

WebSOS: An Overlay-based System For Protecting Web Servers From Denial of Service Attacks^{*}

Angelos Stavrou Debra L. Cook William G. Morein
Angelos D. Keromytis Vishal Misra Dan Rubenstein

*Department of Computer Science
Columbia University in the City of New York*

Abstract

We present WebSOS, a novel overlay-based architecture that provides guaranteed access to a web server that is targeted by a denial of service (DoS) attack. Our approach exploits two key characteristics of the web environment: its design around a human-centric interface, and the extensibility inherent in many browsers through downloadable “applets.” We guarantee access to a web server for a large number of *previously unknown* users, without requiring pre-existing trust relationships between users and the system, by using Reverse Graphic Turing Tests. Furthermore, our system makes it easy for service providers to charge users, providing incentives to a commercial offering of the service. Users can dynamically decide whether to use the WebSOS overlay, based on the prevailing network conditions.

Our prototype requires *no modifications* to either servers or browsers, and makes use of graphical Turing tests, web proxies, and client authentication using the SSL/TLS protocol, all readily supported by modern browsers. We then extend this system with a credential-based micropayment scheme that combines access control and payment authorization in one operation. Turing Tests ensure that malicious code, such as a worm, cannot abuse a user’s micropayment wallet. We use the WebSOS prototype to conduct a performance evaluation over the Internet using PlanetLab, a testbed for experimentation with network overlays. We determine the end-to-end latency using both a Chord-based approach and our *shortcut* extension. Our evaluation shows the latency increase by a factor of 7 and 2 respectively, confirming our simulation results.

Key words: Security and Protection, Denial of Service, Network Topology, Overlay Networks, Security, Reliability, Graphic Turing Tests, Web Proxies, Java.

1 Introduction

The Web is increasingly being used for different kinds of services and interactions with, and between humans. Beyond displaying static content such as home pages or academic papers, the web is actively used for such diverse tasks as e-mail, banking, consumer purchasing, marketing, stock-quote dissemination and trading, and real-time communication. The wide availability of high-quality browsers and servers, as well as programmers' and users' familiarity with the tools and concepts behind web browsing ensure that ongoing creation of additional services.

Such an environment provides a rich set of targets for motivated attackers. This has been demonstrated by the large number of vulnerabilities and exploits against web servers, browsers, and applications. Traditional security considerations revolve around protecting the network connection's confidentiality and integrity, protecting the server from break-in, and protecting the client's private information from unintended disclosure. To that end, several protocols and mechanisms have been developed, addressing these issues individually. However, one area that has long been neglected is that of service availability in the presence of denial of service (DoS) attacks, and their distributed variants (DDoS).

Previous approaches that address the general network DoS problem [1–3] are *reactive*: they monitor traffic at a target location, waiting for an attack to occur. Once the attack is identified, typically via analysis of traffic patterns and packet headers, filters may be established in an attempt to block the offenders. The two main problems with this approach are the accuracy with which legitimate traffic can be distinguished from the DoS traffic, and the robustness of the mechanism for establishing filters deep enough in the network so that the effects of the attack are minimized.

Furthermore, Internet Service Providers (ISPs) seem reluctant to deploy such mechanisms. Investment in the necessary infrastructure and operational support are discouraged because such mechanisms represent a poor value proposition: fundamentally, ISPs cannot charge users for the use of such mechanisms. One possible so-

* This work is supported in part by DARPA contract No. F30602-02-2-0125 (FTN program) and by the National Science Foundation under grant No. ANI-0117738 and CAREER Award No. ANI-0133829, with additional support from Cisco and the Intel IT Research Council. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Email addresses: angel@cs.columbia.edu (Angelos Stavrou), dcook@cs.columbia.edu (Debra L. Cook), wgm2001@cs.columbia.edu (William G. Morein), angelos@cs.columbia.edu (Angelos D. Keromytis), misra@cs.columbia.edu (Vishal Misra), danr@cs.columbia.edu (Dan Rubenstein).

lution would be a system with the ability to both protect against DoS attacks and provide a service payment scheme that would allow ISPs to recoup their costs and support the continued operation and maintenance of this infrastructure. Such incentives would motivate router manufacturers to provide appropriate support in their products. Thus, we wish to explore mechanisms that could support a charging model.

We introduce *WebSOS*, an adaptation of the *Secure Overlay Services (SOS)* architecture [4]. Our intent is to prevent congestion-based DDoS attacks from denying *any* user's access to web servers targeted by those attacks. The novel aspects of WebSOS are (a) its use of *graphic Turing tests* in lieu of (or in addition to) strong client authentication (as was proposed in SOS) to distinguish between human users and automated attack *zombies*, and (b) its transparency to browsers and servers, by taking advantage of browser extensibility. Although WebSOS itself protects only web traffic, it can be used to enable routing of other types of traffic by establishing IPsec tunnels through the overlay; the web-based authentication is leveraged to create a channel for other traffic, as we showed in [5]. We should also add that we envision WebSOS used when an attack is detected, *i.e.*, when a client cannot directly reach the destination web server. Thus, when no attacks are underway, the impact of WebSOS in network performance and other overheads is zero. An additional design goal of WebSOS, which was achieved, was avoiding changes to protocols and network elements (such as routers).

WebSOS protects the portion of the network immediately surrounding attack targets (*i.e.*, the web servers) by high-performance routers that aggressively filter and block all incoming connections from hosts that are not approved, as shown in Figure 1. These routers are “deep” enough in the network that the attack traffic does not adversely impact innocuous traffic, typically in an ISP's Point Of Presence (POP), as we discuss in Sections 2.4 and 3. The identities of the approved nodes is kept secret so that attackers cannot impersonate them to pass through the filter. These nodes (which can be as few as 2 or 3) are picked from a set of nodes that are distributed throughout the wide area network. This superset forms a *secure overlay*: any transmissions that wish to traverse the overlay must first be validated at any of the entry points of the overlay using either a cryptographic authentication mechanism, or a Graphic Turing test to distinguish humans from attack scripts [6]. Once inside the overlay, the traffic is tunneled securely to one of the approved (and secret from attackers) nodes, which can then forward the validated traffic through the filtering routers to the target. Thus, there are two main principles behind our design. The first is the elimination of communication pinch-points, which constitute attractive DoS targets, via a combination of filtering and overlay routing to obscure the identities of the sites whose traffic is permitted to pass through the filter. The second is the ability to recover from random or induced failures within the forwarding infrastructure or the secure overlay nodes. In [4], we gave an analytical model for the resilience of SOS against attackers that concentrate their attacks on the overlay nodes. We summarize these results in Section 2.4.

To demonstrate the potential of WebSOS as a value-added service, we also extended the basic WebSOS architecture to use a credential-based micropayment scheme that combines access control and payment authorization. Our architecture allows ISPs to accurately charge web clients and servers. Clients can dynamically decide whether to use WebSOS, based on the prevailing network conditions (*e.g.*, the client can try to contact the web-server through WebSOS if a direct connection fails, either through manual user intervention or by detecting the connection failure in our client-side proxy). Although practically any micropayment system can be used in our model, we chose a payment system that can inter-operate with WebSOS' distributed architecture and provide the necessary user credentials. OTPchecks [7] encompasses all these properties: it is a simple distributed scheme, intended for general Internet-based micropayments that produces bank-issued users' credentials which can in turn be used to acquire small-valued payment tokens. It has very low transaction overhead and can be tuned to use different risk strategies for different environments making it a suitable payment solution for a wide range of on-line services. Although the pay-per-use component of the architecture is technically independent of the DoS protection guarantees, we feel it is necessary to at least investigate the "market friendliness"¹ of any new proposed anti-DoS system.

WebSOS is the first instantiation of the SOS architecture. We use this instantiation to evaluate the performance of the underlying overlay routing mechanism both in a local area scenario and over the Internet using the PlanetLab testbed [8]. The results show that the average increase in end-to-end latency is a factor of 2 to 3 beyond what is achieved using the standard web infrastructure. We believe this modest increase is an acceptable alternative to providing no service. Such a service can be used on an as-needed basis, and hence need not impact performance when no attack is in progress. These results validate our simulation analyses, where we used real ISP topologies to determine the added average latency imposed by the WebSOS mechanism.

1.1 WebSOS Architectural Scope

DoS attacks can take many forms, depending on the resource the attacker is trying to exhaust. For example, an attacker can try to cause the web server to perform excessive computation, or exhaust all available bandwidth to and from the server. In all forms, the attacker's goal is to deny use of the service to other users. Apart from the annoyance factor, such an attack can prove particularly damaging for time- or life-critical services (*e.g.*, tracking the spread of a real-world epidemic), or when the attack persists over several days². Of particular interest are *link congestion* at-

¹ Of course, market friendliness involves more than the ability to charge a user, which is a necessary but not sufficient condition for deployment.

² In one instance of a persistent DoS attack, a British ISP was forced out of business because it could not provide service to its customers.

tacks, whereby attackers identify “pinch” points in the communications substrate and render them inoperable by flooding them with large volumes of traffic. An example of an obvious attack point is the location (IP address) of the destination that is to be secured, or the routers in its immediate network vicinity; sending enough attack traffic will cause the links close to the destination to be congested and drop all other traffic. *It is such attacks that WebSOS was designed to address. Solving the much harder general denial-of-service problem where attackers could potentially have enough resources to physically partition a network is not addressed in this paper.* Furthermore, we do not consider algorithmic denial of service attacks [9].

We assume that attackers are smart enough to exploit features of the architecture that are made publicly available, such as the set of nodes that form the overlay. However, we do not specifically consider how to protect the architecture against attackers who can infiltrate the security mechanism that distinguishes legitimate traffic from (illegitimate) attack traffic: we assume that communications between overlay nodes remain secure so that an attacker cannot send illegitimate communications, masking them as legitimate. In addition, it is conceivable that more intelligent attackers could monitor communications between nodes in the overlay and, based on observed traffic statistics, determine additional information about the current configuration. We leave it as future work to explore how WebSOS can be used to protect against attacks by such highly specialized and sophisticated attackers. Some work in that direction can be found in [10].

1.2 Paper Organization

The remainder of this paper is organized as follows. Section 2 gives an overview of Secure Overlay Services (SOS) and graphic Turing tests, and discusses the specifics of the WebSOS architecture, as well as an overview of OTPchecks, our micropayment scheme. In Section 3 we present our simulation results, using real ISP topologies. Section 4 presents details of our prototype implementation, while Section 5 contains our performance evaluation. Section 6 discusses other work in DoS detection, prevention, and mitigation. Finally, Section 7 concludes the paper.

2 The WebSOS Architecture

Because our approach is based on the Secure Overlay Services (SOS) [4] architecture, we first highlight its important aspects. We also briefly describe Graphic Turing tests, which implement human-to-overlay authentication, and the microbilling architecture we integrated in WebSOS. We close this section with a description of WebSOS itself, and the sequence of operations in using it.

