David Wetherall, Thomas Anderson, and Stefan Savage, "**Distributed solution for regulating network traffic**", *US patent 7058015*, Filed August 4th 2000, Issued June 6th, 2006.

Thomas Anderson, David Wetherall and Stefan Savage, "**Progressive and distributed regulation of selected network traffic destined for a network node**", *US patent 6801503*, Filed October 9th, 2000, Issued October 5th 2004.

Michael Burrows, Charles G. Nelson, Stefan Savage and Patrick Sobalvarro, "**Detecting concurrency errors in multi-threaded programs**", *US patent 6009269*, Filed March 10th 1997, Issued December 28th, 1999.

John Wilkes and Stefan Savage, "**Frequently-redundant array of independent disks**", *US patent 5720025*, Filed January 18th 1996, Issued February 17th 1998.

Software & Data

WHOIS parser and parsed .com records. Worlds most accurate parser of WHOIS records and easily updatable via labelling (machine learning-based). Also, distributed parsed records for 95% of .com domain names. Code freely available under GPL, data selectively available on request.

BlueSky Cloud-based Log-Structured File System. Code freely available under GPL.

Click Trajectory database. Multi-terabyte dataset mapping registrars, name server hosting, content hosting, sponsoring businesses and acquiring bank data for virtually all counterfeit pharmaceutical, replica and software spam. Not publicly available, but particular records shared with researchers and law enforcement on request.

Malicious URL classifier. Online machine learning system for classifying Web URLs as likely pointing to unwanted sites (e.g., spam-advertised, phishing, malware, etc.) Selectively available.

Judo spam template inference system. System for automatically inferring regular expression from concentrated spam output from a given botnet. Selectively available.

Honeyfarm spam data. Spam output generated from series of botnets operated in ICSI/UCSD botnet studies. Selectively available.

Cumulus backup system. Cloud-based (Amazon S3) backup system optimized for simple storage interface with capacity pricing. Freely available.

Storm botnet crawler and command and control database. Crawler for querying proxy nodes in Storm peer-to-per botnet and populate database. Also many months of complete data on Storm-based spam campaigns based on this data. Made available selectively (primarily to key industrial partners and law enforcement).

UCSD CSE 802.11 traces. Anonymized traces of complete 802.11activity in the UCSE CSE building (across all four floors and all active frequencies) as captured by Jigsaw. Freely available.

Jigsaw Wireless Analysis Tool. System to synchronize and unify hundreds of simultaneous radio traces into a single global view of wireless activity spanning the physical, datalink, network and transport layers. Freely available under GPL.

Potemkin VMM. Copy-on-write extensions to Xen Virtual Machine Monitor – allows highly scalable implementation of virtual machines that share state. Freely available under GPL.

SyncScan. Driver and system modifications to allow wireless 802.11 clients to continuously monitor the signal strength from different access points without disrupting data communication. Freely available to academic and research sites.

Earlybird. Network traffic analysis system that automatically identifies and extracts textual signatures for novel worms and viruses. Exclusively licensed to Netsift Inc -> Cisco Inc.

AutoFocus. Traffic analysis system that converts Cisco Netflow data from routers into concise lists of high-volume traffic clusters according to algorithms in 2003 ACM SIGCOMM paper. Freely available to academic and research sites.

P2P Host Availability Dataset. Week-long active measurement trace of individual host availability for 2400 random hosts deployed Overnet P2P file sharing system. Particularly useful because Overnet maintains globally unique id's and data isn't subject to aliasing. Distributed widely to academic sites and industrial research labs.

Swat 802.11 Vulnerability Testing Tool. Low-level library and interface to sending arbitrary 802.11 frames using undocumented features of Intersil firmware interface. Distribution limited.

DoS Backscatter Dataset. Packet-level header trace from UCSD Network Telescope identifying global denial-of-service attacks utilizing spoofed packet floods via "backscatter analysis". Distributed widely to academic sites and government research labs – the only dataset of its kind publicly available.

Sting. A network measurement tool that estimates one-way path information, such as drop rate, re-ordering, queuing delay, and bandwidth, by manipulating the behavior of the popular TCP protocol. The alpha release, for FreeBSD and Linux, has been distributed to more than 500 sites, including a product effort at Hewlett-Packard/Agilent. Publicly available.

