an agreeing Vote message. Each invitee spends 200 s verifying the PollProof message and 600 s generating the Vote message. An entire successfully concluded poll without repairs costs the poller 1040 s of computation per invitee. With 20 invitees it would take 6 h, which is comparable to the duration of polls in the beta deployment.

6.2 Simulated Loyal LOCKSS Peers

We simulate loyal LOCKSS peers as simple state machines implementing the protocol of Section 4. We set the protocol parameters (see Section 4.1) to values reflecting those in the 60-peer tests of the deployed system ($N = 20$, $Q = 10$, $D = 3$, $A = 3$, $I = 10$, $E = 4$ polls), except $R = 3$ months which we estimate reflects production use.

We set protocol timers to be just long enough for the slowest machine to complete the corresponding protocol step. Peers always consider themselves the fastest. For example, peers who have been invited into a poll give the poller enough time to compute the poll effort proof for N invitees (see Section 4.1.9), assuming that the poller has a memory system five times slower than theirs [Dwork et al. 2003].

Our simulated peers commit to a poll exclusively, for the duration of that poll, even when idly waiting for a protocol message to come back. However, a peer that wishes to call its own poll but is also invited in another poll called by someone else prefers to call its own poll.

All loyal peers in a simulation run have the same nominal rate of random undetected errors that unobtrusively replace the victim's AU replica with random bits.

6.3 Simulated Adversary

In this section we address our simulation of LOCKSS adversaries. We outline how we represent an adversary and his malign peers, and then we describe how we implement within our simulation environment the attack techniques available to him (see Section 5.4).

We simulate an adversary as a multihomed node with as many network interfaces (NICs) as the number of IP addresses, and as many CPUs as the number of nodes controlled by the adversary (i.e., one *humongous* computer). The numbers of NICs and CPUs are parameters of the simulation. An adversary with few NICs and many CPUs has a lot of processing power at his disposal, but is without a great presence in the network. An adversary with many NICs and fewer CPUs has some processing power but a lot of scope for spoofing IP addresses.

To gain a foothold in loyal peers' initial reference lists (Section 5.4.1), the adversary may use his ability to take over some of the LOCKSS peer population (see Section 5.1). We initialize simulation runs at the instant when the takeover is complete. For example, to run a simulation where the adversary subverts 30% of the 1000 peers, but also has 100 extra CPUs at his disposal, we simulate an adversary with $1000 \times 30\% + 100 = 400$ CPUs and only 700 loyal peers.

In our simulations, once the adversary receives a PollProof message via one of his NICs, he considers the number of those NICs via which he has received Poll messages thus far to be M for this poll (see Section 4.1). Then, the adversary divides the list of his own NICs among the M malign peers as which he participates in the inner circle of the poll. When a particular malign peer NIC receives a PollProof message, the adversary waits the appropriate time for the verification of the poll effort proof and then responds with a Nominate message holding the corresponding portion of the list of malign NICs or, if that portion is too large, a random sample of that portion. The adversary thus ensures that the loyal poller inserts into its outer circle the maximum number of distinct malign peer addresses; the adversary cannot do better, without knowledge of the nominations of other loyal inner circle peers.

We simulate delayed commitment (Section 5.4.2) by waiting until the adversary must start computing his first vote before deciding on which AU that first vote will be. At that time, the adversary evaluates the appropriate poll vulnerability criterion, according to the strategy we simulate, decides whether to attack the poll and how, and commits to the appropriate version of the AU.

We simulate the adversary's ability to profile loyal peers (Section 5.4.3) by making all variable protocol parameters known to him. We do not, in this article, otherwise simulate the presence of eavesdroppers near loyal peers.

Finally, we do not simulate in this work the adversary's ability to hijack poll sessions between loyal pollers and their loyal invitees (Section 5.4.4).

6.3.1 Simulated Stealth Modification Adversary. We simulate the effects of an attack by an adversary following the stealth modification strategy (Section 5.5.1) in two sets: *lurking* and *attacking* simulations, corresponding to the lurking and attack phases of the strategy. In lurking simulations, the adversary seeks only to extend his foothold in loyal peers' reference lists. After initially subverting some of the loyal peers, the adversary has malign peers behave exactly as loyal peers do, except for forming their Nominate messages as described in Section 6.3. Lurking simulations last 20 simulated years.

In attacking simulations, malign peers seek not only to extend their foothold in loyal peers' reference lists, but also to change the loyal peers' replicas of the AU with the bad version that the adversary wishes to install throughout the community. Therefore, malign peers also evaluate the vulnerability criterion and decide, as described in the previous section, on which AU to base their votes and their repairs. We initialize the population in an attacking simulation as if a lurking phase preceded the simulation: we initialize the reference lists of loyal peers with a given foothold ratio. Attacking simulations last 10 simulated years, unless an inconclusive poll alarm is raised.

To draw conclusions about entire stealth modification attacks, we must combine the results of a lurking simulation with the results of a *compatible* attacking simulation. We accomplish this by first running a set of lurking simulations for the set of input parameters we seek to study. Based on these runs, we identify how great a foothold ratio the adversary can obtain, for given input parameters. Then we run a set of attacking simulations with input parameters that match the input parameters of the lurking simulations as well as the observed possible foothold ratios gained by the adversary. For example, when studying the stealth modification adversary who begins by subverting 20% of the 1000 initially loyal peers, we run a number of lurking simulations (for different random seeds), from which we conclude that the adversary can obtain average foothold ratios of 40 to 55%. Based on this, we only run attacking simulations for 20% subversion of the 1000 initially loyal peers with initial attack-phase foothold ratios that range between 40 and 55%.

Splitting the strategy into two sets of simulations allows us to explore the choice the adversary makes about the foothold he must achieve before switching from lurking to attacking. In our results, we assign the first possible time at which this foothold is achieved for a given initial subversion as the duration of the lurk phase for that subversion.

6.3.2 Simulated Nuisance Adversary. To simulate the effects of an attack by an adversary following the nuisance strategy (Section 5.5.2) we run simulations similar to the attacking simulations of the stealth modification adversary (Section 6.3.1).

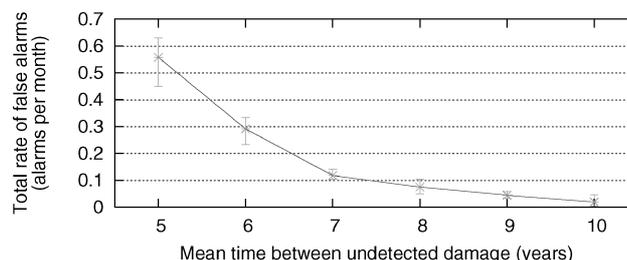


Fig. 2. The total rate of false alarms versus the mean time between random undetected damage at each peer, with no adversary.

Nuisance simulations differ from attacking stealth simulations in two ways. First, the nuisance adversary does not lurk but rather attacks any vulnerable polls immediately. Second, she uses the weaker vulnerability criterion of this strategy.

6.3.3 Simulated Attrition Adversary. We simulate the attrition adversary with unlimited identities but limited resources. He calls useless polls to consume the loyal peers' resources. Unlike the stealth adversary, he neither lurks to degrade reference lists nor attacks polls. Although loyal peers give priority to calling their own polls, those polls become increasingly inoperative when the general busyness in the system increases. A loyal peer raises an interpoll interval alarm if it has not completed a poll in three times the expected interpoll interval.

7. RESULTS

In this section we evaluate the new LOCKSS opinion poll protocol through simulation. We explore how the protocol deals with random storage faults, as well as attacks. We demonstrate the following points:

- Absent an attack, substantial rates of random damage at peers result in low rates of false alarms (Section 7.1).
- The stealth adversary fails when loyal peers bias their reference lists to be composed of one-third friends and the adversary subverts less than one-third of the population. As friend bias decreases or subversion increases, the adversary's chances of causing irrecoverable damage improve gradually (Section 7.2).
- A nuisance adversary whose goal is simply to raise an alarm has to exert significant effort over a long period (Section 7.3).
- An attrition adversary whose goal is to prevent consensus long enough for random damage to corrupt the AU must invest nonnegligible resources in his attack (Section 7.4). Unfortunately, the barrier to entry is relatively low. More work is required against attrition adversaries.