2.1 Overview of SOS

Fundamentally, the goal of the SOS infrastructure is to distinguish between authorized and unauthorized traffic. The former is allowed to reach the destination, while the latter is dropped or is rate-limited. Thus, at a very basic level, SOS requires the functionality of a firewall “deep” enough in the network that the access link to the target is not congested. This imaginary firewall performs access control by using protocols such as IPsec [11]. This generally pre-supposes the presence of authentication credentials (*e.g.*, X.509 [12] certificates) that a user can use to gain access to the overlay. *We consider this one of the the largest drawbacks to SOS, as it precludes casual access to a web server by anonymous, yet benign users.* In WebSOS, we extend the authentication model to include Graphic Turing Tests, thereby allowing previously-unknown humans to access a protected site without allowing automated attacks to do the same. Both types of authentication (cryptographic and GTT based) can be used simultaneously in WebSOS, allowing pre-authorized automated processes to access the overlay in times of attack. Note that when the system is not under attack, automated processes such as Google indexing (as well as human users) can access the site directly.




Fig. 1. **Basic SOS architecture.** SOAP stands for **Secure Overlay Access Point**, and represents an entry point to the SOS overlay. SOS nodes can serve any of the roles of SOAP, Beacon, or Secret Servlet.

Since traditional firewalls themselves are susceptible to DoS attacks, what is really needed is a distributed firewall [13,14]. To avoid the effects of a DoS attack against the firewall connectivity, instances of the firewall are distributed across the network. Expensive processing, such as cryptographic protocol handling, is farmed out to a large number of nodes. However, firewalls depend on topological restrictions in the network to enforce access-control policies. In what we have described so far, an attacker can launch a DoS attack with spoofed traffic purporting to originate from one of these firewalls, whose identity cannot be assumed to remain forever secret. The insight of SOS is that, given a sufficiently large group of such firewalls, one can select a very small number of these as the designated authorized forwarding stations: only traffic forwarded from these will be allowed through the filtering router. In SOS, these nodes are called *secret servlets*. All other firewalls must forward traffic for the protected site to these servlets. Figure 1 gives a high-level overview of a

SOS infrastructure that protects a target node or site so that it only receives legitimate transmissions. Note that the secret servlets can change over time, and that multiple sites can use the same SOS infrastructure.




Fig. 2. Chord-based overlay routing.

To route traffic inside the overlay, SOS uses Chord [15], which can be viewed as a routing service that can be implemented atop the existing IP network fabric, *i.e.*, as a network overlay. Consistent hashing [16] is used to map an arbitrary identifier to a unique destination node that is an active member of the overlay.

In Chord, each node is assigned a numerical identifier (ID) via a hash function in the range $[0, 2^m]$ for some pre-determined value of m . The nodes in the overlay are ordered by these identifiers. The ordering is cyclic (*i.e.*, wraps around) and can be viewed conceptually as a circle, where the next node in the ordering is the next node along the circle in the clockwise direction.

Each overlay node maintains a table that stores the identities of m other overlay nodes. The i^{th} entry in the table is the node whose identifier x equals or, in relation to all other nodes in the overlay, most immediately follows $x + 2^{i-1} \pmod{2^m}$, as shown in Figure 2. When overlay node x receives a packet destined for ID y , it forwards the packet to the overlay node in its table whose ID precedes y by the smallest amount. In the example, if node 7 receives a packet whose destination is the identifier 20, the packet will route from 7 to 16 to 17. When the packet reaches node 17, the next node in the overlay is 22, and hence node 17 knows that 22 is responsible for identifier 20. The Chord algorithm routes packets around the overlay “circle”, progressively getting closer to the desired overlay node. $O(m)$ overlay nodes are visited. Typically, the hash functions used to map nodes to identifiers do not attempt to map two geographically close nodes to nearby identifiers. Hence, it is often the case that two nodes with consecutive identifiers are geographically distant from one another within the network.

The Chord service is robust to changes in overlay membership, and each node's list is adjusted to account for nodes leaving and joining the overlay such that the above properties continue to hold.

SOS uses the IP address of the target (*i.e.*, web server) as the identifier to which the hash function is applied. Thus, Chord can direct traffic from any node in the overlay to the node that the identifier is mapped to, by applying the hash function to the target's IP address. This node, where Chord delivers the packet, is not the target, nor is it necessarily the secret servlet. It is simply a unique node that will be eventually be reached, after up to $m = \log N$ overlay hops, regardless of the entry point. This node is called the *beacon*, since it is to this node that packets destined for the target are first guided. Chord therefore provides a robust and reliable, while relatively unpredictable for an adversary, means of routing packets from an overlay access point to one of several beacons.

Finally, the secret servlet uses Chord to periodically inform the beacon of the secret servlet's identity. Should the servlet for a target change, the beacon will find out as soon as the new servlet sends an advertisement. If the old beacon for a target drops out of the overlay, Chord will route the advertisements to a node closest to the hash of the target's identifier. Such a node will know that it is the new beacon because Chord will not be able to further forward the advertisement. By providing only the beacon with the identity of the secret servlet, traffic can be delivered from any firewall to the target by traveling across the overlay to the beacon, then from the beacon to the secret servlet, and finally from the secret servlet, through the filtering router, to the target. This allows the overlay to scale for arbitrarily large numbers of overlay nodes and target sites. *Unfortunately, this also increases the communication latency, since traffic to the target must be redirected several times across the Internet.* If the overlay only serves a small number of target sites, regular routing protocols may be sufficient.

2.2 *Graphic Turing Tests (GTT)*

In order to prevent automated attacks from breaching the overlay, a CAPTCHA [17] visual test is implemented at the entry point of the overlay to verify the presence of a human user. CAPTCHA (Completely Automated Public Turing test to Tell Computers and Humans Apart) is a program that can generate and grade tests that most humans can pass, but automated programs cannot.

The particular CAPTCHA realization we use is GIMPY, which concatenates an arbitrary sequence of letters to form a word and renders a distorted image of the word as shown in Figure 3. GIMPY relies on the fact that humans can read the words within the distorted image and current automated tools cannot. The human authenticates himself/herself by entering as ASCII text the same sequence of letters

CAPTCHA Implementation for SOS Project



Please type the text you see in the above picture

please hit refresh/reload to view another image

This Captcha library was obtained from [CMU CAPTCHA Project](#)

Fig. 3. WebSOS implementation of user Web Challenge using CAPTCHA. The challenge in this case is “fwtst”.

as what appears in the image. Updating the GIMPY interface to WebSOS can be performed without modifying the other architectural components.

When a user passes the GTT, the access point issues a short-lived X.509 [12] certificate. This certificate is signed by the entity operating the overlay, authorizing the holders to access the web service. The certificate is set to expire after 30 minutes (configurable), and contains the IP address of the client (to avoid reuse by multiple zombies). The overlay securely proxies all traffic from the source to the target via one of the beacons, as before.

Although recent advances in visual pattern recognition [18] can defeat GIMPY, there is no solution to date that can recognize complicated images or relation between images like Animal-PIX. Although for demonstration purposes in our prototype, described in Section 4, we use GIMPY, we can easily substitute it with any other instance of Graphic Turing test. Use of a GTT also minimizes large-scale “wallet depletion” attacks, whereby an attacker that manages to subvert a user’s workstation can spend that user’s money to perform the attack.

2.3 OTPchecks Micropayment System

In addition to the Graphic Turing test, we also integrated an (optional) micropayment mechanism to WebSOS. Our goal was to demonstrate that it is possible to design DoS-protection mechanisms that offer incentives to ISPs (or other entities) to deploy and manage, as they offer a clear way of recouping the associated costs. Here, we provide some background on the microbilling system we used; we explain its use in WebSOS in Section 2.6.

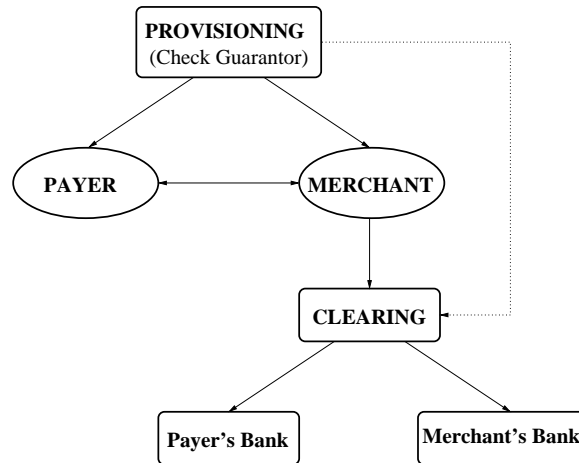


Fig. 4. **Microbilling architecture diagram.** We have the generic terms for each component. The arrows represent communication between the two parties: Provisioning issues credentials to Payers and Merchants; these communicate to complete transactions; Merchants send transaction information to Clearing which verifies the transaction and posts the necessary credits/charges or arranges money transfers. Provisioning and Clearing exchange information on the status of Payer and Merchant accounts.

The general architecture of our microbilling system [19] is shown in figure 4. The Check Guarantor plays the role of *Provisioning*, the Network User plays the role of *Payer*, and the Network Storage Provider (or another NU acting as an NSP) plays the role of the *Merchant*. *Clearing* is done either by a financial institution (if real money is used) or by a selected user of the system (when loyalty points or “play money” are used).

In this system, the *Provisioning* agent issues KeyNote [20] credentials to *Payers* and *Merchants*. These credentials describe the conditions under which a Payer is allowed to perform a transaction, and the fact that a Merchant is authorized to participate in a particular transaction. When a Payer wants to buy something from a Merchant, the Merchant first encodes the details of the proposed transaction into an *offer* which is transmitted to the Payer.

If the Payer wishes to proceed, she must issue to the Merchant a microcheck for this offer. The microchecks are also encoded as KeyNote credentials that authorize payment for a specific transaction. The Payer creates a KeyNote credential signed with her public key and sends it, along with her Payer credential, to the Merchant. This credential is effectively a check signed by the Payer (the Authorizer) and payable to the Merchant (the Licensee). The conditions under which this check is valid match the offer sent to the Payer by the Merchant. Part of the offer is a nonce, which maps payments to specific transactions, and prevents double-depositing of microchecks by the Merchant.

To determine whether he can expect to be paid (and therefore whether to accept the payment), the Merchant passes the action description (the attributes and values

in the offer) and the Payer's key along with the Merchant's policy (that identifies the Provisioning key), the Payer credential (signed by Provisioning) and the microchecks credential (signed by the Payer) to his local KeyNote compliance checker. If the compliance checker authorizes the transaction, the Merchant is guaranteed that Provisioning will allow payment. The correct linkage among the Merchant's policy, the Provisioning key, the Payer key, and the transaction details follow from KeyNote's semantics [20].

If the transaction is approved, the Merchant should give the item to the Payer and store a copy of the microcheck along with the payer credential and associated offer details for later settlement and payment. If the transaction is not approved because the limits in the payer credentials have been exceeded, then, depending on their network connectivity, either the Payer or the Merchant can request a transaction-specific credential that can be used to authorize the transaction. Observe that this approach, if implemented transparently and automatically, provides a continuum between online and offline transactions tuned to the risk and operational conditions.