Eraser. A dynamic race detection tool that automatically detects errors in multi-threaded software using binary rewriting on the Alpha platform. Owned by DEC/Compaq/Hewlett Packard. This tool led to DEC/Compaq/HP's *Visual Threads* software debugging package.

SPIN. An extensible operating system written in Modula-3 for Alpha and Intel platforms. Publicly available.

Real-Time Mach. A real-time extension of the Mach operating system with support for multiple real-time scheduling and synchronization policies and a framework to support CPU reservation and priority inheritance. Publicly available.

Exhibit 1025

Declaration of Dr. Sylvia Hall-Ellis

Customer No.: 000027683

Haynes and Boone, LLP
IP Section
2323 Victory Avenue, Suite 700
Dallas, Texas 75219
Telephone [214] 651.5000
Fax [214] 200.0853

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE CENTRAL REEXAMINATION UNIT

UNIFIED PATENTS, LLC

Third Party Requester

Ex Parte Reexamination of
U.S. Patent 7,523,497

DECLARATION OF SYLVIA HALL-ELLIS, PH.D.

I, Sylvia D. Hall-Ellis, Ph.D., declare as follows:

I. INTRODUCTION

1. My name is Sylvia D. Hall-Ellis. I have been retained as an expert by Unified Patents, LLC who I am informed is a petitioner to this IPR proceeding ("the Petitioner").

2. I have written this declaration at the request of the Petitioner to provide my expert opinion regarding the public availability and authenticity of a number of publications identified in Section V below. My Declaration sets forth my opinions in detail and provides the basis for my opinions regarding the authenticity and public availability of these publications. If called to testify in the above-captioned matter, I will testify with regard to the opinions and bases set forth below.

3. I reserve the right to supplement or amend my opinions, and bases for them, in response to any additional evidence, testimony, discovery, argument, and/or other additional information that may be provided to me after the date of this Declaration.

4. As of the preparation and signing of this declaration, many libraries across the nation are closed or permit only limited access due to the COVID-19 virus. However, were the libraries open, I would expect to be able to obtain paper copies of at least some of the documents in this declaration. Additionally, it is my typical practice to obtain a paper copy of each publication to further confirm my opinions

that the documents were available prior to the alleged priority date of the patent at issue. I reserve the right to supplement my declaration when the libraries reopen to provide such information.

5. I am being compensated for my time spent working on this matter at my normal consulting rate of $300 per hour, plus reimbursement for any additional reasonable expenses. My compensation is not in any way tied to the content of this report, the substance of my opinions, or the outcome of this proceeding. I have no other interests in this proceeding or with any of the parties.

6. All of the materials that I considered and relied upon are discussed explicitly in this declaration.

II. QUALIFICATIONS

7. I am currently an Adjunct Professor in the School of Information at San José State University in San José, California. I obtained a Master of Library Science from the University of North Texas in 1972 and a Ph.D. in Library Science from the University of Pittsburgh in 1985. Over the last fifty years, I have held various positions in the field of library and information resources. I was first employed as a librarian in 1966 and have been involved in the field of library sciences since, holding numerous positions.

8. I am a member of the American Library Association (ALA) and its Association for Library Collections & Technical Services (ALCTS) Division, and I

Cloudflare, Inc. Exhibit 1003 Page 1320

served on the Committee on Cataloging: Resource and Description (which wrote the new cataloging rules) and as the chair of the Committee for Education and Training of Catalogers and the Competencies and Education for a Career in Cataloging Interest Group. I also served as the Chair of the ALCTS Division's Task Force on Competencies and Education for a Career in Cataloging. Additionally, I have served as the Chair for the ALA Office of Diversity's Committee on Diversity, as a member of the REFORMA National Board of Directors, and as a member of the Editorial Board for the ALCTS premier cataloging journal, *Library Resources and Technical Services*. Currently I serve as a Co-Chair for the Library Research Round Table of the American Library Association.

9. I have also given over one hundred presentations in the field, including several on library cataloging systems and Machine-Readable Cataloging ("MARC") standards. My current research interests include library cataloging systems, metadata, and organization of electronic resources.

10. My full curriculum vitae is attached hereto as **Attachment A**.[1]

III. PRELIMINARIES

A. Scope of Declaration and Legal Standards

11. I am not an attorney and will not offer opinions on the law. I am,

[1] I understand that my CV is also being submitted as Exhibit 1026 in this IPR proceeding.