7.1 Rate of False Positives

Without an adversary, but with peers subject to random damage they do not themselves detect, Figure 2 shows that false alarms occur rarely. We simulate

20 years with *every* peer suffering random undetected damage at mean intervals varying from 5 to 10 years. Over 20 runs, we show the minimum, average, and maximum total rates of false alarms raised at *any* peer in the entire system. With undetected damage at each peer every 5 years, in the worst case the average rate of false alarms in the system is 44 days, that is, every 44 days some peer in the system sees an alarm. The average peer sees an alarm once in about 120 years.

The rates of random undetected damage we simulate are vastly higher than we observe in practice. Our peers typically lack reliability features such as ECC memory. Yet in over 200 machine-years of the test deployment, we have observed only one peer in which such errors affected polls. Our simulations that follow assume this 1 in 200 probability of a random undetected error per peer year.

7.2 Stealth Adversary

We show that the probability of a stealth adversary causing irrecoverable damage remains very low even for an initial subversion of one-third, and then increases gradually. Conservatively, we deem damage irrecoverable if the initially subverted (malign) and the damaged (loyal) peers form more than 50% of the population. For the following simulations, the adversary has infinite CPUs and as many NICs as necessary to gain the maximum possible foothold ratio during 20 years of lurking.³ We vary friend bias from 10% to 40% and subversion from 2% to 40%. Because the definition and term for friend bias in the reference list has changed since the conference version of this article, we give here a quick equivalence of the two: 2% churn corresponds to friend bias of 8%, and 10% churn corresponds to friend bias of 40%.

For every initial subversion and friend bias level, we run all compatible attack phases lasting up to 10 years for all foothold ratios (20% and up) achieved during the lurking runs (see Section 6.3.1). We run every combination of parameters described above with 20 different random seeds.

Figure 3 shows the minimum time taken by the lurking phase to a foothold for different levels of friend bias. The x axis shows the proportion of total peers that are initially subverted. The y axis shows the minimum time it takes the adversary to deteriorate the loyal peers' reference lists to various foothold ratios. The adversary can achieve greater footholds when loyal peers use lower friend bias. This is because reference list biasing resists the attempts of the adversary to create a foothold in reference lists that is much greater than the initial subversion he has obtained; less friend bias means that the adversary can expand his foothold more effectively using additional identifiers. Note that in our simulations, friends are as likely as other peers to be subverted by the initial subversion.

Figures 4 and 5 focus on our defenses against the stealth adversary from the point of view of two different adversary subgoals: time to inconclusive alarm

³Please note that simulations in an earlier version of this work [Maniatis et al. 2003] contained an error that, at very low levels of subversion, did not allow the adversary as many NICs as he could exploit. This caused lurking to take him longer.

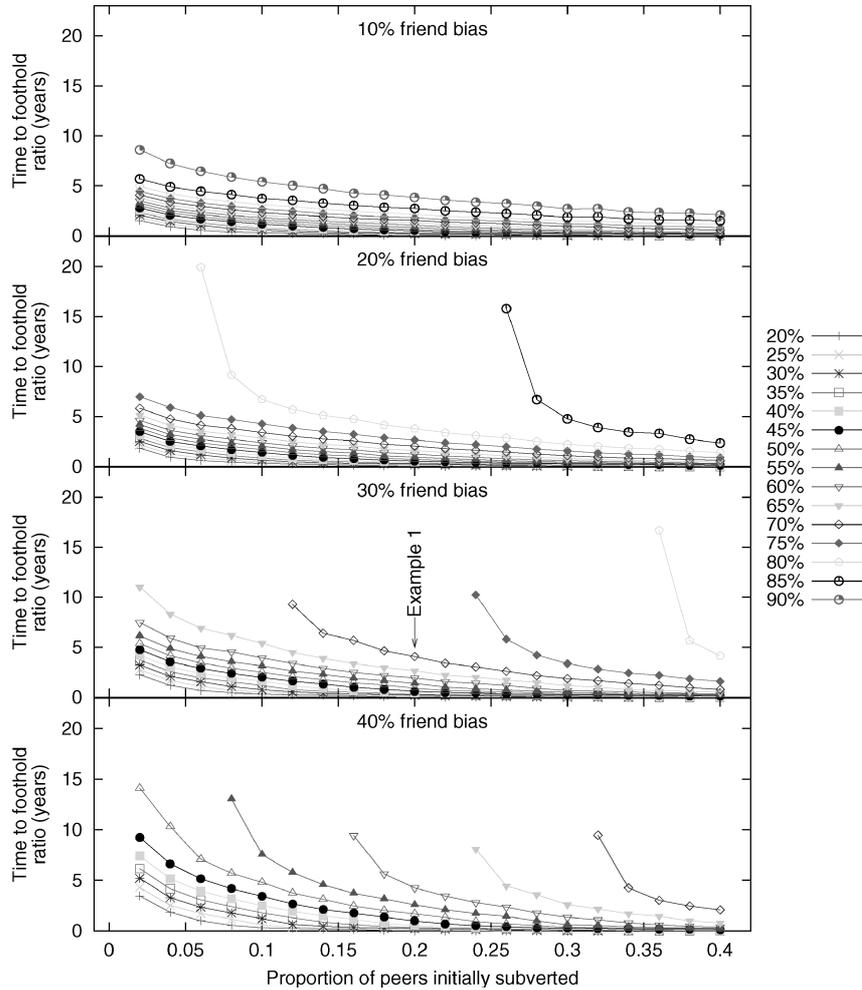


Fig. 3. Minimum time for the stealth adversary to achieve various foothold ratios, starting with various proportions of initially subverted peers. We show graphs for 10, 20, 30, and 40% friend bias.

and damage achieved, both of which the adversary wishes to maximize. These two adversary subgoals are, however, conflicting since the greater the damage inflicted on loyal peers' AU replicas, the greater the chance an inconclusive alarm may be raised, and conversely, avoiding inconclusive alarms may mean limiting damage inflicted on loyal peers' AU replicas. Figure 6 combines the two subgoals for the adversary, maximizing the access failure he can inflict. We now describe these graphs in detail.

Figure 4 shows how long the attack phase lasts before it is detected. For each foothold ratio at the beginning of the attack phase, we show the quartile distribution of times until the first alarm. Some runs do not raise an inconclusive poll alarm; they damage very few loyal peers. At the top of each distribution is the percentage of such runs.

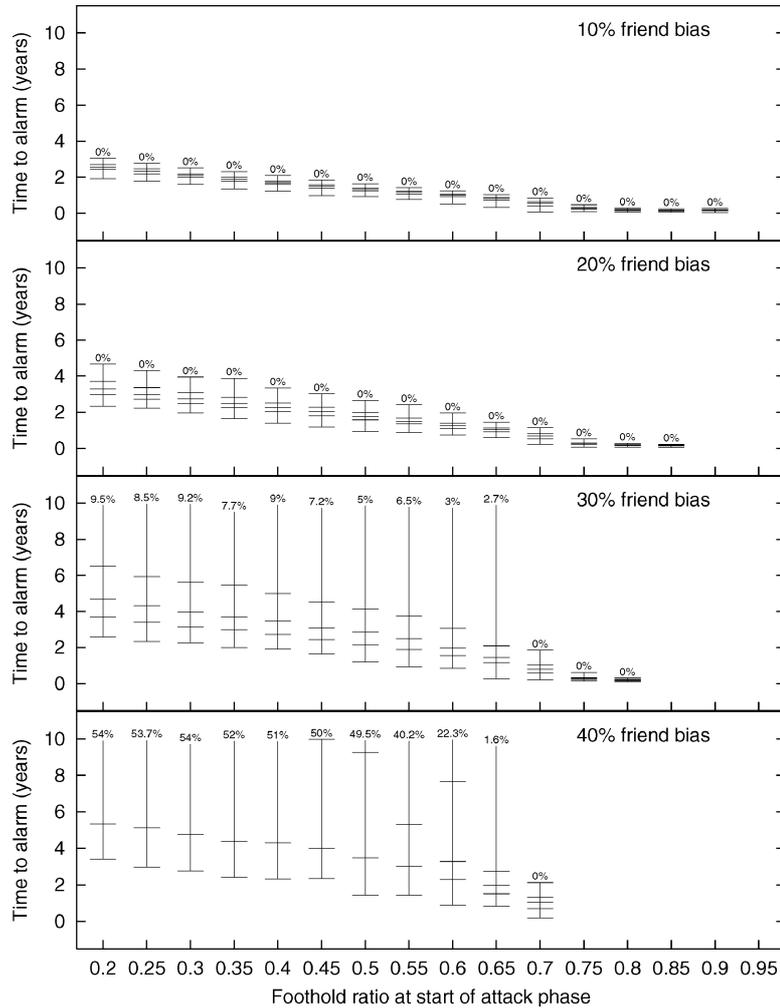


Fig. 4. The time from the start of the attack phase (in the stealth strategy) to the time of detection for different starting reference list foothold ratios. Ticks split the value distributions into quartiles. Percentages above the distributions indicate runs that did not generate an alarm. We show graphs for 10, 20, 30, and 40% friend bias.