Periodically, the Merchant will 'deposit' the microchecks (and associated transaction details) it has collected to the *Clearing and Settlement Center (CSC)*. The CSC may or may not be run by the same company as the Provisioning, but it must have the proper authorization to transmit billing and payment records to the Provisioning for the customers. The CSC receives payment records from the various Merchants; these records consist of the Offer, and the KeyNote microcheck and credential from the payer sent in response to the offer. In order to verify that a microcheck is good, the CSC goes through the same procedure as the Merchant did when accepting the microcheck. If the KeyNote compliance checker approves, the check is accepted. Using her public key as an index, the payer's account is debited for the amount of the transaction. Similarly, the Merchant's account is credited for the same amount.

The central advantage of this architecture is the ability to encode risk management rules for micropayments in user credentials. Other electronic systems have focused on preventing fraud and failure, rather than on managing it. In many cases with such systems, the prevention mechanisms can be too expensive for micropayments, making the risk management approach particularly attractive.

2.4 *Sequence of Operations in WebSOS*

To illustrate the use of the WebSOS architecture by servers and clients, we describe the steps both sides must undertake to protect their communication channel. For simplicity, we omit the micropayment component in this description; we describe the steps involved in a client using the micropayment scheme following this discussion.

- A site (target) installs a filter on a router in its immediate vicinity and then se-

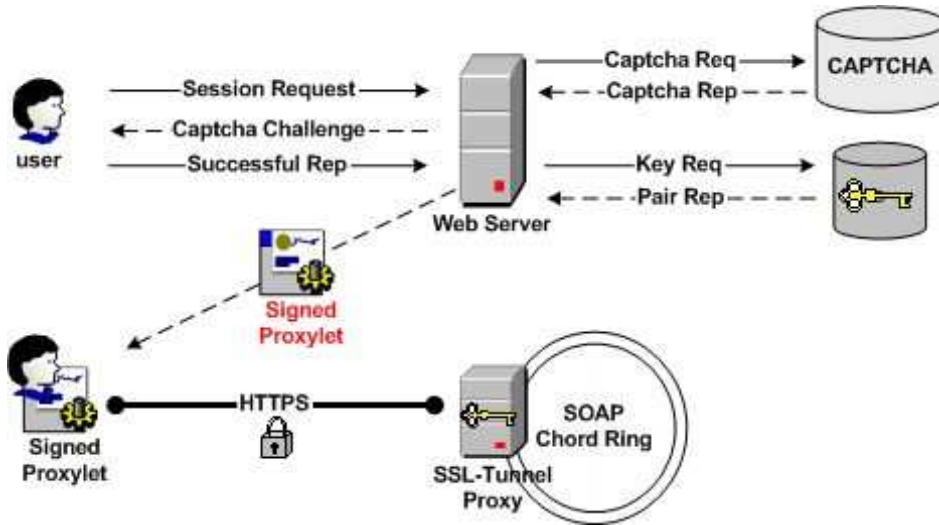


Fig. 5. WebSOS client session initiation diagram.

lects a number of WebSOS nodes to act as “secret servlets” that are allowed to forward traffic through the filter to the target site. Routers at the perimeter of the site are instructed to only allow traffic from these servlets to reach the internal of the site’s network. These routers are powerful enough to do filtering using only a small number of rules on incoming traffic without adversely impacting their performance. In order to make guessing the identity of a secret servlet for a particular target harder for the attacker, the filtering mechanism uses packet fields with potentially high entropy. For example, only GRE [21] packets from a particular source (the secret servlet) containing a specific 32-bit value in the GRE Key field [22]. An attacker trying to slip attack traffic through the filter must guess not only the current servlet’s IP address, but the correct 32-bit key as well. Although we expect 32 bits to be sufficient for this application, we can easily use larger keys to avoid brute-force attacks.

- When a WebSOS node is informed that it will act as a secret servlet for a site (and after verifying the authenticity of the request, by verifying the certificate received during the SSL exchange), it computes the key k for a number of well-known consistent hash functions, based on the target site’s network address. Each of these keys will identify a number of overlay nodes that will act as beacons for that web server, for redundancy purposes.
- Having identified the beacons, the servlets or the target will contact all of them, notifying them of the servlets’ association with a particular target. Beacons will store this information and use it for traffic-forwarding purposes.
- A source that wants to communicate with the target contacts a random overlay node, the Secure Overlay Access Point (SOAP). After authenticating and authorizing the request via the CAPTCHA test, the overlay node securely proxies all traffic from the source to the target via one of the beacons. The SOAP (and all subsequent hops on the overlay) can proxy the HTTP request to an appropriate beacon in a distributed fashion using Chord, by applying the appropriate hash

function(s) to the target’s IP address to identify the next hop on the overlay³. To minimize delays in future requests, the client is issued a short-lived X.509 certificate, bound to the SOAP and the client’s IP address, that can be used to directly contact the proxy-server component of the SOAP without requiring another CAPTCHA test. The overall interaction of clients with the overlay is shown graphically in Figure 5. We shall explain the role of the signed proxylet in Section 4.

This scheme is robust against DoS attacks because if an access point is attacked, the confirmed source point can simply choose an alternate access point to enter the overlay. Any overlay node can provide all different required functionalities (SOAP, Chord routing, beacon, secret servlet). If a node within the overlay is attacked, the node simply exits the overlay and the Chord service self-heals, providing new paths over the re-formed overlay to (potentially new sets of) beacons. Furthermore, no node is more important or sensitive than others — even beacons can be attacked and are allowed to fail. Finally, if a secret servlet’s identity is discovered and the servlet is targeted as an attack point, or attacks arrive at the target with the source IP address of some secret servlet, the target can choose an alternate set of secret servlets.

Use of GRE for encapsulating the traffic between the secret servlet and the filtering router can offer an additional benefit, if we also use transparent proxies and IPsec for packet encapsulation between the proxies (replacing SSL). In that implementation scenario, as far as the target web server is concerned the HTTP/HTTPS connection from the browser was received directly. Thus, any return TCP traffic will be sent directly to the browser’s IP address. Following our discussion in Section 2.5, this asymmetric connection routing will considerably improve the end-to-end latency and reduce the load on the overlay network (less traffic to proxy). While asymmetric routing was once considered potentially harmful, empirical studies show that most of the long-haul traffic (*e.g.*, non-local traffic) over the Internet exhibits high asymmetry [23]. Most of the arguments against this asymmetry arise from the difficulty of configuring packet classification mechanisms, which preclude stateful filtering and required synchronized configuration of multiple nodes (those the traffic may traverse). This would not be a problem in our case, as the asymmetry is exhibited far enough in the network (beyond the filtering router) that the local administrative tasks, such as configuring a firewall, remain unaffected. IPsec and transparent proxying techniques are well-known and (in the case of transparent proxies) widely used, thus we believe such an implementation is feasible. For the purposes of this paper, we decided to implement the straight-forward version of WebSOS; development of the optimized version remains in our plans for future

³ As we shall see in Section 5, the SOAP actually uses the Chord ring to determine the identity of the secret servlet, and then proxies the traffic directly to that node. The SOAP queries in parallel all beacons for a particular destination for the identity of the secret servlet.

work.

In [4], we performed a preliminary analysis using simple networking models to evaluate the likelihood that an attacker is able to prevent communications to a particular target by launching denial of service attacks against overlay nodes, causing them to become unreachable. This likelihood was determined as a function of the aggregate bandwidth obtained by an attacker through the exploitation of compromised systems. The analysis included an examination of the capabilities of static attackers who focus all their attack resources on a fixed set of nodes, as well as attackers who adjust their attacks to “chase after” the repairs that the SOS system implements when it detects an attack. We demonstrated that even attackers that are able to launch massive attacks are very unlikely to prevent successful communication. For instance, attackers capable of launching debilitating attacks against 50% of the nodes in the overlay have roughly one chance in one thousand of stopping a given communication from a client who can access the overlay through a small subset of overlay nodes. For more details on the analysis, see [4].

Finally, it is worth estimating the volume of attacks that a WebSOS-protected system can withstand. Since the Internet (and ISPs’) backbones are well provisioned, the limiting factors are going to be the links close to the target of the attack. Consider the common POP structure used by ISPs shown in Figure 8. The aggregate bandwidth for the POP is on the order of 10 to 20 Gbps. If the aggregate bandwidth of the attack plus the legitimate traffic is less than or equal to the POP capacity, then legitimate traffic will not be affected, and the POP routers can drop the attack traffic (by virtue of dropping any traffic that did not arrive from a SOS secret servlet). Unfortunately, there do not exist good data on DDoS attack volumes; network telescopes [24] tend to underestimate their volume, since they only detect response packets to spoofed attack packets. However, we can attempt a simple calculation of the effective attack bandwidth available to an attacker that controls X hosts that are (on average) connected to an ADSL network, each with 256Kbps uplink capacity. Assuming an effective yield (after packet drops, self-interference, and lower capacity than the nominal link speed) of 50%, the attacker controls $128 \times X$ Kbps of attack traffic. If the POP has an OC-192 (10 Gbps) connection to the rest of the ISP, then an attacker needs 78,125 hosts to saturate the POP’s links. If the POP has a capacity of 20 Gbps, then the attacker needs 156,256 hosts. Although we have seen attack clouds of that magnitude, they are not very common. Thus, a SOS-protected system should be able to withstand the majority of DDoS attacks even when deployed within a single ISP. If attacks of that magnitude are a concern, we can expand the scope of the filtering region to neighboring POPs of the same ISP (and their routers); this would increase the link capacity of the filtered region significantly, since each of the neighboring POPs see only a fraction of the attack traffic. Our discussion is not meant as a proof of security against DDoS attacks, but as an exploration of the limits of our mechanism. Additional work is needed to determine the practical limits of the system, although we are encouraged by our findings to date.

2.5 Forwarding Specifics

WebSOS uses SSL to provide two layers of security. First, messages are encrypted end-to-end, so that only the end-points of the exchange (user and web-server) can view the data actually being transmitted. Additionally, WebSOS uses SSL over each hop of the overlay as a means of verifying the authenticity of the previous hop, both when a forwarding session is established and for the forwarded traffic, to avoid attacks from adversaries impersonating as legitimate WebSOS nodes. Note that we do not actually need message confidentiality for the internal communications of WebSOS; however, the additional cost of encryption with the RC4 algorithm is small, since we need to provide peer and data authentication (we shall see the total overhead of WebSOS in Section 5). No special functionality is required by the overlay nodes to perform these tasks; the user browser simply has to be supplied with the appropriate certificate(s) from the WebSOS administrator.

In the original SOS architecture, the path established from the user to the target through the overlay was unidirectional. Traffic in the reverse direction could also traverse the overlay, by reversing the roles of user and target. In that case, the path taken by requests and responses would be different. Alternatively, traffic from the target to the user could be sent directly (without using the overlay); this is usually not a problem, since most communication channels are full-duplex and, in the event of a DDoS attack, only the downstream portion (to the target) is congested. An additional benefit of this asymmetric approach is reduced latency, since most client/server traffic (especially in web environments) is highly asymmetric (*i.e.*, clients receive a lot more information than they transmit). This was possible because routing decisions in SOS are made on a per-packet basis.