Cloudflare, Inc. Exhibit 1003 Page 1321

however, rendering my expert opinion on the authenticity of the documents referenced herein and on when and how each of these documents was disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located the documents before the dates discussed below with respect to the specific documents.

12. I am informed by counsel that a printed publication qualifies as publicly accessible as of the date it was disseminated or otherwise made available such that a person interested in and ordinarily skilled in the relevant subject matter could locate it through the exercise of ordinary diligence.

13. While I understand that the determination of public accessibility under the foregoing standard rests on a case-by-case analysis of the facts particular to an individual publication, I also understand that a printed publication is rendered "publicly accessible" if it is cataloged and indexed by a library such that a person interested in the relevant subject matter could locate it (*i.e.*, I understand that cataloging and indexing by a library is sufficient, though there are other ways that a printed publication may qualify as publicly accessible). One manner of sufficient indexing is indexing according to subject matter category. I understand that the cataloging and indexing by a single library of a single instance of a particular printed publication is sufficient, even if the single library is in a foreign country. I

understand that, even if access to a library is restricted, a printed publication that has been cataloged and indexed therein is publicly accessible so long as a presumption is raised that the portion of the public concerned with the relevant subject matter would know of the printed publication. I also understand that the cataloging and indexing of information that would guide a person interested in the relevant subject matter to the printed publication, such as the cataloging and indexing of an abstract for the printed publication, is sufficient to render the printed publication publicly accessible.

14. I understand that routine business practices, such as general library cataloging and indexing practices, can be used to establish an approximate date on which a printed publication became publicly accessible. I also understand that the indicia on the face of a reference, such as printed dates and stamps, are considered as part of the totality of the evidence.

B. Persons of Ordinary Skill in the Art

15. I am informed by counsel that the subject matter of this proceeding generally relates to network data transmission.

16. I have been informed by counsel that a "person of ordinary skill in the art at the time of the invention" is a hypothetical person who is presumed to be familiar with the relevant field and its literature at the time of the invention. This hypothetical person is also a person of ordinary creativity, capable of understanding

the scientific principles applicable to the pertinent field.

17. I am informed by counsel that a person of ordinary skill in this subject matter or art would typically be someone who has a bachelor's degree in electrical engineering, computer engineering, computer science, physics, or a related subject, and one to two years of work experience in network data transmission. I have further been informed by counsel that a person of ordinary skill in the art would have been familiar with and able to understand the information known in the art relating to these fields, including the publications discussed in this declaration.

C. Use of Authoritative Databases

18. In preparing this report, I used authoritative databases, such as the OCLC bibliographic database and the Library of Congress Online Catalog, to confirm citation details of the publication discussed. Unless I note otherwise below in reference to a specific serial publication, it is my expert opinion that this standard protocol was followed for the publications discussed in Section V below.

19. *Indexing*. A researcher may discover material relevant to his or her topic in a variety of ways. One common means of discovery is to search for relevant information in an index of periodical and other publications. Having found relevant material, the researcher will then normally obtain it online, look for it in libraries, or purchase it from the publisher, a bookstore, a document delivery service, or other provider. Sometimes, the date of a document's public accessibility will involve both

Cloudflare, Inc. Exhibit 1003 Page 1324

indexing and library date information.

20. Indexing services use a wide variety of controlled vocabularies to provide subject access and other means of discovering the content of documents. The formats in which these access terms are presented vary from service to service.

21. Online indexing services and digital repositories commonly provide bibliographic information, abstracts, and full-text copies of the indexed publications, along with a list of the documents cited in the indexed publication. These services also often provide lists of publications that cite a given document. A citation of a document is evidence that the document was publicly available and in use by researchers no later than the publication date of the citing document.