Figure 5 illustrates that the adversary must subvert a greater number of loyal peers to change the AU irreversibly as the friend bias increases; friend bias of 10% permits irrecoverable damage for all levels of subversion we are showing, whereas friend bias of 40% permits no irrecoverable damage even when 40% of the population is subverted. The graph shows the distribution of the maximum proportion of bad replicas (including those at subverted peers) caused as a function of initial subversion. The number above the distribution shows the time in years needed to achieve the maximum damage with that subversion. For example, consider a scenario in which the friend bias is 30% and the adversary initially subverts 20% of the population (marked “Example 1”

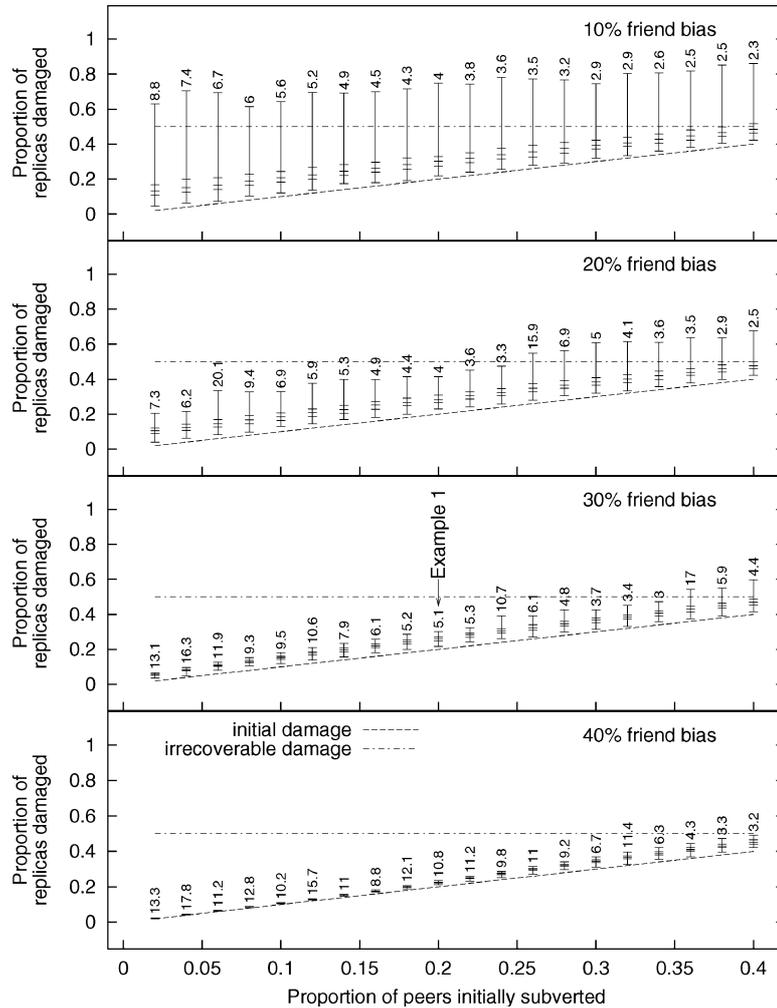


Fig. 5. The percentage of bad replicas as a function of initial subversion. Ticks split the value distributions into quartiles. Numbers above the distributions show the time in years needed by the adversary to cause maximum damage. The diagonal line shows the damage due to peer subversion. The horizontal line shows the threshold for irrecoverable damage.

in Figures 3 and 5). To inflict the greatest damage, the adversary must lurk until he achieves a 70% foothold ratio, which takes him at best a bit more than 4 years, and then must attack for longer than another year to damage an additional 10% of the replicas beyond those replicas at subverted peers; from start to finish, this best case for the adversary in this particular scenario takes 5.1 years. As subversion increases, the adversary is able to damage more loyal peers. With a 30% friend bias, the adversary does not cause irrecoverable damage with less than 36% initial subversion; beyond that level of subversion, he causes irrecoverable damage in no more than 13% of the runs in any particular subversion scenario (see the 30% friend bias curve of Figure 8).

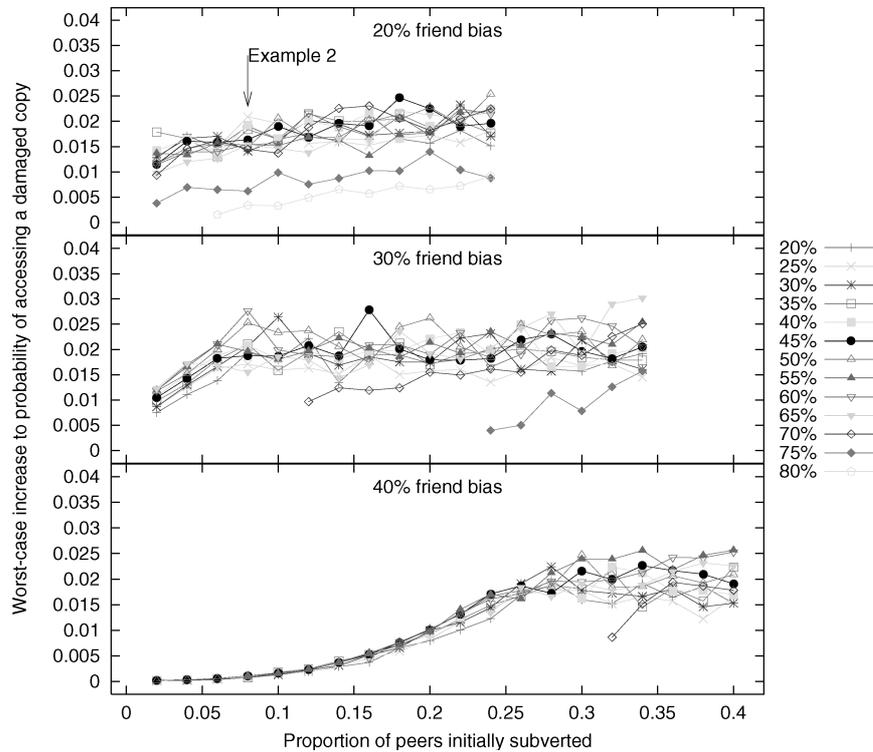


Fig. 6. Worst-case increase in the probability of accessing a bad replica due to the attack as a function of initial subversion. Each curve corresponds to runs with a different foothold ratio at the start of the attack phase. We only show values for subversions at which the adversary did not cause irrecoverable damage. We show graphs for 20, 30, and 40% friend bias.

Figure 6 summarizes the effect on readers of attacks, isolating the benefit that the adversary obtains with his strategy from what he is given through the initial subversion. We omit subversions for which irrecoverable damage is possible; in those cases, the transient effect of the attack on readers is irrelevant compared to the permanent loss of content that the adversary causes. These excluded cases are included in Figure 8.

On the x axis, we show the initial subversion. On the y axis, we show the worst-case probability (due to the attack) that a reader of the AU finds a damaged copy, that is, the expected value of the maximum fraction of damaged (not initially subverted) AUs during the lurk and attack phases. We graph a curve for each foothold at which the adversary starts the attack phase, and we show a graph for friend bias of 20, 30, and 40%; at friend bias of 10%, the adversary always has a chance of causing irrecoverable damage. Each point in every curve represents the worst possible case over many different simulation runs for that foothold and initial subversion, so it is not surprising that the curves are not smooth and frequently intersect each other. Interestingly, the adversary's best strategy is not to lurk for as long as possible: readers are most likely to see a bad AU when he lurks up to a foothold ratio of 40 to 65% at a friend bias of

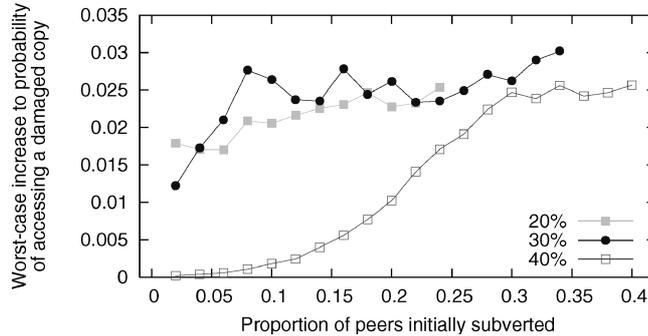


Fig. 7. Worst-case increase in the probability of accessing a bad replica due to the attack as a function of initial peer subversion. Each curve corresponds to runs with a different level of friend bias. We only show values for subversions at which the adversary did not cause irrecoverable damage for each friend bias level.