In WebSOS, routing decisions are made on a per-connection basis. Any subsequent requests over the same connection (when using HTTP 1.1) and any responses from the web server can take the reverse path through the overlay. While this makes the implementation simpler, it also introduces increased latency, as the bulk of the traffic will also traverse the overlay. We give some thoughts on how to address this issue in Section 5.

2.6 Pay-Per-Use Mechanism

We now examine the use of a payment mechanism in our architecture. We begin by describing the necessary hardware and software a service provider needs in order to deploy our mechanism, and then examine its implications for clients.

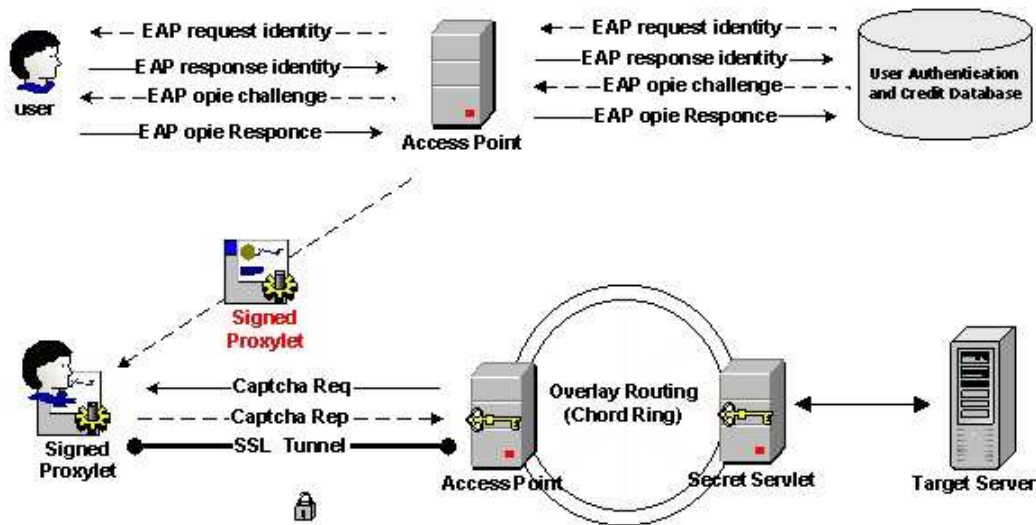


Fig. 6. Pay-per-use DoS protection system operation overview. The user is connected to an access point which in turn authenticates the user credentials and issues an X.509 certificate and a signed proxylet that allows the user to connect securely to the web service for a limited amount time. The Access Point can act as a proxy for the EAP authentication between the authentication database and the user.

2.6.1 ISP Provisioning

The ISP first creates an overlay network of WebSOS access points ('servlets'). Per our previous discussion, the routers at the perimeter of the site are instructed to allow traffic only from these servlets to reach the interior of the site's network. These routers are powerful enough to do filtering using only a small number of rules on incoming traffic without adversely impacting their performance.

For a payment scheme, we chose the OTPchecks system because of its inherent flexibility to accommodate different services and its ability to interoperate with a distributed system like WebSOS. Refer to the roles presented in the OTPchecks functional description, in Figure 4; the Payer is the client connecting to the access points, the Merchant is the ISP providing the DoS protection service, and the web service provider (Target) is the clearing entity. The web service provider controls the usage of the service provided via the ISP's network by having the access points delegate payment credentials to each of the clients. In this manner, the service payment can be charged either to the client or to the web service provider. The ISP, using the same transaction information, charges the site providing the web service. The web service itself may charge the user at the same or even a higher rate for the DoS protection and possibly for other Internet commodities (bandwidth, time *etc.*) using the data presented by the access points. The overall system is presented in Figure 6.

2.6.2 *Buying One-Time Coins*

Whenever a new client host wants to access a service that the ISP protects from DoS attacks, the access point attempts to run the Extensible Authentication Protocol (EAP) [25] Over LAN (EAPoL) protocol [26,27] with the client. The status of the client is kept unauthenticated as long as the client fails to authenticate through EAPoL. In our case, we provide unauthenticated clients limited access so that they can buy one-time “coins” (OTC), modeled after One-Time Passwords (OTP) [28], used for the actual EAPoL level authentication (see below).

2.6.3 *Using One-Time Coins*

Once the Client has acquired a set of one-time coins, it runs the standard EAPoL protocol towards the local access point. The protocol run is illustrated in Figure 6.

Upon connection, the access point requests a user identifier from the client. The client answers with a string that identifies the microcheck used for buying the one-time coins, and the web service the coins were bought for. This allows the access point to contact the correct back-end authenticator, the web service provider (Target). The microcheck fingerprint identifies the relevant unused OTC pile.

Once the back-end authenticator receives the identity response, it checks the OTC pile and sends a request for the next unused one-time password, *i.e.*, an one-time coin. The Client responds with the next unused coin, H_{i+1} . The back-end authenticator checks the coin, records it as used, and replies with an EAP SUCCESS message. As the access point receives the EAP SUCCESS message from the back-end authenticator, it changes the status of the client into authenticated, and passes the message to the client. Shortly before the OTC is used up, the back-end authenticator sends a new request and a GTT to the client.

For the client to continue, it has to reply with the next OTC, and the user must answer correctly the CAPTCHA challenge. This gives us the ability to have a strong protection against malicious code, such as a worm or a zombie process, using a user’s micropayment wallet. The lifetime of a coin can be easily configured by the service provider. We expect to prompt the user with a CAPTCHA challenge every 30 to 45 minutes, depending on the service. The client needs to both pay and pass the CAPTCHA challenge before being issued an X.509 certificate (per our discussion in Section 2.2), which authorizes it to access the WebSOS overlay.

On the other hand, if the client does not want to continue access for any reason, it simply does not respond to the request. Thus, if the client goes off-line, the access point automatically changes the status of the client’s address into unauthenticated once the coin has been used up.




Fig. 7. **Overlay nodes serving regions of a coordinate-space.**

3 Simulation

To understand the impact of the overlay network on the routing of packets between the source and target nodes, we have applied the WebSOS algorithm to two models of ISP networks [29]. One model, indicative of a U.S. topology, is derived from AT&T’s U.S. backbone network. The other, indicative of a European topology, is derived from Worldcom’s (now MCI’s) European backbone network. Remote access points were excluded from the AT&T model, as were connections from Worldcom’s European POPs to points outside the geographical area. For each model, two algorithms for routing traffic through the overlay were tested, one based on Chord, which uses a random ordering of the overlay nodes, and a heuristic variation of CAN that uses geographical ordering of the overlay nodes. In both cases, we tested variations on how the beacons and servlets were chosen in relation to each other, the target, and the source, *e.g.*, requiring some minimum distance between the servlet and target.

We first give a brief description of CAN [30], and then discuss the specifics of the simulation environment, such as ISP structure, the distribution of overlay nodes across ISP Points of Presence (POPs), and the selection strategies for beacons and secret servlets.

3.1 CAN

Like Chord, CAN uses a hash function to map overlay nodes to identifiers. However, a CAN identifier maps a node to a region within a d -dimensional space. Each overlay node contains a table of overlay nodes responsible for neighboring areas in the coordinate space. As shown in Figure 7, overlay node 7 would contain pointers to nodes 3, 6, 8, and 11. In its basic form, CAN does not assume any relationship between node positions of the coordinate space and their geographical positions

in the real world. A variation suggested in [30] that assigns positions within the coordinate space being representative of the geography provided the basis for the heuristic used in the model.

3.2 Network Layout

A POP-level representation of the ISP was used, where each POP is assumed to consist of a hierarchy of routers as shown in Figure 8. At the top level are routers with links to other POPs. At the lowest level are links to client networks.




Fig. 8. ISP POP structure used in the simulation.

Latencies between POPs were estimated from a subset of known latencies. Distances between POPs were estimated using airline miles. Three routers were included at the second level and twelve at the lowest level of each POP; however, for the statistics computed, the exact number of routers within a POP was not relevant, only the latency from the time a packet entered a router within a POP to the time it left the POP was needed.

The model assumes that there is ample bandwidth between POPs and that the choke points are the links to clients. All latencies and distances to clients to their local

POP are assigned the same value.

There were 19 POPs in the US model and 18 in the Europe model. Overlay nodes participating in the overlay were evenly distributed across POPs, meaning each POP served the same number of client nodes eligible to be overlay nodes. In the cases where servlets and beacons were randomly chosen, this allowed each POP to be equally likely to have a client site that was a servlet or beacon. In the cases where the servlet and beacon nodes were not randomly chosen, there were more eligible nodes per POP than utilized and the even distribution did not impact selection. A node was not allowed to serve more than one purpose for a specific source-target pair, for example, a node could not be both a beacon and a servlet for the same target. Removing the restriction would result in shorter routes on average because some scenarios tested would pick the same node for both the servlet and beacon.

In each case, two client nodes served by each POP were included in the overlay. Since each source / target pair was tested individually, at most two nodes per POP would be selected to serve the functions of beacon and servlet. When ordering the overlay nodes according to the geographic heuristic described below, designating more than two nodes per POP could only change a route between a source and target by possibly passing through a different client on a given POP. When ordering the overlay nodes randomly and using Chord as the routing algorithm for the overlay, the probability that a client on a specific POP was picked as a beacon or servlet, or was at a certain position in the overlay impacted the route. Since it was assumed overlay eligible nodes were evenly distributed across all POPs, having 2 versus 100 overlay nodes per POP would not impact the probabilities and thus would not affect the results. The node for the source was chosen to be a client on the same POP as the source. The impact due to it being served by a different POP than the source would be to add the cost of the normal route between the source and SOAP to the cost of the route between the SOAP and target.

3.3 Routing Algorithms

In WebSOS, traffic from a source to a target utilizes a route which contains the following sequence of nodes in order: source, access point, beacon, servlet and target. Normal routing is used to reach the SOAP. Also, since the beacon knows the specific servlet for the target, and the servlet knows the location of the target, normal routing is used between the beacon and servlet, and between the servlet and target. An overlay route is used between the SOAP and beacon. The increase in the route length over that of the normal route between the source and target is due not only to the requirement that the route pass through specific nodes, but also due to the need to route through an overlay network between the SOAP and beacon as opposed to using the normal route between the two nodes. For normal routing, each node in the model contained a routing table populated via Dijkstra's

algorithm, using minimum hops as the criteria for shortest path. Each node in the overlay network also contained a table with the destination address and overlay node id of a subset of overlay nodes. The table was populated based on the routing algorithms described below.

A routing algorithm for use in overlays is required to send traffic between the SOAP and beacon. The Chord algorithm was utilized in the first set of experiments. The overlay nodes were randomly ordered. The tables within each overlay node were populated using the method described previously involving powers of 2. The size of a node's table is $O(\log n)$, where n is the size of the overlay.