22. *Internet Archive*. The *Internet Archive* is a non-profit digital library founded in 1996. The *Internet Archive* maintains an archive of webpages collected from the Internet using software called a crawler. Crawlers automatically create a snapshot of webpages as they existed at a certain point in time. The WayBack Machine is an application using a crawler created by the Internet Archive to search its archive of Web page URLs and to represent, graphically, the date of each crawler capture. The Internet Archive captures data that is publicly available. Some sites are "not archived because they were password protected, blocked by robots.txt, or otherwise inaccessible to our automated systems. Site owners might have also requested that their sites be excluded from the WayBack Machine." Many Internet

Archive captures made by the WayBack Machine have a banner at the top with the capture date prominently displayed. Other dates when captures of the same URL have been made are indicated to the right and left of the date provided in the banner. Some captures may lack this banner. In any case, the URL for the capture begins with the identification of the Internet Archive page (e.g., http://web.archive.org/web/) followed by information that dates and time stamps the capture as follows: year in yyyy, month in mm, day in dd, time code in hh:mm:ss (e.g., 20071120082013, or November 20, 2007 at 8:20:13 a.m.). These elements are then followed by the URL of the original capture site. When links are active, the WayBack Machine is programmed to produce the archived file with the closest available date (not the closest available prior date) to the page upon which the link appeared and was clicked. I and other librarian professionals are familiar with the Internet Archive and the Wayback Machine.[2]

D. Summary of Opinions

23. I am informed by counsel that the alleged priority date for the patent at issue is November 16, 2000. As I will explain below, it is my opinion that the publications discussed in my Declaration were publicly accessible before the alleged November 16, 2000 priority date.

---

[2] For more information about the Internet Archive see the WayBack Machine FAQ, https://archive.org/about/faqs.php#The_Wayback_Machine.

IV. LIBRARY CATALOGING PRACTICES

A. MARC Records

24. I am fully familiar with the library cataloging standard known as the MARC standard, which is an industry-wide standard method of storing and organizing library catalog information. MARC was first developed in the 1960s by the Library of Congress. A MARC-compatible library is one that has a catalog consisting of individual MARC records for works made available at that library.

25. Since at least the early 1970s and continuing to the present day, MARC has been the primary communications protocol for the transfer and storage of bibliographic metadata in libraries.[3] As explained by the Library of Congress:

> You could devise your own method of organizing the bibliographic information, but you would be isolating your library, limiting its options, and creating much more work for yourself. Using the MARC standard prevents duplication of work and allows libraries to better share bibliographic resources. Choosing to use MARC enables libraries to acquire cataloging data that is predictable and reliable. If a library were to develop a "home-grown"

[3] A complete history of the development of MARC can be found in *MARC: Its History and Implications* by Henrietta D. Avram (Washington, DC: Library of Congress, 1975) and available online from the Hathi Trust (https://babel.hathitrust.org/cgi/pt?id=mdp.39015034388556;view=1up;seq=1; last visited March 23, 2021).

system that did not use MARC records, it would not be taking advantage of an industry-wide standard whose primary purpose is to foster communication of information.

Using the MARC standard also enables libraries to make use of commercially available library automation systems to manage library operations. Many systems are available for libraries of all sizes and are designed to work with the MARC format. Systems are maintained and improved by the vendor so that libraries can benefit from the latest advances in computer technology. The MARC standard also allows libraries to replace one system with another with the assurance that their data will still be compatible.

Why Is a MARC Record Necessary? LIBRARY OF CONGRESS, http://www.loc.gov/marc/umb/um01to06.html#part2 (last visited March 23, 2021).

26. Thus, almost every major library in the world is MARC-compatible. *See, e.g.*, *MARC Frequently Asked Questions (FAQ)*, LIBRARY OF CONGRESS, https://www.loc.gov/marc/faq.html (last visited March 23, 2021) ("MARC is the acronym for MAchine-Readable Cataloging. It defines a data format that emerged from a Library of Congress-led initiative that began nearly fifty years ago. It provides the mechanism by which computers exchange, use, and interpret bibliographic information, and its data elements make up the foundation of most library catalogs used today."). MARC is the ANSI/NISO Z39.2-1994 standard

Cloudflare, Inc. Exhibit 1003 Page 1328

(reaffirmed in 2016) for Information Interchange Format. The full text of the standard is available from the Library of Congress at http://www.loc.gov/marc/bibliographic/ (last visited March 23, 2021).