30%. Lurking less results in weaker attack phases. Lurking more means that peers supply readers with good AUs for longer; for instance, lurking up to a foothold ratio of 75 or 80% is consistently the worst choice when friend bias is 20%. This behavior is also consistent for the mean probability (not shown). The “best” foothold ratio at which to start attacking is lower for lower levels of friend bias; for example, there is a scenario (8% initial subversion, marked “Example 2” in the figure) for which starting the attack at a foothold ratio of 25% is best when the friend bias is 20%. This is because the system’s resistance to the adversary’s attempts to gain a foothold in reference lists is weaker when friend bias is lower, that is, when fewer friends are brought into the reference list to counter the adversary’s biased statistics.

Note also that for friend bias of 20% or more, despite an initial subversion of up to 40% of the peers by an adversary with unlimited computational power, unlimited identities, complete knowledge of the protocol parameters, and an attack lasting up to decades, the probability of a reader accessing a bad AU is never more than 3 percentage points greater than it is immediately after the initial subversion. The system resists further damage effectively despite the subversion of more than one-third of its peers.

In Figure 7, we summarize how different levels of friend bias affect the worst-case probability of accessing a bad AU over all choices of foothold ratio at the start of the attack phase. Thus the curve for each friend bias level represents the upper envelope of the corresponding subfigure of Figure 6. Again, we only show data for subversions at which irrecoverable damage does not occur. The graph also shows that increasing the friend bias raises the initial subversion the adversary needs before he can cause irrecoverable damage—recall that the line ends where irrecoverable damage begins—and reduces the probability of accessing a bad AU replica.

Figure 8 shows that biasing the reference list toward friends is an invaluable tool in thwarting the adversary. On the y axis, we show the probability that the adversary causes irrecoverable damage in our simulations, given different initial subversions on the x axis, one curve per friend bias level. Increasing

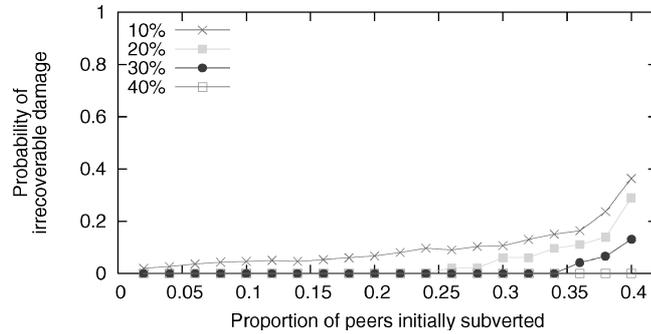


Fig. 8. Probability that the adversary causes irrecoverable damage. Each curve corresponds to a different level of friend bias.

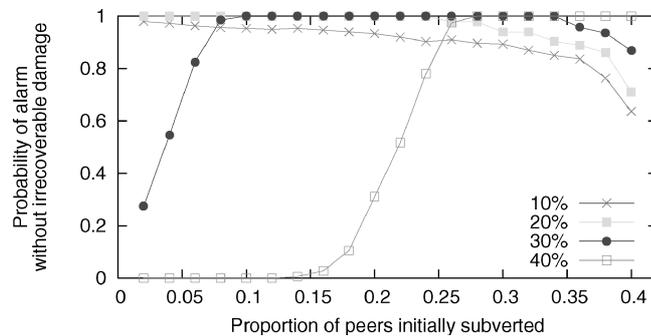


Fig. 9. Probability that the adversary is detected before causing irrecoverable damage. Each curve corresponds to a different level of friend bias.

the friend bias increases the initial subversion needed to make irrecoverable damage possible. Beyond this critical subversion level, the system suffers a gradual increase in the probability of irrecoverable damage. Note, however, that even with low friend bias and up to 40% initial subversion, the adversary has no more than a 36% chance of causing irrecoverable damage.

Finally, we evaluate the adversary's "experience" trying to cause irrecoverable damage with Figures 9 and 10. Figure 9 plots the probability that the adversary is detected *before* having caused irrecoverable damage. For the two lowest levels of friend bias, the adversary is almost guaranteed to be caught trying. For friend bias of 30%, he is likely to get away with trying out a stealth modification attack only for initial subversions at which he doesn't have the slightest chance of causing irrecoverable damage; he is again guaranteed to be caught trying for any level of subversion from 10% and higher, except for those few chances he has of causing irrecoverable damage. Even for the highest friend bias level of 40%, at which the adversary caused no irrecoverable damage in any of our simulations, his probability of being detected becomes significant as his initial subversion increases beyond 20%.

Figure 10 summarizes this behavior, showing that scenarios in which "nothing happens"—the adversary neither causes irrecoverable damage nor is he detected trying—only occur when he has no chance of success. The figure

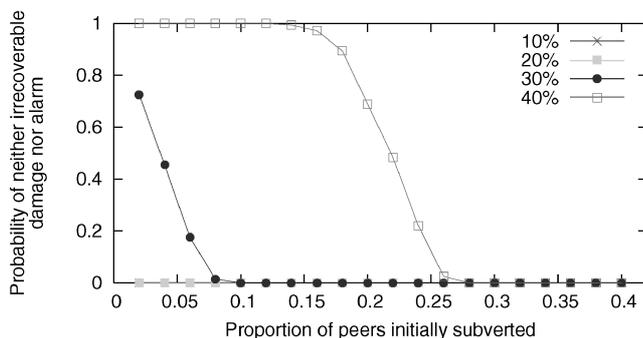


Fig. 10. Probability that the adversary plays the game but neither wins nor loses. Each curve corresponds to a different level of friend bias.

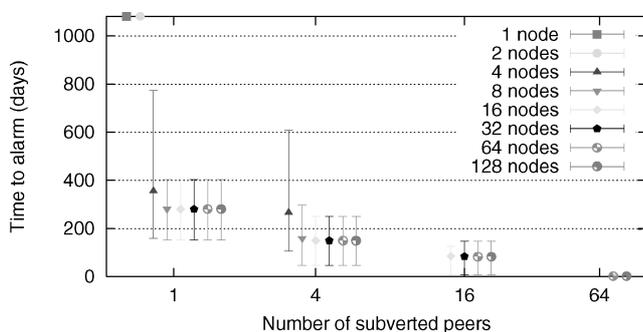


Fig. 11. Time from start of nuisance attack to first alarm at any peer (y axis) versus the number of peers subverted (x axis), for varying adversary resource sizes. Some attacks with few subverted peers do not cause alarms on our timescale.

illustrates that the stealth modification adversary is playing a high-risk, high-reward game: to have any chances of success (no more than 36% at best), he has to suffer very high levels of risk. This is another manifestation of the bimodal behavior of the system, making it virtually impossible for an adversary to get a guaranteed *and* safe win over his victims.

7.3 Nuisance Adversary

Figure 11 shows the effect of a nuisance adversary, with 1 to 128 nodes' worth of computing effort available, who subverts 1 to 64 peers, then attempts to raise an alarm at *any* peer. The simulation ends after 3 years or at the first such alarm. The error bars show minimum, average, and maximum times to the first alarm, over 20 runs per data point, with some attacks not generating an alarm in the first 3 years.

If the nuisance adversary subverts only a few peers, irrespective of her computing resources, she takes about 6 months to raise an alarm, or (see Figure 2) the equivalent of a random damage rate of once every 6–7 years. If she takes over a nonnegligible number of peers (16 or 64), irrespective of her computing resources, the alarm happens quickly. This seems to be suitable behavior, in that large compromises should be detected.

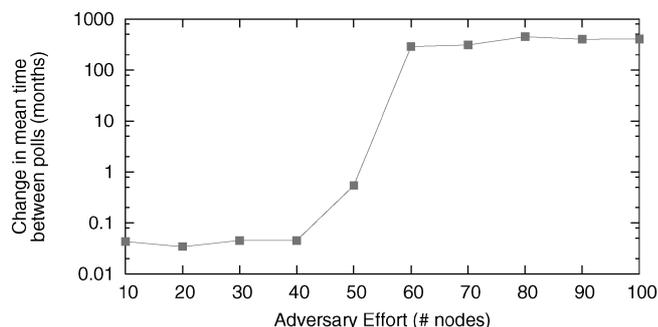


Fig. 12. Effect of the attrition adversary in increasing mean time between polls as a function of his effort. The y axis is in logarithmic scale.