The second set of experiments used a heuristic which divided the POPs into geographical areas. This method is based on modifications suggested to the basic algorithm for CAN. For a specific area, A, a node n_A was chosen as the area's representative. Each n_A was an entry in each overlay node's table. In addition, if n_i is an overlay node in area A, n_i 's table would include entries for each n_j in A, $i \neq j$. Thus an overlay node maintained pointers to every other overlay node in the same geographical area and to one overlay node in each other geographical area. For an overlay of size n , the size of a node's table is $O(n/\#areas) + \#(areas)$, which is $O(n/\#areas)$ when n is large compared to the number of areas, assuming each area contains roughly the same portion of overlay nodes. The US model involved 6 areas, one contained 2 POPs and the other contained 3 or 4 POPs each. The Europe model contained 4 areas with 4 to 5 POPs each.

3.4 Beacon/Servlet Selection Scenarios

Seven source-target pairs were chosen in each of the two models. They were selected to represent a variation in source-target relations. Factors considered when selecting the pairs included the distance between cities, whether they were served by neighboring POPs and the level of connectivity for the POP. In all cases a servlet and beacon for a specific target were not permitted to be the same node and neither could serve as a SOAP .

For each model and each routing algorithm, the normal route between each source-target pair was computed then the following four scenarios were tested on each pair. In the scenarios, minimizing the number of hops refers to the number of hops as calculated by normal routing.

- (1) Randomly select the servlet and beacon (100 trials per source-target pair were run).
- (2) Select the servlet to minimize the number of hops between the servlet and target, then select the beacon to minimize the number of hops between the beacon and servlet, with the restriction that the servlet and beacon not be served by the same POP.

- (3) Select the servlet to minimize the number of hops between the servlet and target, then select the beacon to minimize the number of hops between the beacon and source.
- (4) Select a servlet randomly from those approximately X miles from the target then select a beacon randomly from those approximately X miles from the servlet, where X was 1000 in the US model and 500 in the Europe model. In the case of the Europe model, a few POPs did not have neighbors within this distance, in which case the next closest available overlay node was used.

The first scenario was used to obtain an understanding of the impact when no selection criteria was utilized for the servlet and beacon. This would be the simplest version to implement. The second and third scenarios were aimed at keeping the intermediate nodes in the route near the end points to determine if the route between the source and target would then be similar to the normal route. These two scenarios using minimum distance instead of hops were tested on the US version, but the results were not noticeably different from the scenarios using hops. The fourth scenario was used to understand the impact of selecting the servlet and beacon so they would be served by different POPs than the target, which may be desired for diversity, but at the same time guaranteeing they would be relatively close in an attempt to avoid an unnecessarily long route.

Table 1

Average ratio: latency with WebSOS vs. normal routing.

	US	US	Europe	Europe
	Chord	CAN	Chord	CAN
scenario				
1 random selection	4.51	4.16	5.69	4.11
2 min hops	3.45	2.4	3.25	2.54
3 min hops	7.19	1.75	6.77	1.74
4 diversity	5.18	4.08	5.6	2.88

3.5 Results

Results are presented in terms of the ratio of the measurement for the WebSOS route to that of the normal route between the source and target. The measurements are for one direction only, source to target, and are averaged over 100 simulation runs. Table 1 shows the ratio of the latency using WebSOS to the latency expected when using normal routing. The scenario number corresponds to the previous list. These were averaged over all source-target pairs. The worst case from all source-target pairs is shown in Table 2. Table 3 indicates the increase in the number of ISP POPs involved in a route compared to that of the normal route.

Table 2

Worst-case ratio: latency with WebSOS vs. normal routing.

	US/Chord	US/CAN	Europe/Chord	Europe/CAN
scenario				
1 random selection — worst individual source-target average over 100 trials	8.76	6.05	8.05	5.81
2 min hops	7.57	3.76	4.74	3.26
3 min hops	10.9	2.14	11.29	2.14
4 diversity	10.57	6.24	8.1	3.57

Table 3

Numbers of POPs in WebSOS routing vs. normal routing.

	US/Chord	US/CAN	Europe/Chord	Europe/CAN
scenario				
1 random selection — worst individual source-target average over 100 trials	4	3	4	2.5
2 min hops	2	1.5	2	1.5
3 min hops	5	1	4.2	1
4 diversity	3.5	2.5	4.2	2

When using scenario 3 with the geographic heuristic, the servlet was always selected from a node on the same POP as the target and the beacon was selected from a node on the same POP as the source and SOAP because there were eligible nodes at every POP. This resulted in the WebSOS route being identical to the normal route with the addition of a few detours to clients within the first and last POPs in the route, thus it was expected to produce the best results in terms of latency.

The results reported for random selection are averaged over 100 trials run per source-target pair. The actual increase in latency may be much higher depending on the specific servlet and beacon chosen. The greatest increase occurs when the source and target are close together. The overlay route may involve points geographically far from the source and target, turning a normally short route into one that may traverse every POP in the ISP at least once. Among all trials involving random selection, the worst case in the Europe model was an increase in latency 15 times that of the normal route between London and Paris when using Chord and 9.5 times when using the geographical heuristic. In the US model, the worst case also involved a latency 15 times normal between NY and Philadelphia when using Chord and 8.86 times when using the geographical heuristic. For NY to Philadelphia, the worst case increase using the geographical heuristic is approximately the

same as the average (8.76) when using Chord. The worst cases from all trials involved latencies of 378ms using Chord and 230ms using the geographical heuristic.

The number of POPs serves as a measure of the complexity of the route but does not necessarily imply a physically long route because several POPs may be geographically close. In scenario 3, the beacon would be selected on the same POP as the SOAP. The ratio for scenario 3 using Chord is high due to a couple of source-target pairs in which the beacon's overlay id was just prior to that of the SOAP's id, resulting in routing through several overlay nodes in the path between the SOAP and beacon.

When using Chord, other variations for populating the overlay node's tables using powers of 3 and $i + x_j$, where x_j is the j^{th} number in the Fibonacci series, for $j = 3, 4, 5, \dots$, were tested on a subset of source-target pairs but had no noticeable impact on the length of the route between the SOAP and beacon. A geographic ordering of the overlay nodes was also tested while maintaining the Chord routing. Nodes that were geographically close were assigned IDs placing them close together on the overlay network. While this shortened the route in cases where nodes X and Y were physically close, a packet was being routed from X to Y using the overlay and X was assigned a lower overlay id than Y; it resulted in a worst case scenario when Y was assigned the overlay id just prior to X's because the packet would route to $O(\log n)$ overlay nodes before reaching the one that knew about X.

4 Implementation

While the simulation results are encouraging, we felt that experimentation in real networks was necessary to validate our approach. To that end, we developed a prototype of WebSOS, consisting of three main modules. The components are a communications module, a WebSOS routing module, and an overlay routing module running on each node in the WebSOS overlay. The interactions of these are shown in Figure 9.

The communications module is responsible for forwarding HTTP requests and responses among the nodes in the WebSOS overlay. When a new proxy request (in the form of a new TCP connection) is received, the communications module calls the WebSOS routing module with the target's destination address to obtain the address of the next hop in the overlay. It then opens a new TCP connection to that node and relays the received HTTP request. Any traffic received in the opposite direction (*i.e.*, the HTTP responses and web content) are relayed back to the source. Authentication of the requesting node by the access point (SOAP) and by internal nodes is accomplished through SSL. Authorized users and WebSOS overlay nodes are issued X.509 [12] certificates signed by the SOAP, once the user has succeeded in the CAPTCHA authentication.




Fig. 9. **Software modules for the WebSOS implementation.**

The main *WebSOS routing module* receives requests from the communications module and responds with the IP address of the next node in the WebSOS overlay to which the request should be forwarded. The module first checks whether the current node serves a specific purpose (*i.e.*, whether is it a beacon or secret servlet for that target). If the node serves no such purpose, the module calls the overlay routing module to determine the next hop in the WebSOS overlay and passes the reply onto the communications module. Presently, the WebSOS routing module is initialized with configuration data at startup indicating which nodes serve specific purposes. We are working on an administrative module with increased flexibility to avoid this static provisioning.

The overlay routing module is a general routing algorithm for overlay networks. An implementation of Chord was written for the initial tests. However, this module can be replaced with any other routing algorithm, *e.g.*, CAN [30]. It receives queries containing a destination IP address (the web server's) and responds with the IP address of the next node in the overlay to which the request should be forwarded. For maintenance of its own routing algorithm, the Chord implementation also communicates with other overlay nodes to determine their status, as described in [15].

When a request is issued by the browser, it is tunneled through a series of SSL-

encrypted links to the target, allowing the entire transmission between the requester and target to be encrypted. The SSL connections between WebSOS nodes are dynamically established, as new requests are routed. One problem we ran into while developing the WebSOS prototype is that web browsers do not provide support for the actual proxy request to be encrypted. To solve this problem, we wrote a port forwarder that runs on the user's system, accepts plaintext proxy requests locally, and forwards them using SSL to the access point node. This is implemented as a Java applet that runs inside the browser itself. The Java applet is not considered part of the WebSOS overlay and is not trusted to perform any access control decisions; it is simply a "helper" application.

Thus, to use WebSOS, an authorized user simply has to access any SOAP, successfully respond to the CAPTCHA challenge, download the applet, and set the browser's proxy settings to the localhost, as shown in Figure 5. Java applets typically cannot communicate with any host other than the one they were downloaded from, but this is not a problem in our case. If the user is successful in his/her reply, then the web server connects to a DBMS system (local or remote) and associates a pair of RSA keys (a private key and a certificate) with the host. This set of keys are unique per IP and have an expiration time that can be configured by the system administrator. The user is prompted to download a signed applet that runs locally using one browser window and contacts the Web Server via a temporary HTTPS connection to fetch the X.509 certificate. Since we expect overlay membership (including SOAPs) to change infrequently, we can publish the list of available nodes in several different venues (*e.g.*, DNS, "flat" text file available from several locations, *etc.*), exploiting the ease of replicating static content to prevent DoS attacks on that part of the system. Using an outdated list of SOAPs only increases the initial latency for contacting WebSOS; once a functioning SOAP is located, the applet can download the most recent membership list.

The applet then starts listening for HTTP/HTTPS connections on a local port (*e.g.*, 8080) and establishes an SSL-tunnel connection with the proxy server running on the SOAP (or elsewhere, since the signed applet has the ability to connect to any server by changing the Java Policy files on the users' machine). The proxy server matches the X.509 certificate and the IP from client to the private key obtained from the DBMS system and allows the connection to be proxied. The only imposition on the user is that he/she must change the Proxy settings of the local browser to point to the socket that listens for the applets. The applet may also directly proxy a connection to the webserver and, if that fails, switch to using WebSOS. The applet is also responsible for determining which SOAP to use, and switching to a different SOAP if the one in use becomes unresponsive (*e.g.*, due to an attack).