27. A MARC record comprises several fields, each of which contains specific data about the work. Each field is identified by a standardized, unique, three-digit code corresponding to the type of data that follow. *See, e.g.*, http://www.loc.gov/marc/umb/um07to10.html (last visited March 23, 2021); http://www.loc.gov/marc/bibliographic/ (last visited March 32, 2021). For example, a work's title is recorded in field 245, the primary author of the work is recorded in field 100, a work's International Standard Book Number ("ISBN") is recorded in field 020, a work's International Standard Serial Number ("ISSN") is recorded in field 022, and the publication date is recorded in field 260 under the subfield "c." *Id.*[4] If a work is a periodical, then its publication frequency is recorded in field 310, and the publication dates (e.g., the first and last publication) are recorded in field 362, which is also referred to as the enumeration/chronology field. *See*

[4] In some MARC records, field 264 is used rather than field 260 to record publication information. *See* http://www.loc.gov/marc/bibliographic/bd264.html (last visited March 23, 2021) ("Information in field 264 is similar to information in field 260 (Publication, Distribution, etc. (Imprint)). Field 264 is useful for cases where the content standard or institutional policies make a distinction between functions").

http://www.loc.gov/marc/bibliographic/bd3xx.html (last visited March 23, 2021).[5]

28. The library that initially created the MARC record is reflected in field 040 in subfield "a" with that library's unique library code. *See, e.g.*, http://www.loc.gov/marc/umb/um07to10.html (last visited March 23, 2021); http://www.loc.gov/marc/bibliographic/ (last visited March 23, 2021). Once a MARC record for a particular work is originally created by one library, other libraries can use that original MARC record to then create their own MARC records for their own copies of the same work. These other libraries may modify or add to the original MARC record as necessary to reflect data specific to their own copies of the work. However, the library that created the original MARC record would still be reflected in these modified MARC records (corresponding to other copies of the same work at other libraries) in field 040, subfield "a". The modifying library (or

[5] Upwards of two-thirds to three-quarters of book sales to libraries come from a jobber or wholesaler for online and print resources. These resellers make it their business to provide books to their customers as fast as possible, often providing turnaround times of only a single day after publication. Libraries purchase a significant portion of the balance of their books directly from publishers themselves, which provide delivery on a similarly expedited schedule. In general, libraries make these purchases throughout the year as the books are published and shelve the books as soon thereafter as possible in order to make the books available to their patrons. Thus, books are generally available at libraries across the country within just a few days of publication.

Cloudflare, Inc. Exhibit 1003 Page 1330

libraries) is reflected in field 040, subfield "d". *See* http://www.loc.gov/marc/bibliographic/bd040.html (last visited March 23, 2021).

29. I consulted the Directory of OCLC Libraries (http://www.oclc.org/contacts/libraries.en.html; last visited March 23, 2021) in order to identify the institution that created or modified the MARC record. Moreover, when viewing the MARC record online via Online Computer Library Center's ("OCLC") bibliographic database, which I discuss further below, hovering over a library code in field 040 with the mouse reveals the full name of the library. I also used this method of "mousing over" the library codes in the OCLC database to identify the originating and modifying libraries for the MARC records discussed in this Declaration.

30. MARC records also include one or more fields that show information regarding subject matter classification. For example, 6XX fields are termed "Subject Access Fields." *See* http://www.loc.gov/marc/bibliographic/bd6xx.html (last visited March 23, 2021). Among these, for example, is the 650 field; this is the "Subject Added Entry – Topical Term" field. *See* http://www.loc.gov/marc/bibliographic/bd650.html (last visited March 23, 2021). The 650 field is a "[s]ubject added entry in which the entry element is a topical term." *Id.* These entries "are assigned to a bibliographic record to provide access according to generally accepted thesaurus-building rules (e.g., *Library of Congress*

Subject Headings (LCSH), *Medical Subject Headings* (MeSH))." *Id.*

31. Further, MARC records can include call numbers, which themselves contain a classification number. For example, a MARC record may identify a 050 field, which is the "Library of Congress Call Number." *See* http://www.loc.gov/marc/bibliographic/bd050.html (last visited March 23, 2021). A defined portion of the Library of Congress Call Number is the classification number, and "source of the classification number is *Library of Congress Classification* and the *LC Classification-Additions and Changes*." *Id.* Thus, the 050 field may be used to show information regarding subject matter classification. As another example, a MARC record may identify a 082 field, which is the "Dewey Decimal Classification Number." *See* http://www.loc.gov/marc/bibliographic/bd082.html (last visited March 23, 2021). A defined portion of the Dewey Decimal Classification Number is the classification number, and "source of the classification number is the Dewey Decimal Classification and Relative Index," which is published online. *Id.* Thus, the 082 field may be used to show information regarding subject matter classification consistent with the Dewey Decimal Classification System.