7.4 Attrition Adversary

Figure 12 shows the increase in the average interpoll time at loyal peers as a function of the effort expended by the attrition adversary. Without sufficient resources, he has little effect on the system. If he can deploy sufficient resources, in our case about 60 machines, he can make peers busy enough to make all subsequently called polls inoperative. After three average interpoll intervals (9 months with our current operating parameters), affected peers raise an interpoll interval alarm which we did not simulate.

Assuming that once detected the attack can be suppressed and any accumulated damage repaired, a 9-month interruption in polls cannot affect the survivability of an AU, although it can affect the probability of a reader accessing a damaged AU replica. It is unclear, however, how easy it is to suppress an attack mounted and maintained by only 60 machines around the world. Further work is needed to understand the capabilities of an attrition adversary and the additional defensive needs of the LOCKSS system.

8. RELATED WORK

In common with the Byzantine-fault-tolerance (BFT) literature (e.g., Castro and Liskov [1999]; Chor and Dwork [1989]; Malkhi and Reiter [1998]; Reiter [1994]), our voting protocol derives an apparent prevailing opinion among a set of peers, some of whom are malicious. There are many differences; our population size is too large for BFT's global communication, we degrade gradually rather than mask the effects of failure or attack, and because we cannot assume an absolute upper bound on the malicious peers' resources we have to consider the possibility of being overwhelmed. We use sampling to avoid global knowledge or communication, rate-limiters to prevent our adversary's unlimited resources from overwhelming the system quickly, and integrated intrusion detection to preempt irrecoverable failure.

Our work has similarities with the antientropy protocol forming part of Bimodal Multicast [Birman et al. 1999], a reliable multicast protocol in which peers send digests of their message histories to randomly chosen other peers. Peers receiving these messages can detect omissions and request repairs from

the peer that sent the digest. The system's name comes from its bimodal distribution of delivery probability, which is similar to our distribution of poll results absent an attack. As in our case, it exploits the properties of random graphs. As the authors acknowledge, the lack of voting among the peers leaves the antientropy protocol vulnerable to malign peers.

Our work also shares a goal with some of the first peer-to-peer systems including Freenet [Clarke et al. 2000], FreeHaven [Dingledine et al. 2000], and the Eternity Service [Anderson 1996], namely, to make it hard for a powerful adversary to damage or destroy a document in the system. The other key goal of these systems is to provide anonymity for both publishers and readers of content, which we do not share. It would make our system both illegal and unworkable, since we often preserve content that must be paid for.

Several studies have proposed a persistent, peer-to-peer storage service including Intermemory [Chen et al. 1999], CFS [Dabek et al. 2001], Oceanstore [Kubiatowicz et al. 2000], PAST [Rowstron and Druschel 2001], and Tangler [Waldman and Mazières 2001]. Some (e.g., Oceanstore) implement access control by encrypting the data and thus do not solve our preservation problem, merely reducing it from preserving and controlling access to the content, to preserving and controlling access to the encryption key. Others (e.g., PAST) implement access control based on long-term secrets and smartcards or a key management infrastructure. Neither is appropriate for our application. Some (e.g., Intermemory) use cryptographic sharing to proliferate n partial replicas, from any $m < n$ of which the file can be reconstituted. Others (e.g., PAST) replicate the entire file, as we do, but do not allow control over where the replicas are located. The goal of the LOCKSS system is to allow librarians to take custody of the content to which they subscribe. This requires that each library keep its own copy of a document it has purchased, not share the load of preserving a small number of copies.

Moore et al. [2001] reported interesting measurements on packet-level denial-of-service attacks, types, and frequencies. For example, most attacks are relatively short, with 90% lasting less than an hour. Rarely do such attacks span multiple days. The relevance of this data to application-level denial-of-service attacks is questionable, but our simulated attacks require attention spans from the attackers several orders of magnitude longer.

The practice of making an attack costlier to the attacker than to the victim via provable effort has been proposed as a defense against otherwise cheap yet powerful attacks, such as email spam or DoS against cryptographic protocols. The idea first appeared under the term *pricing via processing* [Dwork and Naor 1992]. In our work, we measure cost by memory cycles, via a scheme described by Abadi et al. [2003]. Dwork et al. [2003] described a specific, practical application of that scheme to defending against email spam. In other work, CPU cycles [Back 2002; Dean and Stubblefield 2001; Dwork and Naor 1992] or even Turing tests⁴ have been used as the cost primitive. We favor the memory-bound approach for our purposes, because the variance in memory system performance is much smaller than that of CPU performance in today's off-the-shelf

⁴"Take Control of Your Inbox," by Spam Arrest, LLC. Visit the Web site <http://spamarrest.com>.

systems. Greater variance, that is, a greater disparity between the slowest and the fastest “producer of effort,” can hinder the effectiveness of pricing via processing [Douceur 2002].

9. FUTURE WORK

We have two immediate goals: to deploy an implementation to our test sites, and to improve the protocol’s performance against the attrition adversary. Before deployment, we need to emulate the initial version’s handling of many practical details, especially the “divide-and-conquer” search for damage in an AU formed of many documents. Our simulated attrition adversary can currently prevent 1000 loyal peers from running polls with about 60 malign nodes fully committed to the attack. We have identified several avenues for improving this inadequate performance, including using additional peer state to identify attrition attacks and improving the economic model to account for commitment of time as well as effort [Rosenthal et al. 2003]. We have described our current research agenda toward impeding attrition attackers elsewhere [Maniatis et al. 2004].

Our current adversary model starts by making a proportion of the peers malign; these peers remain malign for the duration. Other peers may have their AUs damaged by an attack, but they remain loyal for the duration. We need to enhance the model to account for peers becoming malign as they fall victim to vulnerabilities, and becoming loyal as their administrators patch and repair them.

We have analyzed potential adversaries from two points of view, their goals and the opportunities offered them by their roles in protocol implementations. But we have yet to show that our simulated adversary strategies are optimal. We have shown that the system’s cost structure and other parameters can be set appropriately for a realistic application, but we have yet to explore alternative cost structures and parameter values, or how to set them optimally. These investigations will be aided when we are able to validate the simulator against a deployed implementation.

Friend bias is an integral defense against adversaries, but we have given intuitive reasons why it should not be “too high.” However, we have not quantified those reasons well enough to allow a systematic identification of the optimal friend bias. We hope to develop adversary strategies that attack peers by first attacking their commonly invited voters, their friends most certainly. This will help us to quantify the impact of higher levels of friend bias.

10. CONCLUSION

Storage alone will not solve the problem of digital preservation. Academic materials have many enemies beyond natural bit rot: ideologies, governments, corporations, and inadequate budgets. It is essential that sound storage and administration practices are complemented with the institution of communities acting together to thwart attacks that are too strong or too extrinsic for such practices to protect against.

In a novel thrust toward this goal, we have built a new opinion poll protocol for the LOCKSS system, applying the design principles in Section 2:

- Cheap storage is unreliable.* We replicate all persistent storage across peers, audit replicas regularly, and repair any damage they find. Peer state is soft and rebuilt by normal system operations if it is lost.
- Use no long-term secrets.* Our peers need keep secrets only for the duration of a single poll. Without long-term secrets, attackers may spoof the identity of peers, but by requiring evidence of recent effort we reduce the time during which stability of identity matters to a few poll durations, and we use short-term secrets to reduce spoofing during a poll.
- Use inertia.* We provide the system with an analog of inertia by making *all* of its operations inherently expensive and by limiting the rate of possible change in the system. Because even the operations involved in failure are inherently time-consuming, it is very hard for attackers to overwhelm the system quickly, which provides time for humans to prevent an attack from resulting in catastrophic failure.
- Avoid third-party reputation.* Third-party reputation is vulnerable to slander and subversion of previously reliable peers, especially in the absence of strong identities. Further, we do not use accumulated evidence of a peer's good behavior. We instead require evidence of substantial recent effort to allow a peer to influence the outcome of a poll.