Initial prototyping of the communications module used Apache, whose proxy module was modified to query the routing module for the next hop. This worked well when unencrypted HTTP requests were issued by the browser. However, when we encountered the requirement for end-to-end authentication and encryption, we

changed the implementation to use a stand-alone proxy server instead of Apache.

comment We intend to expand the implementation to include additional modules addressing the administration and maintenance of the overlay. A centralized administration module will be used to set node characteristics in real time, including assigning specific roles (beacon, SOAP, secret servlet) to nodes, and changing the operational status of nodes. A maintenance module running on each node will check the status of all nodes in the WebSOS overlay and provide updates to both the main and overlay routing modules in order for routing to be adjusted. This module will also serve as the interface to centralized administration by receiving updates regarding a node's function and status, and passing the updates to the appropriate routing module.

To improve performance, we also explored an adaptation of the initial approach: rather than transporting the request and response through the full overlay network, only routing information travels through the overlay. As before, the requester makes a proxy request to the SOAP. At that point, the SOAP sends a UDP message into the overlay, specifying the target. The message is routed to the beacon, which responds directly to the SOAP with information on the secret servlet for that target. The SOAP then connects to the servlet, which proxies the request as before, in effect creating a *shortcut* through the overlay.

The SOAP caches the servlet information for use in future requests. That information is timed out after a period of time to allow for changes to propagate correctly. The same basic UDP protocol is used by servlets to announce their presence to (and periodically update) the beacons for the various targets.

5 Experimental Evaluation

In order to quantify the overhead associated with use of WebSOS, we created a simple topology running on the local network (100 Mbit fully-switched Ethernet). For our local-area network overlay, we used 10 commodity PCs running Linux Redhat 7.3. We measured the time-to-completion of https requests. That is, we measured the elapsed time starting when the browser initiates the TCP connection to the destination or the first proxy, to the time all data from the remote web server have been received. We ran this test by contacting 3 different SSL-enabled sites: *login.yahoo.com*, *www.verisign.com*, and the Columbia course bulletin board web service (at <https://www1.columbia.edu/sec/bboard>). For each of these sites, we measured the time-to-completion for a different number of overlay nodes between the browser and the target (remote web server). Our measurements sum up all WebSOS overheads, including network latency, software processing at each overlay node, and encryption (specifically, double SSL encryption, as discussed in Section 2.5). Note that a measurement involving n nodes implies a Chord

overlay of 2^n nodes, since for an overlay of size x , it takes approximately $\log(x)$ hops to reach a destination. Thus, our experiments with 10 nodes represent an overlay with 1024 nodes.

The browser was located on a separate ISP. The reason for this configuration was to introduce some latency in the first-hop connection (from the browser to the SOAP), thus simulating (albeit using a real network) an environment where the browsers have slower access links to the SOAPs, relative to the links connecting the overlay nodes themselves (which may be co-located with core routers). By locating all the overlay nodes in the same location, we effectively measure the aggregate overhead of the WebSOS nodes in the optimal (from a performance point of view) case.

Table 4 shows the results for the case of 0 (browser contacts remote server directly), 1, 4, 7, and 10 overlay nodes. The times reported are in seconds, and are averaged over several HTTPS GET requests of the same page, which was not locally cached. For each GET request, a new TCP connection was initiated by the browser. The row labeled “Columbia BB (2nd)” shows the time-to-completion of an HTTPS GET request that uses an already-established connection through the overlay to the web server, using the HTTP 1.1 protocol.

As the figure shows, WebSOS increases the end-to-end latency between the browser and the server by a factor of 2 to 3. These results are consistent with our simulations of using SOS in an ISP topology, where the latency between the different overlay nodes would be small, as discussed in Section 3. The increase in latency can be primarily attributed to the network-stack processing overhead and proxy processing at each hop. It may be possible to use TCP splicing [31] or similar techniques to reduce connection handling overhead, since WebSOS performs routing on a per-request basis. Also, in the experiments we ran, we did not make use of the asymmetric routing option possible with the use of GRE as both a filtering and an encapsulation mechanism, as discussed in Section 2.4.

Table 4

Latency (in seconds) when contacting various SSL-enabled web servers directly and with different numbers of (intermediate) overlay nodes over the local-Ethernet network.

Server	Direct	1 node	4 nodes	7 nodes	10 nodes
Yahoo!	1.39	2.06	2.37	2.79	3.33
Verisign	3.43	4.22	5.95	6.41	9.01
Columbia BB	0.64	0.86	1.06	1.16	1.21
Columbia BB (2nd)	0.14	0.17	0.19	0.20	0.25

Furthermore, there is an SSL-processing overhead for the inter-overlay communications. A minor additional cryptographic overhead, relative to the direct access case, is the certificate validation that the SOAPs have to perform, to determine the client’s authority to use the overlay, and the SSL connection between the proxy

Table 5

Latency (in seconds) when contacting various SSL-enabled web servers directly and with different numbers of (intermediate) overlay nodes using the PlanetLab network.

Server	Direct	1 node	4 nodes	7 nodes	10 nodes
Yahoo!	1.39	3.15	5.53	10.65	14.36
Verisign	3.43	5.12	7.95	14.95	22.82
Columbia BB	0.64	1.01	1.45	3.14	5.07
Columbia BB (2nd)	0.14	0.23	0.28	0.57	0.72

running on the user’s machine and the SOAP. As shown in [32], such overheads are typically dominated by the end-to-end communication overheads. Use of cryptographic accelerators can further improve performance in that area. One further optimization is to maintain persistent SSL connections between the overlay nodes. However, this will make the task of the communication module harder, as it will have to parse HTTP requests and responses arriving over the same connection in order to make routing decisions.

Table 5 shows the same experiment using PlanetLab [8], a wide-area overlay network testbed. The PlanetLab nodes are distributed in academic institutions across the country, and are connected over the Internet⁴. We deployed our WebSOS proxies PlanetLab and ran the exact same tests. Naturally, the direct-contact case remains the same. We see that the time-to-completion in this scenario increases by a factor of 2 to 10, depending on the number of nodes in the overlay. In each case, the increase in latency over the local-Ethernet configuration can be directly attributed to the delay in the links between the WebSOS nodes. While the PlanetLab configuration allowed us to conduct a much more realistic performance evaluation, it also represents a worst-case deployment scenario for WebSOS: typically, we would expect WebSOS to be offered as a service by an ISP, with the (majority of) WebSOS nodes located near the core of the network. Using PlanetLab, the nodes are distributed in (admittedly well-connected) end-sites. We would expect that a more commercial-oriented deployment of WebSOS would result in a corresponding decrease in the inter-overlay delay. On the other hand, it is easier to envision end-site deployment of WebSOS, since it does not require any participation from the ISPs.

Finally, while the additional overhead imposed by WebSOS can be significant, we have to consider the alternative: no web service while a DoS attack against the server is occurring. While an increase in end-to-end latency by a factor of 5 (or even 10, in the worst case) is considerable, we believe it is more than acceptable in certain environments and in the presence of a determined attack.

Table 6 shows the results when the shortcut implementation was tested on the Plan-

⁴ In fact, the majority of PlanetLab nodes are inside Internet2. This can give somewhat better results than the current Internet, since connectivity is better over Internet2.

etLab testbed, with a graphical depiction in Figure 10. This variant provides significant performance improvements, particularly on subsequent requests for the same site, because of the caching. To simulate the effects of an attack on individual nodes in the overlay, we simply brought down specific nodes. The system healed itself within 10 seconds.

Table 6

Latency (in seconds) when contacting various SSL-enabled web servers directly and while using the shortcut implementation of the WebSOS system. The testing was performed on a 76 node subset of the PlanetLab testbed using the Chord overlay. The hops to the beacon ranged from 4 to 8 and did not have a significant effect on latency. The *cached* column refers to subsequent requests using the same SOAP, whereupon the Secret Servlet information has been cached.

Server	Direct	Original Request	Cached Re-requests
Yahoo!	1.39	4.15	3.67
Verisign	3.43	7.33	6.77
Columbia BB	0.64	3.97	3.43
Columbia BB (2nd)	0.14	0.55	0.56




Fig. 10. WebSOS Latency overhead for different SSL-enabled services when using the shortcut routing mechanism.

To complete the overhead analysis, we measured the number of public key verifications an access point can perform, which indicates how many microchecks it can validate in unit time. We used a 3GHz Pentium4 processor machine running Linux with the OpenSSL V 0.9.7c library for the measurements. The contribution of the micropayment system to the overall system latency overhead is minimal even, when we issue 1024-bit RSA certificates for the client credentials, as shown in Table 7.

These measurements show that the impact of the user verification process on the access points is minimal: since users only need to use a microcheck every half hour or so, a single node can handle 18 million users per hour, even without hardware accelerators [33].

Table 7

Signing and verification times for 1024-bit RSA keys.

Sign	Verify	Signatures/sec	Verifications/sec
0.0037 secs	0.0002 secs	270.0	5055.9

6 Related Work

As a result of its increased popularity and usefulness, the Internet contains both interesting targets and enough malicious and ignorant users that DoS attacks are simply not going to disappear on their own; indeed, although the press has stopped reporting such incidents, recent studies have shown a surprisingly high number of DoS attacks occurring around the clock throughout the Internet [24]. Worse, the Internet is increasingly being used for time-critical applications. A further compounding factor is the susceptibility of the basic protocols (*i.e.*, IP and TCP) to denial of service attacks [34–36].

The need to protect against or mitigate the effects of DoS attacks has been recognized by both the commercial and research world. Some work has been done toward achieving these goals, *e.g.*, [1–3,37–41]. These mechanisms focus on detecting the source of DoS attacks in progress and then countering them, typically by “pushing” some filtering rules on routers as far away from the target of the attack (and close to the sources) as possible. The motivation behind such approaches has been twofold: first, it is conceptually simple to introduce a protocol that will be used by a relatively small subset of the nodes on the Internet (*i.e.*, ISP routers), as opposed to requiring the introduction of new protocols that must be deployed and used by end-systems. Second, these mechanisms are fairly transparent to protocols, applications, and legitimate users. Unfortunately, these reactive approaches by themselves are not adequate, since large-scale coordination across multiple administrative domains is not always practical.

The D-WARD system [42] monitors outgoing traffic from a given source network and attempts to identify attack traffic by comparing against models of reasonable congestion control behavior. The amount of throttling on suspicious traffic is proportional to its deviation from the expected behavior, as specified by the model. In COSSACK [43], participating agents at edge networks exchange information about observed traffic and form multicast cliques to coordinate attack suppression. An interesting approach is that of [44], which proposes an IP hop-count-based filter to weed out spoofed packets. The rationale is that most such packets will not have

a hop-count (TTL) field consistent with the IP addresses being spoofed. In [45], the authors use a combination of techniques that examine packet contents, transient ramp-up behavior and spectral analysis to determine whether an attack is single- or multi-sourced, which would help focus the efforts of a hypothetical anti-DoS mechanism.