32. Each item in a library has a single classification number. A library selects a classification scheme (*e.g.*, the Library of Congress Classification scheme just described or a similar scheme such as the Dewey Decimal Classification scheme) and uses it consistently. When the Library of Congress assigns the

classification number, it appears as part of the 050 field, as discussed above. For MARC records created by libraries other than the Library of Congress (e.g., a university library or a local public library), the classification number may appear in a 09X (e.g., 090) field. *See* http://www.loc.gov/marc/bibliographic/bd09x.html (last visited March 23, 2021).

33. When a MARC-compatible library acquires a work, it creates a MARC record for its copy of the work in its computer catalog system in the ordinary course of its business. This MARC record (for the copy of a work available at the particular library) may be later accessed by researchers in a number of ways. For example, many libraries, including the Library of Congress, make their MARC records available through their website. As an example, the MARC record for the copy of *The Unlikely Spy*, by Daniel Silva,[6] available at the Library of Congress can be viewed through the Library of Congress website, at https://catalog.loc.gov/vwebv/staffView?searchId=20265&recPointer=1&recCount=25&bibId=2579985 (last visited March 23, 2021). One could, of course, always physically visit the library at which the work is available, and request to see that library's MARC record for the work. Moreover, members of the Online Computer Library Center ("OCLC") can access the MARC records of other member

[6] *The Unlikely Spy* is a 1996 novel written by Daniel Silva, who happens to be one of my favorite authors.

Cloudflare, Inc. Exhibit 1003 Page 1333

institutions through OCLC's online bibliographic database, as I explain further below.

B. OCLC

34. The OCLC was created "to establish, maintain and operate a computerized library network and to promote the evolution of library use, of libraries themselves, and of librarianship, and to provide processes and products for the benefit of library users and libraries, including such objectives as increasing availability of library resources to individual library patrons and reducing the rate of rise of library per-unit costs, all for the fundamental public purpose of furthering ease of access to and use of the ever-expanding body of worldwide scientific, literary and educational knowledge and information."[7] Among other services, OCLC and its members are responsible for maintaining the WorldCat database (http://www.worldcat.org/; last visited March 23, 2021), used by independent and institutional libraries throughout the world. All libraries that are members of OCLC are MARC-compatible. *See, e.g.*, https://help.oclc.org/Metadata_Services/OCLC-MARC_records/About_OCLC-MARC_records (last visited March 23, 2021)

[7] Third Article, Amended Articles of Incorporation of OCLC Online Computer Library Center, Incorporated (available at https://www.oclc.org/content/dam/oclc/membership/articles-of-incorporation.pdf; last visited March 23, 2021).

Cloudflare, Inc. Exhibit 1003 Page 1334

("OCLC-MARC records describes records produced since November 1993."); https://www.oclc.org/support/services/worldcat/documentation/cataloging/electronicresources.en.html (last visited March 23, 2021) ("Like the two superseded OCLC documents, this revised set of guidelines is intended to assist catalogers in creating records for electronic resources in WorldCat, the OCLC Online Union Catalog. These guidelines pertain to OCLC-MARC tagging (that is, content designation). Cataloging rules and manuals (such as AACR2) govern the content of records. You should implement these guidelines immediately.").

35. When an OCLC member institution acquires a publication, like the other MARC-compatible libraries discussed above, it creates a MARC record for this publication in its computer catalog system in the ordinary course of its business. MARC records created at the Library of Congress are tape-loaded into the OCLC database through a subscription to MARC Distribution Services daily or weekly. Once the MARC record is created by a cataloger at an OCLC member library or is tape-loaded from the Library of Congress, the MARC record is then made available to any other OCLC members online, and thereby made available to the public. Accordingly, once the MARC record is created by a cataloger at an OCLC member library or is tape-loaded from the Library of Congress, any publication corresponding to the MARC record has been cataloged and indexed according to its subject matter such that a person interested in that subject matter could, with

reasonable diligence, locate and access the publication through any library with access to the OCLC bibliographic database or through the Library of Congress.