To the extent to which a peer does maintain a history of another peer's behavior, that history is very simple, derived from direct observation, and acts only as a hint. The system survives the loss or corruption of this memory at peers.
- Reduce predictability.* The reference lists are the only mechanism by which an attacker can damage a loyal peer's AU. Friend bias deprives the attacker of complete control of them, slowing an attack and increasing the risk of detection.
- Integrate intrusion detection intrinsically.* Random damage to individual replica AUs stored by peers is incoherent, resulting in polls that are either landslide agreement or landslide disagreement. An attacker attempting to change a preserved AU requires coherent damage to replica AUs across peers, which results in polls that are more closely contested. Contested polls, and thus coherent damage, raise an inconclusive poll alarm and lead to detection of the attack.
- Assume a strong adversary.* We allow for an adversary who can call on unlimited computing resources and unlimited identities, who can subvert or spoof a large proportion of the peers, and who has information about the parameters of each poll that a real attacker would be unlikely to possess.

We have shown that a combination of massive replication, rate limitation, inherent intrusion detection and costly operations can produce a peer-to-peer system with remarkable ability to resist attacks by some extraordinarily powerful adversaries over decades. Its lack of dependence on long-term secrets and stable identities blocks many of the paths by which systems are typically attacked. Although we developed the new LOCKSS protocol for an application with unusual

characteristics, especially its goal of preventing change, we nonetheless believe that the concepts and principles underlying the protocol will be useful in the design of other long-term large-scale applications operating in hostile environments.

The LOCKSS project is hosted by SourceForge, where the current implementation can be obtained. The Narses simulator is also available from SourceForge. Both carry BSD-style Open Source licenses.

APPENDIX: ECONOMIC CONSIDERATIONS

We use a memory-bound function (MBF) scheme due to Dwork et al. [2003]. Here we briefly describe it and compute appropriate costs to impose on poll initiation and voting. Recall that S is the size of the AU in cache lines.

A.1 Overview of a Memory-Bound Function Scheme

The goal of the MBF is to cause the *prover* of the necessary effort to incur a number C of cache misses and thus RAM accesses. If each of these takes t seconds, the prover must have used $C \cdot t$ seconds on a real computer. Memory bandwidths vary significantly less among commonly available architectures than CPU speeds do, making MBFs superior in provable effort to the CPU-bound functions previously proposed [Dean and Stubblefield 2001; Dwork and Naor 1992].

The scheme we use has two adjustable parameters, the cost, l , of verifying an effort proof and the ratio, E , between l and the cost of constructing the proof. We measure all costs in cache misses, so a proof costs $E \cdot l$ cache misses to construct and l cache misses to verify.

Dwork et al. [2003] described an MBF scheme that uses an incompressible fixed public data set T larger than any cache one is likely to meet currently. In our case, a gigabyte would be practical. An effort prover who must expend effort $E \cdot l$ is given as challenge a nonce n (so that the prover cannot reuse older effort proofs) and the values of E and l . In response, the prover must perform a series of pseudorandom walks in the table T . For each walk, the prover starts from a different position s of his choosing and computes a one-way value A based on n , s and the encountered elements of table T . The walk is cryptographically dependent on n and s ; it is constructed so that the number of encountered elements is l , and fetching each encountered element causes an L1 cache miss. Each walk, therefore, causes l cache misses.

The prover stops when he computes a value A that has 0 bits in its least significant $\log_2 E$ positions. With the given MBF scheme, it is expected the prover will try E walks with different starting positions s before finding an appropriate starting position s' ; this costs the prover $C = E \cdot l$ cache misses.

The s' that yielded the appropriate A is the effort proof. The verifier need only perform the random walk on T starting with the n he chose and the s' sent by the prover; this costs the verifier $V = l$ cache misses. If the resulting value A from this walk has the proper 0 bits in its last $\log_2 E$ positions, the verifier accepts the effort proof as valid.

We first describe how we use the MBF scheme in the construction and verification of votes. We then describe how we impose costs on vote construction and verification and how we choose appropriate parameters for each of these.

A.2 Vote Construction and Verification

A.2.1 Construction. As described in Sections 4.1.5 and 4.1.11, vote construction is divided into rounds. Each round consists of two parts: the construction of an MBF proof of effort and the hashing of a portion of the document. r denotes the total number of rounds for a given AU.

We design vote construction in a way that ensures the order of computation for each stage; specifically, we wish to ensure that voters cannot precompute or parallelize different stages of the construction of a vote (although individual stages of MBF proof construction may be parallelizable [Dwork et al. 2003]).

A vote contains the MBF proofs and content hashes computed during the two stages of every vote construction round. We use $[s_1, \dots, s_r]$ to denote the list of MBF effort proofs for the r rounds of vote construction and $[H_1, \dots, H_r]$ to denote the list of the corresponding content hashes.

At round i , the MBF proof computation stage takes as input the nonce n_i (we explain how we build this nonce below) and the proof parameters E_i and l_i , returning the proof s_i and the output value A_i (see Section A.1). This is followed by the content hash stage of round i , which computes the hash $h(s_i \parallel A_i \parallel \text{content}_i)$, where \parallel denotes bit string concatenation, h is our cryptographic hash function, and content_i denotes the portion of the content that we hash during round i (to be determined in Section A.3). The output of the hashing stage is the hash value H_i .

Because the proof and output of the MBF proof computation are included in the input to the cryptographic hash function h , the vote constructor cannot precompute H_i before having determined the MBF proof s_i and the corresponding output value A_i . We include both s_i and A_i in the hash (as opposed to only the proof s_i) to ensure that it is hard for the vote constructor to precompute H_i values for all possible s_i values. Instead, s_i and A_i together come from a large enough range of possible values that precomputing all likely hashes H_i by brute-force is intractable.

The nonce n_i input into the MBF proof effort computation of the i th round must be such that precomputing an effort proof for round i before the MBF proof effort computation or content hashing stage of round $i - 1$ is intractable. For the first round of vote construction, n_1 must be such that the voter cannot start computing its vote until it receives the poller's challenge. Consequently, we set $n_1 = h(\text{pollerChallenge} \parallel \text{pollID} \parallel \text{voterAddress})$, and, for $i > 1$, we similarly set $n_i = h(s_{i-1} \parallel A_{i-1} \parallel H_{i-1} \parallel \text{pollerChallenge} \parallel \text{pollID} \parallel \text{voterAddress})$.

A.2.2 Verification. Vote verification is also divided into rounds, with MBF and hashing stages in each round. The poller, guided by the list of MBF effort

proofs and hashes contained in the Vote it receives, verifies the computations of the voter using its local copy of the document.

At round i , the MBF proof verification stage takes as input the nonce n_i (see Section A.2.1), the proof parameters E_i and l_i , and the proof s_i (included in the message) and constructs the output value A_i (see Section A.1). If this value ends in $\log_2 E_i$ 0 bits, the verifier accepts the effort proof. Otherwise, the verifier deems the vote invalid.

If the verifier has yet to deem the vote disagreeing, it proceeds with the content hashing stage of the vote verification. Specifically, it computes the hash of its appropriate local content portion: $H'_i = h(s_i \| A_i \| \text{content}_i)$. If the resulting hash H'_i is different from the value H_i contained in the vote, the verifier deems the vote disagreeing and only verifies MBF efforts in the remaining rounds.

Note that, as with vote construction, the verifier sets n_1 using the poller's challenge, poll identifier, and the voter's identity, and all subsequent n_i 's using, again, the challenge, poll identifier, and voter's identity, as well as the effort proof s_{i-1} from the vote, the value A_{i-1} computed during the previous round's MBF stage, and the voter's hash H_{i-1} from the vote.

In our continuing implementation efforts, we are exploring a slightly re-ordered verification algorithm, in which all MBF proofs are verified before any hashing of content is performed. Since an invalid MBF proof is the only portion of a vote verification where the process may stop, this seems to be a strict improvement of the algorithm. We are currently evaluating the implications of this change on the analysis of our provable effort scheme.

A.3 Choice Of Voting Effort Sizes

Let $C_{MBF}(i)$ and $V_{MBF}(i)$ be the costs of the i th round MBF construction and verification, respectively, measured in L1 cache misses. Let $H(i)$ be the cost of hashing operations in the i th round, again measured in L1 cache misses.