A variant of the packet marking approaches creates probabilistically unique path-marks on packets without requiring router coordination; end-hosts or firewalls can then easily filter out packets belonging to a path that exhibits anomalous behavior [46]. Although this approach avoids many of the limitations of the pure marking schemes, it requires that core routers “touch” packets (rather than simply switch them), and assumes that the limited resource is the target’s CPU cycles, rather than the available bandwidth (*i.e.*, preventing the DoS attack is “simply” a matter of quickly determining which packets the server should ignore). In our work, we assume that the scarce resource is bandwidth. Collins and Reiter [47] present an empirical analysis of several different anti-DoS techniques (including Pi [46] and Hop-Count Filtering [44]) that use filters near the target of an attack, using traces of real DDoS attacks to simulate the impact of the filters on the attack traffic.

[48] describes filtering out source-spoofed packets inside the Internet core, and discusses the effectiveness of this approach. The authors suggest piggybacking on BGP to propagate the necessary information. DDoS attacks using real IP addresses are not affected by this scheme.

[49] proposes using Class-Based Queuing on a web load-balancer to identify misbehaving IP addresses and place them in lower priority queues. However, many of the DDoS attacks simply cause congestion to the web server’s access link.

Another approach to mitigating DoS attacks against information carriers is to massively replicate the content being secured around the entire network, *e.g.*, [50]. To prevent access to the replicated information, an attacker must attack all replication points throughout the entire network — a task that is considerably more difficult than attacking a small number of, often co-located, servers. Replication is a promising means to preserve information that is relatively static, such as news articles. However, there are several reasons why replication is not always an ideal solution. For instance, the information may require frequent updates complicating large-scale coherency (especially during DoS attacks), or may be dynamic by its very nature (*e.g.*, a live web-cast). Another concern is the security of the stored information: engineering a highly-replicated solution without leaks of information is a challenging endeavor. Finally, much of the web-served content is semi-dynamic in nature, *i.e.*, it is created on the fly using a database back-end and business-logic middleware such as PHP. This trend decreases the scope for static content replication as a comprehensive solution to DDoS.

An extension of the ideas of SOS [4,51,52] appears in [53]. There, the two main

facets of the SOS architecture: filtering and overlay routing, are explored separately, and several alternative mechanisms are considered. It is observed that in some cases, the various security properties offered by SOS can still be maintained using mechanisms that are simpler and more predictable. However, some second-order properties, such as the ability to rapidly reconfigure the architecture in anticipation of or in reaction to a breach of the filtering identity (*e.g.*, identifying the secret servlet) are compromised. In most other respects, the two approaches are very similar.

Gligor [54] proposes the use of a server that can produce tickets at line speeds. Clients must obtain a ticket from this server before they are allowed to access a protected service. The approach is geared towards application-level DoS protection, with some other mechanism, such as SOS or Pushback, used to address network-level DoS attacks. Anderson *et. al* [55] subsequently proposed a similar system for use at the network layer of an Internet-like architecture designed with a clean slate, assuming a distributed token server architecture and rate-limiting/filtering traffic on routers based on these tokens. Another similar idea appears in [56].

The NetBouncer project [57] considers the use of client-legitimacy tests for filtering attack traffic. Such tests include packet-validity tests (*e.g.*, source address validation), flow-behavior analysis, and application-specific tests, including Graphic Turing Tests. However, since their solution is end-point based, it is susceptible to large link-congestion attacks. [58] is the first system to create stateless flow filtering by having each router add “capabilities” to packets that traverse them; the receiver of these packets is then responsible for sending these capabilities to its peers, which will allow them to send traffic at higher rates (privileged traffic). Unprivileged traffic is limited to a fraction of the available bandwidth; thus, although a DoS attack can prevent new connections from being established (by overloading the control channel used to communicate these capabilities), existing connections will be unharmed.

[59] examines several different DDoS mitigation technologies and their interactions. Among their conclusions, they mention that requiring the clients to do some work, *e.g.*, [60], can be an effective countermeasure, provided the attacker does not have too many resources compared to the defender. Wang and Reiter [61] introduced the idea of a puzzle auction as a way to ease some of the practical deployment difficulties, *e.g.*, selecting the appropriate hardness for the puzzles. Their intuition is to let clients bid for the resources by tuning the difficulty of the puzzles they solve. When the server is attacked, legitimate clients gradually increase their bids (puzzle difficulty), eventually bringing the cost outside the adversary’s capabilities. The authors envision combining their scheme with some anti-DoS mechanism that counteracts volume-based attacks [1,46,4].

6.1 *Micropayment Systems*

Although we use a particular micropayment system [62,19], other schemes can also be used, including digital cash systems [63–66], scrip-based micropayments [67,68], and offline micropayment protocols [69–72]. MiniPay [73] is particularly attractive, since it was developed primarily for use with a web browser, with considerable effort gone into the user interface aspect. Risk management is implemented as a decision to perform an online check with the billing server based on the total spending by the customer that day, and some parameter set by the merchant. We believe that general transactional payment schemes [74–77] may prove too heavy-weight for our purposes.

7 **Conclusions**

We presented WebSOS, an architecture that allows legitimate users to access a web server in the presence of a denial of service attack. The architecture uses a combination of Graphic Turing tests, cryptographic protocols for data origin authentication, packet filtering, overlay networks, and consistent hashing to provide service to casual web-browsing users. Furthermore, our architecture is the first pay-friendly DoS protection mechanisms, furnishing ISPs with a better value proposition for deploying anti-DoS systems: a way to turn DoS protection into a commodity. We discussed our prototype implementation, which uses standard web proxying and authentication mechanisms built in all browsers. Our architecture requires no changes to web servers, browsers, or existing protocols.

We conducted a performance evaluation of WebSOS over both a local area network and over the Internet using PlanetLab, a testbed for experimentation with network overlays and similar technologies. Our experiments show that, in a realistic but worst-case deployment scenario, the end-to-end communication latency between browser and server increases on the average by a factor of 7, with a worst case of 10. We also implemented and evaluated a shortcut optimization, which reduced the latency to a factor of 2. These results are consistent with our simulations. We also discussed other optimizations. However, we believe that even at its current level, the overhead imposed is acceptable for many critical environments and applications.

Future work plans include completion and long-term deployment of the WebSOS prototype on PlanetLab, development of the IPsec-enabled prototype that allows for transparent proxying and asymmetric traffic routing for improved performance, and more comprehensive performance measurements, over a longer period of time and for a wider set of users and web sites.

8 Acknowledgements

Alexander Konstantinou's NetCallback was used as a basis for the forwarding code in the communications module. Abhinav Kamra wrote the Chord implementation used for overlay routing.

References

- [1] J. Ioannidis, S. M. Bellovin, Implementing Pushback: Router-Based Defense Against DDoS Attacks, in: Proceedings of the Network and Distributed System Security Symposium (NDSS), 2002.
- [2] D. Dean, M. Franklin, A. Stubblefield, An Algebraic Approach to IP Traceback, in: Proceedings of the Network and Distributed System Security Symposium (NDSS), 2001, pp. 3–12.
- [3] S. Savage, D. Wetherall, A. Karlin, T. Anderson, Network Support for IP Traceback, ACM/IEEE Transactions on Networking 9 (3) (2001) 226–237.
- [4] A. D. Keromytis, V. Misra, D. Rubenstein, SOS: Secure Overlay Services, in: Proceedings of ACM SIGCOMM, 2002, pp. 61–72.
- [5] M. C. Benvenuto, A. D. Keromytis, EasyVPN: IPsec Remote Access Made Easy, in: Proceedings of the 17th USENIX Systems Administration Conference (LISA), 2003, pp. 87–93.
- [6] S. Dietrich, N. Long, D. Dittrich, Analyzing Distributed Denial of Service Tools: The Shaft Case, in: Proceedings of USENIX LISA XIV, 2000.
- [7] J. Ioannidis, S. Ioannidis, A. D. Keromytis, V. Prevelakis, Fileteller: Paying and Getting Paid for File Storage, in: Proceeding of Financial Cryptography (FC) Conference, 2002, pp. 282–299.
- [8] L. Peterson, D. Culler, T. Anderson, T. Roscoe, A Blueprint for Introducing Disruptive Technology into the Internet, in: Proceedings of the 1st Workshop on Hot Topics in Networks (HotNets-I), 2002.
URL citeseer.nj.nec.com/peterson02blueprint.html
- [9] S. A. Crosby, D. S. Wallach, Denial of Service via Algorithmic Complexity Attacks, in: Proceedings of the 12th USENIX Security Symposium, 2003, pp. 29–44.
- [10] A. D. Keromytis, J. Parekh, P. N. Gross, G. Kaiser, V. Misra, J. Nieh, D. Rubenstein, S. Stolfo, A Holistic Approach to Service Survivability, in: Proceedings of the ACM Survivable and Self-Regenerative Systems Workshop, 2003, pp. 11–22.
- [11] S. Kent, R. Atkinson, Security Architecture for the Internet Protocol, RFC 2401, IETF (November 1998).
URL [ftp://ftp.isi.edu/in-notes/rfc2401.txt](http://ftp.isi.edu/in-notes/rfc2401.txt)

- [12] CCITT, X.509: The Directory Authentication Framework, International Telecommunications Union, Geneva (1989).
- [13] S. M. Bellovin, Distributed Firewalls, *login*: magazine, special issue on security (1999) 37–39.
- [14] S. Ioannidis, A. Keromytis, S. Bellovin, J. Smith, Implementing a Distributed Firewall, in: Proceedings of Computer and Communications Security (CCS), 2000, pp. 190–199.
- [15] I. Stoica, R. Morris, D. Karger, F. Kaashoek, H. Balakrishnan, Chord: A Scalable Peer-To-Peer Lookup Service for Internet Application, in: Proceedings of ACM SIGCOMM, 2001.
- [16] D. Karger, E. Lehman, F. Leighton, R. Panigrahy, M. Levine, D. Lewin, Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web, in: Proceedings of ACM Symposium on Theory of Computing (STOC), 1997, pp. 654–663.
URL citeseer.nj.nec.com/karger97consistent.html
- [17] L. von Ahn, M. Blum, N. J. Hopper, J. Langford, CAPTCHA: Using Hard AI Problems For Security, in: Proceedings of EUROCRYPT, 2003.
- [18] G. Mori, J. Malik, Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA, in: Computer Vision and Pattern Recognition CVPR'03, 2003.
- [19] M. Blaze, J. Ioannidis, S. Ioannidis, A. D. Keromytis, P. Nikander, V. Prevelakis, TAPI: Transactions for Accessing Public Infrastructure, in: Proceedings of the 8th IFIP Personal Wireless Communications (PWC) Conference, 2003, pp. 90–10.
- [20] M. Blaze, J. Feigenbaum, J. Ioannidis, A. D. Keromytis, The KeyNote Trust Management System Version 2, RFC 2704 (September 1999).
- [21] D. Farinacci, T. Li, S. Hanks, D. Meyer, P. Traina, Generic Routing Encapsulation (GRE), RFC 2784, IETF (March 2000).
URL <http://www.rfc-editor.org/rfc/rfc2784.txt>
- [22] G. Dommety, Key and Sequence Number Extensions to GRE, RFC 2890, IETF (September 2000).
URL <http://www.rfc-editor.org/rfc/rfc2890.txt>
- [23] L. Amini, H. Schulzrinne, A. Lazar, Observations from Router-level Internet Traces, in: DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling, 2002.
- [24] D. Moore, G. Voelker, S. Savage, Inferring Internet Denial-of-Service Activity, in: Proceedings of the 10th USENIX Security Symposium, 2001, pp. 9–22.
- [25] J. V. L. Blunk, PPP Extensible Authentication Protocol (EAP), RFC 2284, IETF (March 1998).
URL <http://www.ietf.org/rfc/rfc2284.txt>