36. When a MARC-compatible library creates an original MARC record for a work, the library records the date of creation of that MARC record in **field 008**, characters 00 through 05, in the ordinary course of its business. *See* http://www.loc.gov/marc/bibliographic/bd008a.html (last visited March 23, 2021). For OCLC member institutions that use OCLC software to create original MARC records, the date of creation in field 008 is automatically supplied by the OCLC software. The MARC record creation date in field 008 thus reflects the date on which, or shortly after which, a work was first acquired and cataloged by the library that created the original MARC record.

37. When other MARC-compatible libraries subsequently acquire their own copies of the same work, as mentioned, they create MARC records in their own computer catalog systems for their copies in the ordinary course of business.[8] They may use a MARC record previously created for that work (by another MARC-compatible library) to create their own MARC records for their own copies of that

[8] Initial contributions to the bibliographic database for a work are called "master records."

Cloudflare, Inc. Exhibit 1003 Page 1336

same work.[9] The previously created MARC record used by subsequently-acquiring libraries to create MARC records for their own copies may be obtained through the OCLC bibliographic database, as described above. If, when creating a MARC record to represent its own copy of the work, the subsequently-acquiring library uses the master MARC record in its original form, the subsequently-acquiring library cannot reenter data into the 008 field; therefore, the date in the 008 field will continue to reflect the date the MARC record was initially created by the originating library. On the other hand, if the subsequently-acquiring library modifies the previously created MARC record when creating its own MARC record for its own copy of the work, the subsequently-acquiring library may enter into the 008 field of its own MARC record the date its own MARC record was created.[10] But the library that created the original MARC record used by the subsequently-acquiring library would still be reflected in the MARC record of the subsequently-acquiring library in field

---

[9] When a local library uses a master record in OCLC and produces (or downloads) it to the in-house system, the three-character symbol for the subsequent library is added to the holdings for the work.

[10] This practice is not required by, but is nevertheless consistent with, the MARC standard. Many MARC records exist today whose 008 fields indicate when the first original MARC record for a work was created, rather than when a derivative record was created based on the original MARC record by a subsequently-acquiring library for its own computer catalog system.

Cloudflare, Inc. Exhibit 1003 Page 1337

040, subfield "a". Thus, the work identified by any MARC record possessed by any MARC-compatible library would have been accessible to the public at least as of the date shown in the 008 field, or shortly thereafter, either from the library that possesses the MARC record itself, or from the originating library indicated in field 040, subfield "a". As discussed, a MARC-compatible library in the ordinary course of its business creates a MARC record in its own catalog system for a work when it acquires a copy of that work.

38. Moreover, when a MARC record is created by a library for its own copy of a work, **field 005** is automatically populated with the date that MARC record was created in year, month, day format (YYYYMMDD). *See* http://www.loc.gov/marc/bibliographic/bd005.html (last visited March 23, 2021).[11] Thereafter, the library's computer system may automatically update the date in field 005 every time the library updates the MARC record (*e.g.*, to reflect that an item has been moved to a different shelving location within the library). *Id.*[12] Thus, the work identified by any MARC record possessed by any MARC-compatible library would

[11] Some of the newer library catalog systems also include hour, minute, second (HHMMSS).

[12] Field 005 is visible when viewing a MARC record via an appropriate computerized interface. But when a MARC record is printed directly to hardcopy from the OCLC database, the "005" label is not shown. The date in the 005 field instead appears next to the label "Replaced."

have been accessible to the public at least as of the date shown in the 005 field, or shortly thereafter, from the library that possesses the MARC record itself. As noted, because the 005 field may be updated each time the library updates its MARC record, the work identified by the MARC record may, in fact, have been accessible to the public from that library much earlier than the date indicated in the 005 field.

39. Based on my personal experience as a professional librarian using the MARC and OCLC resources, it has been my experience that both of these resources were continuously operational and available since at least 1992. Indeed, in the course of my work, I have extensively used both of these resources over the past 30+ years, and I have consistently found the information contained within these resources to be complete and reliable. I have never found the date of accessibility as indicated in fields 008 or 005 to be incorrect. And in only a minute number of cases have I found any errors at all in these records – none of which affected my ability to render an accurate opinion as to accessibility, indexing, or subject headings.