We require (Section 4.2.1) that the cost of constructing a vote be greater than or equal to the cost of verifying that vote, even if the vote is disagreeing, corrupt, or malformed, for example, in round i its MBF verification fails or its content hash disagrees. Let i be the first round for which the vote hash and the verifier's hash disagree; this may be because the voter genuinely disagrees with the poller or because the voter maliciously constructs a vote whose hash portions are zero cost, arbitrarily chosen values. Let j be the first round for which the vote effort proof is invalid. From the definition of the vote construction/verification mechanisms, $i < j$, since once the poller detects an invalid MBF proof, it stops the vote verification process and, therefore, all subsequent hash operations. We define $C_v(i, j)$ to denote the cost of producing a vote that causes the verifier to stop checking hashes after round i , and that causes the verifier to declare the vote invalid after round j . Similarly, $V_v(i, j)$ denotes the cost of verifying such a vote. By convention, $C_v(i, j)$ denotes the cost of producing a valid vote when $j > r$, and of a loyally produced valid vote when $i, j > r$. To assist in this convention, we set $H(i) = 0$, $C_{MBF}(i) = 0$, and $V_{MBF}(i) = 0$ for $i < 1$ or $i > r$.

We wish to establish that

$$\begin{aligned}
& C_v(i, j) \geq V_v(i, j) \\
\Leftrightarrow & \sum_{k=1}^{i-1} [C_{MBF}(k) + H(k)] + \sum_{k=i}^{j-1} C_{MBF}(k) \geq \sum_{k=1}^i [V_{MBF}(k) + H(k)] + \sum_{k=i+1}^j V_{MBF}(k) \\
& \Leftrightarrow \sum_{k=1}^{i-1} C_{MBF}(k) + \sum_{k=i}^{j-1} C_{MBF}(k) \geq \sum_{k=1}^i V_{MBF}(k) + H(i) + \sum_{k=i+1}^j V_{MBF}(k) \\
& \Leftrightarrow \sum_{k=1}^{j-1} C_{MBF}(k) \geq \sum_{k=1}^{j-1} V_{MBF}(k) + V_{MBF}(j) + H(i) \\
& \Leftrightarrow \sum_{k=1}^{j-1} [C_{MBF}(k) - V_{MBF}(k)] \geq V_{MBF}(j) + H(i). \tag{7}
\end{aligned}$$

Intuitively, Inequality 7 captures the fact that the accrued premium that the voter pays round after round by computing effort proofs must “cover” the possible extra work that a poller performs to detect misbehavior. This extra work includes a round’s worth of hashing and an effort verification at the last round; the former is done before the poller detects a disagreement and the latter is done before the poller detects that the proof effort is invalid.

The inequality does not hold when the entire vote is bogus, that is, $j = 1$. In that case, the poller always pays the right hand side ($V_{MBF}(1)$ with no hashing), which is always greater than 0. However, to be invited to send this bogus vote, the malicious voter had to be in the poller’s reference list, which cost at least $C_v(1, r) = \sum_{k=i}^r C_{MBF}(k)$, far more than the $V_{MBF}(1)$ imposed on the poller. We explore the implications of this inequality for all other cases.

A.3.1 Invalid Vote. Inequality 7 describes exactly the balance of efforts between a malicious voter and a loyal poller when the voter produces an invalid vote.

A.3.2 Valid Agreeing Vote. We wish to ensure that $C_v(i, j) \geq V_v(i, j)$, where $i, j > r$. Inequality 7 becomes $\sum_{k=1}^r [C_{MBF}(k) - V_{MBF}(k)] \geq 0$, which is trivially true from the definition of the MBF provable effort scheme.

A.3.3 Valid Ostensibly Disagreeing Vote. In this case, the poller cannot be certain that the voter performed any hashing after the first disagreeing round hash it detects. A loyal voter does perform all hashing operations, but a malicious one performs no more than necessary, which requires strictly less effort than that put in by a loyal disagreeing voter.

For this case, we wish to ensure $C_v(i, j) \geq V_v(i, j)$, where $j > r$ but $i \in [1, r]$, which from Inequality 7 means $\sum_{k=1}^r [C_{MBF}(k) - V_{MBF}(k)] \geq H(i)$. Intuitively, in this case the inequality ensures that the total difference between provable effort computation and verification across all rounds is greater than the cost of the hashing operations of the most expensive round.

A.3.4 *Instantiation of the Scheme.* To obtain an appropriate instantiation of all parameters of our scheme, we simplify Inequality 7:

$$\begin{aligned}
\sum_{k=1}^{j-1} [C_{MBF}(k) - V_{MBF}(k)] &\geq V_{MBF}(j) + H(i) \\
\Leftrightarrow \sum_{k=1}^{j-1} [E_k l_k - l_k] &\geq l_j + H(i) \\
\Leftrightarrow \sum_{k=1}^{j-1} (E_k - 1) l_k &\geq l_j + H(i).
\end{aligned}$$

We make all E_i parameters constant and equal to E and set $l_i = H(i) = 2^{i-1}b$, where b is the number of cache lines in an AU block. These choices are arbitrary, but designed to simplify the task of instantiating these parameters. Other choices may be possible.

$$\begin{aligned}
\sum_{k=1}^{j-1} (E_k - 1) l_k &\geq l_j + H(i) \\
\Leftrightarrow (E - 1) \sum_{k=1}^{j-1} 2^{k-1} &\geq 2^{j-1} + 2^{i-1} \\
\Leftrightarrow (E - 1)(2^{j-1} - 1) &\geq 2^{j-1} + 2^{i-1}.
\end{aligned}$$

However, for all $E \geq 4$ and $j > 1$ this holds (recall that $i < j$ and that E must be a power of 2).

Higher values for E might be desirable for smaller content sizes to make rate limitation more effective.

If S is the size of the content in cache lines and $E = 4$, the cost of constructing a vote correctly is $C_v(r + 1, r + 1)$ (by convention):

$$\begin{aligned}
C_v(r + 1, r + 1) &= \sum_{k=1}^r [C_{MBF}(k) + H(k)] \\
&= \sum_{k=1}^r [E \times 2^{k-1}b + 2^{k-1}b] \\
&= (E + 1)b \sum_{k=1}^r 2^{k-1} \\
&= 52^r b \\
&= 5S.
\end{aligned}$$

Through similar algebraic manipulation, the cost of verifying an agreeing vote is $2S$ and the cost of verifying a disagreeing vote ranges between S and $2S$.

A.4 Designation of Poll Initiation Effort Sizes

We want the cost C_p of poll initiation per invitee to be at least the induced cost on the invitee:

$$C_p \geq V_p + C_v(r + 1, r + 1), \quad (8)$$

where V_p is the cost of verifying the poll initiation effort.

Based on the analysis of the voting costs above, this means

$$\begin{aligned} E_p l_p &\geq l_p + 5S \\ \Leftrightarrow (E_p - 1)l_p &\geq 5S. \end{aligned} \quad (9)$$

One choice for the MBF parameters is $E_p = 4$ and $l_p = (5/3)S$. The poller must expend $C_p = (20/3)S$ cache misses per invitee, and each invitee must spend $V_p = (5/3)S$ cache misses to verify the invitation.

ACKNOWLEDGMENTS

We have benefited from discussions with colleagues who provided helpful feedback, test code, and early paper manuscripts on related work or tools. In particular, we would like to express our gratitude to Yanto Muliadi, who participated and greatly contributed in earlier versions of this work, as well as to Cynthia Dwork, Kevin Lai, Aris Gionis, Katerina Argyraki, Lada Adamic, Sergio Marti, Ed Swierk, Hector Garcia-Molina, John Heidemann, the anonymous reviewers, and the LOCKSS engineering team and beta sites. We especially thank Prashanth Bungale for uncovering a critical bug in our simulations. Finally, Vicky Reich, the director of the LOCKSS program, immeasurably helped us to stay on track and remain nourished before deadlines.

REFERENCES

- ABADI, M., BURROWS, M., MANASSE, M., AND WOBBER, T. 2003. Moderately hard, memory-bound functions. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium* (San Diego, CA).
- ANDERSON, R. J. 1996. The Eternity Service. In *Proceedings of the 1st International Conference on the Theory and Applications of Cryptology* (PRAGOCRYPT 1996, Prague, Czech Republic).
- ARL—ASSOCIATION OF RESEARCH LIBRARIES. 2001. ARL Statistics 2000-01. Available online at <http://www.arl.org/stats/arlstat/01pub/intro.html>.
- BACK, A. 2002. Hashcash—a denial of service counter measure. Available online at <http://www.hashcash.org/hashcash.pdf>.
- BIRMAN, K., HAYDEN, M., OZKASAP, O., XIAO, Z., BUDI, M., AND MINSKY, Y. 1999. Bimodal Multicast. *ACM Trans. Comput. Syst.* 17, 2, 41–88.
- CASTRO, M. AND LISKOV, B. 1999. Practical Byzantine fault tolerance. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*. USENIX Association (New Orleans, LA). 173–186.
- CHEN, Y., EDLER, J., GOLDBERG, A., GOTTLIEB, A., SOBTI, S., AND YIANILOS, P. 1999. A prototype implementation of archival intermemory. In *Proceedings of the International Conference on Digital Libraries* (Berkeley, CA). 28–37.
- CHOR, B. AND DWORK, C. 1989. Randomization in Byzantine agreement. *Advan. Comput. Res.* 5, 443–497.
- CLARKE, I., SANDBERG, O., WILEY, B., AND HONG, T. W. 2000. Freenet: A distributed anonymous information storage and retrieval system. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, H. Federrath, Ed. Lecture Notes in Computer Science, vol. 2009. Springer, Berlin, Germany, 46–66.