- [26] IEEE Draft P802.1X/D11: Standard for Port based Network Access Control (March 2001).
- [27] P. Nikander, Authorization and charging in public wlans using freebsd and 802.1x, in: Proceedings of the Annual USENIX Technical Conference, Freenix Track, 2002.
- [28] N. Haller, C. Metz, P. Nesser, M. Straw, A One-Time Password System, RFC 2289, IETF (February 1998).
URL <http://www.ietf.org/rfc/rfc2289.txt>
- [29] D. Cook, Analysis of Routing Algorithms for Secure Overlay Service, Computer Science Department Technical Report CUCS-010-02, Columbia University (April 2002).
URL <http://www.cs.columbia.edu/~library/TR-repository/repor%ts/reports-2002/cucs-010-02.pdf>
- [30] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Shenker, A Scalable Content-Addressable Network, in: Proceedings of ACM SIGCOMM, 2001.
- [31] A. Cohen, S. Rangarajan, J. H. Slye, On the Performance of TCP Splicing for URL-Aware Redirection, in: USENIX Symposium on Internet Technologies and Systems, 1999.
URL citeseer.nj.nec.com/cohen99performance.html
- [32] S. Miltchev, S. Ioannidis, A. D. Keromytis, A Study of the Relative Costs of Network Security Protocols, in: Proceedings of USENIX Annual Technical Conference, Freenix Track), 2002, pp. 41–48.
- [33] A. D. Keromytis, J. L. Wright, T. de Raadt, The Design of the OpenBSD Cryptographic Framework, in: Proceedings of the USENIX Annual Technical Conference, 2003, pp. 181–196.
- [34] C. Schuba, I. Krsul, M. Kuhn, E. Spafford, A. Sundaram, D. Zamboni, Analysis of a Denial of Service Attack on TCP, in: IEEE Security and Privacy Conference, 1997, pp. 208–223.
- [35] L. Heberlein, M. Bishop, Attack Class: Address Spoofing, in: Proceedings of the 19th National Information Systems Security Conference, 1996, pp. 371–377.
- [36] S. Savage, N. Cardwell, D. Wetherall, T. Anderson, TCP Congestion Control with a Misbehaving Receiver, ACM Computer Communications Review 29 (5) (1999) 71–78.
- [37] S. Savage, D. Wetherall, A. Karlin, T. Anderson, Practical Network Support for IP Traceback, in: Proceedings of the 2000 ACM SIGCOMM Conference, 2000, pp. 295–306.
- [38] M. T. Goodrich, Efficient Packet Marking for Large-Scale IP Traceback, in: Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS), 2002, pp. 117–126.
- [39] A. Snoeren, C. Partridge, L. Sanchez, C. Jones, F. Tchakountio, S. Kent, W. Strayer, Hash-Based IP Traceback, in: Proceedings of ACM SIGCOMM, 2001.

- [40] R. Stone, CenterTrack: An IP Overlay Network for Tracking DoS Floods, in: Proceedings of the USENIX Security Symposium, 2000.
- [41] J. Li, M. Sung, J. Xu, L. Li, Large-Scale IP Traceback in High-Speed Internet: Practical Techniques and Theoretical Foundation, in: Proceedings of the IEEE Symposium on Security and Privacy, 2004.
- [42] P. Reiher, J. Mirkovic, G. Prier, Attacking DDoS at the source, in: Proceedings of the 10th IEEE International Conference on Network Protocols, 2002.
- [43] C. Papadopoulos, R. Lindell, J. Mehringer, A. Hussain, R. Govindan, COSSACK: Coordinated Suppression of Simultaneous Attacks, in: Proceedings of DISCEX III, 2003, pp. 2–13.
- [44] C. Jin, H. Wang, K. G. Shin, Hop-Count Filtering: An Effective Defense Against Spoofed DoS Traffic, in: Proceedings of the 10th ACM International Conference on Computer and Communications Security (CCS), 2003, pp. 30–41.
- [45] A. Hussain, J. Heidemann, C. Papadopoulos, A Framework for Classifying Denial of Service Attacks, in: Proceedings of ACM SIGCOMM, 2003, pp. 99–110.
- [46] A. Yaar, A. Perrig, D. Song, Pi: A Path Identification Mechanism to Defend against DDoS Attacks, in: Proceedings of the IEEE Symposium on Security and Privacy, 2003.
- [47] M. Collins, M. Reiter, An empirical analysis of target-resident DoS filters, in: Proceedings of the IEEE Symposium on Security and Privacy, 2004.
- [48] K. Park, H. Lee, On the Effectiveness of Route-based Packet Filtering for Distributed DoS Attack Prevention in Power-law Internets, in: Proceedings of ACM SIGCOMM, 2001, pp. 15–26.
- [49] F. Kargl, J. Maier, M. Weber, Protecting web servers from distributed denial of service attacks, in: World Wide Web, 2001, pp. 514–524.
URL citeseer.nj.nec.com/444367.html
- [50] A. Stavrou, D. Rubenstein, S. Sahu, A Lightweight, Robust P2P System to Handle Flash Crowds, IEEE Journal on Selected Areas in Communications (JSAC) 22 (1) (2004) 6–17.
- [51] D. L. Cook, W. G. Morein, A. D. Keromytis, V. Misra, D. Rubenstein, WebSOS: Protecting Web Servers From DDoS Attacks, in: Proceedings of the 11th IEEE International Conference on Networks (ICON), 2003, pp. 455–460.
- [52] W. G. Morein, A. Stavrou, D. L. Cook, A. D. Keromytis, V. Misra, D. Rubenstein, Using Graphic Turing Tests to Counter Automated DDoS Attacks Against Web Servers, in: Proceedings of the 10th ACM International Conference on Computer and Communications Security (CCS), 2003, pp. 8–19.
- [53] D. G. Andersen, Mayday: Distributed Filtering for Internet Services, in: 4th USENIX Symposium on Internet Technologies and Systems USITS, 2003.

- [54] V. D. Gligor, Guaranteeing Access in Spite of Distributed Service-Flooding Attacks, in: Proceedings of the Security Protocols Workshop, 2003.
- [55] T. Anderson, T. Roscoe, D. Wetherall, Preventing Internet Denial-of-Service with Capabilities, in: Proceedings of the 2nd Workshop on Hot Topics in Networks (HotNets-II), 2003.
- [56] K. Lakshminarayanan, D. Adkins, A. Perrig, I. Stoica, Taming IP Packet Flooding Attacks, in: Proceedings of the Second Workshop on Hot Topics in Networks (HotNets-II), 2003.
- [57] R. Thomas, B. Mark, T. Johnson, J. Croall, NetBouncer: Client-legitimacy-based High-performance DDoS Filtering, in: Proceedings of DISCEX III, 2003, pp. 14–25.
- [58] A. Yaar, A. Perrig, D. Song, An Endhost Capability Mechanism to Mitigate DDoS Flooding Attacks, in: Proceedings of the IEEE Security and Privacy Symposium, 2004.
- [59] W. J. Blackert, D. M. Gregg, A. K. Castner, E. M. Kyle, R. L. Hom, R. M. Jokerst, Analyzing Interaction Between Distributed Denial of Service Attacks and Mitigation Technologies, in: Proceedings of DISCEX III, 2003, pp. 26–36.
- [60] D. Dean, A. Stubblefield, Using client puzzles to protect TLS, in: Proceedings of the 10th USENIX Security Symposium, 2001.
- [61] X. Wang, M. K. Reiter, Defending Against Denial-of-Service Attacks with Puzzle Auctions (Extended Abstract), in: Proceedings of the IEEE Symposium on Security and Privacy, 2003.
- [62] M. Blaze, J. Ioannidis, A. D. Keromytis, Offline Micropayments without Trusted Hardware, in: Proceedings of the Fifth International Conference on Financial Cryptography, 2001, pp. 21–40.
- [63] D. Chaum, Achieving Electronic Privacy, *Scientific American* (1992) 96–101.
- [64] D. Chaum, Blind signatures for untraceable payments, in: *Advances in Cryptology: Crypto '82 Proceedings*, Plenum Press, 1982.
- [65] G. Medvinsky, C. Neuman, NetCash: A design for practical electronic currency on the internet, in: Proceedings of the Second ACM Conference on Computer and Communication Security, 1994.
- [66] M. Bellare, J. Garay, C. Jutla, M. Yung, VarietyCash: a Multi-Purpose Electronic Payment System, in: Proceedings of the Third USENIX Workshop on Electronic Commerce, USENIX, 1998.
- [67] T. Poutanen, H. Hinton, M. Stumm, NetCents: A Lightweight Protocol for Secure Micropayments, in: Proceedings of the Third USENIX Workshop on Electronic Commerce, USENIX, 1998.
- [68] M. S. Manasse, The Millicent protocols for electronic commerce, in: Proceedings of the First USENIX Workshop on Electronic Commerce, USENIX, 1995.

- [69] L. Tang, A Set of Protocols for MicroPayments in Distributed Systems, in: Proceedings of the First USENIX Workshop on Electronic Commerce, USENIX, 1995.
- [70] C. Jutla, M. Yung, Paytree: amortized signature for flexible micropayments, in: Proceedings of the Second USENIX Workshop on Electronic Commerce, USENIX, 1996.
- [71] R. Rivest, A. Shamir, PayWord and MicroMint, *CryptoBytes* 2 (1) (1996) 7–11.
- [72] R. Hauser, M. Steiner, M. Waidner, Micro-payments based on ikp, in: Proceedings of the 14th Worldwide Congress on Computer and Communication Security Protection, 1996.
- [73] A. Herzberg, Safeguarding Digital Library Contents, *D-Lib Magazine*.
- [74] C. Neuman, G. Medvinsky, Requirements for network payment: The Netcheque prospective, in: Proceedings of IEEE COMCON, 1995.
- [75] M. Bellare, J. Garay, A. Herzberg, H. Krawczyk, M. Steiner, G. Tsudik, M. Waidner, iKP – A Family of Secure Electronic Payment Protocols, in: Proceedings of the First USENIX Workshop on Electronic Commerce, USENIX, 1995.
- [76] E. Foo, C. Boyd, A Payment Scheme Using Vouchers, in: Proceedings of the Second International Conference on Financial Cryptography, no. 1465 in *Lecture Notes in Computer Science*, Springer-Verlag, 1998, pp. 103–121.
- [77] B. Cox, D. Tygar, M. Sirbu, NetBill security and transaction protocol, in: Proceedings of the First USENIX Workshop on Electronic commerce, USENIX, 1995.