- CONSERVATIONONLINE. 2003. Electronic storage media. Available online at <http://palimpsest.stanford.edu/bytopic/electronic-records/electronic-storage-media/>.
- DABEK, F., KAASHOEK, M. F., KARGER, D., MORRIS, R., AND STOICA, I. 2001. Wide-area Cooperative storage with CFS. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles* (Chateau Lake Louise, Banff, AB, Canada). 202–215.
- DAVIS, D. 1996. Compliance defects in public-key cryptography. In *Proceedings of the 6th USENIX Security Symposium* (San Jose, CA). 171–178.
- DEAN, D. AND STUBBLEFIELD, A. 2001. Using client puzzles to protect TLS. In *Proceedings of the 10th USENIX Security Symposium* (Washington, DC).
- DIFFIE, W. 2003. Perspective: Decrypting the secret to strong security. Available online at <http://news.com.com/2010-1071-980462.html>.
- DIFFIE, W. AND HELLMAN, M. E. 1976. New directions in cryptography. *IEEE Trans. Inform. Theor.* *IT-22*, 6 (Nov.), 644–654.
- DINGLEDINE, R., FREEDMAN, M. J., AND MOLNAR, D. 2000. The Free Haven Project: Distributed anonymous storage service. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, H. Federrath, Ed. Lecture Notes in Computer Science, vol. 2009. Springer, Berlin, Germany, 67–95.
- DOUCEUR, J. 2002. The Sybil attack. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems* (Boston, MA). 251–260.
- DWORK, C., GOLDBERG, A., AND NAOR, M. 2003. On memory-bound functions for fighting spam. In *Proceedings of the 23rd Annual International Cryptology Conference* (Santa Barbara, CA).
- DWORK, C. AND NAOR, M. 1992. Pricing via processing. In *Proceedings of the 12nd Annual International Cryptology Conference* (Santa Barbara, CA). 139–147.
- ERDOS, M. AND CANTOR, S. 2002. Shibboleth Architecture DRAFT v05. Available online at <http://shibboleth.internet2.edu/docs/draft-internet2-shibboleth-arch-v05.pdf>. Work in progress.
- FRIEDMAN, E. J. AND RESNICK, P. 2001. The social costs of cheap pseudonyms. *J. Econom. Managem. Strat.* *10*, 2 (Summer), 173–199.
- FULLMER, C. 2002. Storage and multimedia: The facts and more. Available online at http://www.cse.ucsc.edu/classes/cmpe003/Fall102/L11_ch6.pps.
- GIULI, T. AND BAKER, M. 2002. Narses: A scalable, flow-based network simulator. Tech. rep. arXiv:cs.PF/0211024. Computer Science Department, Stanford University, Stanford, CA.
- HORLINGS, J. 2003. Cd-r's binnen twee jaar onleesbaar. Available online at <http://www.pc-active.nl/toonArtikel.asp?artikelID=508>. <http://www.cdfreaks.com/news/7751>.
- KELLER, M., REICH, V., AND HERKOVIC, A. 2003. What is a library anymore, anyway? *First Monday* *8*, 5 (May). Available online at http://www.firstmonday.org/issues/issue8_5/keller/index.html.
- KUBIATOWICZ, J., BINDEL, D., CHEN, Y., CZERWINSKI, S., EATON, P., GEELS, D., GUMMADI, R., RHEA, S., WEATHERSPOON, H., WEIMER, W., WELLS, C., AND ZHAO, B. 2000. OceanStore: An architecture for global-scale persistent storage. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems* (Cambridge, MA). 190–201.
- MALKHI, D. AND REITER, M. 1998. Byzantine quorum systems. *J. Distrib. Comput.* *11*, 4 (Oct.), 203–213.
- MANIATIS, P., GIULI, T., ROUSSOPOULOS, M., ROSENTHAL, D. S. H., AND BAKER, M. 2004. Impeding attrition attacks in P2P systems. In *Proceedings of the 11th ACM SIGOPS European Workshop* (Leuven, Belgium).
- MANIATIS, P., ROUSSOPOULOS, M., GIULI, T., ROSENTHAL, D. S. H., BAKER, M., AND MULIADI, Y. 2003. Preserving peer replicas by rate-limited sampled voting. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles* (Bolton Landing, NY). 44–59.
- MANNIS, B. AND SHAHANI, C. J. 2003. Longevity of CD media research at the Library of Congress. Available online at <http://www.loc.gov/preserv/study%20of%20CD%20longevity.pdf>.
- MICHALAKIS, N., CHIU, D.-M., AND ROSENTHAL, D. S. H. 2003. Long term data resilience using opinion polls. In *Proceedings of the 22nd IEEE International Performance Computing and Communications Conference* (Phoenix, AZ).
- MOGGE, D. 1999. Seven years of tracking electronic publishing: The ARL Directory of Electronic Journals, Newsletters and Academic Discussion Lists. *Library Hi Tech* *17*, 1, 17–25.

- MOORE, D., VOELKER, G. M., AND SAVAGE, S. 2001. Inferring Internet denial-of-service activity. In *Proceedings of the 10th USENIX Security Symposium*.
- PATTERSON, D. A., GIBSON, G., AND KATZ, R. H. 1988. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (Chicago, IL). 109–116.
- QUIRKE, C. 2002. Hard drive data corruption. Available online at <http://users.iafrica.com/c/cq/cquirke/baddata.htm>.
- REITER, M. 1994. Secure agreement protocols: Reliable and atomic group multicast in Rampart. In *Proceedings of the 2nd ACM Conference on Computer and Communications Security* (Fairfax, VA). 68–80.
- RHEA, S., GEELS, D., ROSCOE, T., AND KUBIATOWICZ, J. 2004. Handling Churn in a DHT. In *Proceedings of the Usenix Annual Technical Conference* (Boston, MA).
- ROSENTHAL, D. S. H. 2003. A digital preservation network appliance based on OpenBSD. In *Proceedings of BSDcon 2003* (San Mateo, CA).
- ROSENTHAL, D. S. H. AND REICH, V. 2000. Permanent Web publishing. In *Proceedings of the USENIX Annual Technical Conference, Freenix Track* (San Diego, CA). 129–140.
- ROSENTHAL, D. S. H., ROUSSOPOULOS, M., MANIATIS, P., AND BAKER, M. 2003. Economic measures to resist attacks on a peer-to-peer network. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems* (Berkeley, CA).
- ROWSTRON, A. AND DRUSCHEL, P. 2001. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles* (Chateau Lake Louise, Banff, AB, Canada). 188–201.
- STANIFORD, S., PAXSON, V., AND WEAVER, N. 2002. How to Own the Internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium* (San Francisco, CA). 149–167.
- TENOPIR, C. 2004. Online scholarly journals: How many? *The Library Journal* 2 (Feb.). Available online at <http://www.libraryjournal.com/index.asp?layout=articlePrint&articleID=C%A374956>.
- U. S. GOVERNMENT. 2004. Web design and usability guidelines. Available online at <http://usability.gov/guidelines/softhard.html>.
- VENEMA, W. 1996. Murphy's Law and computer security. In *Proceedings of the 6th USENIX Security Symposium* (San Jose, CA).
- WALDMAN, M. AND MAZIÈRES, D. 2001. Tangler: A censorship-resistant publishing system based on document entanglements. In *Proceedings of the 8th ACM Conference on Computer and Communications Security* (Philadelphia, PA). 126–135.
- WILLIAMSON, M. 2002. Throttling viruses: Restricting propagation to defeat malicious mobile code. In *Proceedings of the 18th Annual Computer Security Applications Conference* (Las Vegas, NV).
- WINGFIELD, N. 2002. EBay's figurine scandal: Auction site merchant disappears with the goods. *Wall Street J.*, Feb 22.

Received October 2003; revised July 2004; accepted May 2004