V. PUBLICATIONS IN THIS PROCEEDING

A. "Practical Network Support for IP Traceback" by Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson ("Savage paper") [Exhibit 1005]

40. Exhibit 1005 is a copy of a conference paper published in the volume titled *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication* on pages 295-306. The

conference paper is titled "Practical Network Support for IP Traceback" by Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson (hereafter "Savage paper"). The ACM SIGCOMM 2000 Conference titled Applications, Technologies, Architectures, and Protocols for Computer Communication was held on August 28 through September 1, 2000, in Stockholm, Sweden.[13] The conference program indicated that the Savage paper was presented during Session 8, "Tracing and Measurement," on Friday, September 1, 2000.[14] I obtained this copy of the Savage paper (Exhibit 1005) from the University of Cincinnati Library (Cincinnati, Ohio). Exhibit 1040 attached to my declaration is a date-stamped copy of an excerpt of *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication*, containing the Savage paper (Exhibit 1005) at pages 295-306. The text of the Savage paper (Exhibit 1005) is complete; no pages are missing, and the text on each page appears to flow seamlessly from one page to the next; further, there are no visible alterations to the document. Exhibit 1040 (including the Savage paper (Exhibit 1005)) was found within the custody of a library – a place where, if authentic, a copy of this conference proceedings volume would likely be. Exhibit 1040 is a true and correct copy in a condition that creates no suspicion about authenticity.

---

[13] http://conferences.sigcomm.org/sigcomm/2000/

[14] http://conferences.sigcomm.org/sigcomm/2000/conf/techprog.htm

41. Attached hereto as Attachment 1a is a true and correct copy of the MARC record for this conference proceedings volume (Exhibit 1040) from the University of Cincinnati Library online catalog. The library ownership is indicated by the presence of the library's code (CIN) in the fourth subfield "d" of the 040 field. The library continues to update this MARC record and enhanced the MARC record to meet current cataloging rules. I personally identified and retrieved the library catalog record which is Attachment 1a.

42. Based on finding a print copy of Exhibit 1040 (including the Savage paper (Exhibit 1005)) in the University of Cincinnati Library and MARC record in its online library catalog attached as Attachment 1a, it is my opinion that the book *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication* was publicly available in the University of Cincinnati Library as of October 14, 2000, as shown in field 008 ("001014").

43. Attached hereto as Attachment 1b is a true and correct copy of the MARC record for the volume *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication* obtained from the OCLC bibliographic database. I personally identified and retrieved the MARC record that is Attachment 1b. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Attachment 1b, that library code is "OGE," which means that the MARC record for this conference

proceedings volume was cataloged at the University of California – San Diego Library. As can be seen in the "Entered" field in the MARC record for this exhibit, a cataloger at the University of California – San Diego created OCLC record number 45172567 on October 14, 2000, as shown in the "Entered" field ("20001014"). The library continues to update this MARC record the MARC record to meet current cataloging rules. The most recent enhancement to the MARC record occurred on March 11, 2019, as shown in the "Replaced" field ("20190311").

44. Attachment 1b further includes an entry in field 050 ("TK5105.5 $b .A26 2000")—as described above, a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system) and an entry in field 082 ("004.6")—as described above, a subject matter classification number consistent with the Dewey Decimal classification system. Attachment 1b also shows that Exhibit 1040 was catalogued with three descriptor terms reading "Computer networks $v Congresses" (see Attachment 1c, Library of Congress subject heading sh2008101222), "Computer network protocols $v Congresses" (see Attachment 1d, Library of Congress subject heading sh2008101491), and "Data transmission systems $v Congresses" (see Attachment 1e, Library of Congress subject heading sh2008102044) in the 650 fields. Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 1b was indexed according to its subject matter by

virtue of at least three independently sufficient classifications: the field 050 entry, the field 082 entry, and the field 650 entries. Further, as of October 14, 2000, the MARC record attached hereto as Attachment 1b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the proceedings volume, which means that the corresponding publication was publicly available on or before that same date through any library with access to the OCLC bibliographic database or through an individual library.

45. Attachment 1b indicates that the volume *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication* as cataloged at the University of Cincinnati Library is currently available from 25 libraries ("25 other holdings"). In view of the above, the volume *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication* was publicly available on or shortly after October 14, 2000, because by that date it had been received, cataloged, and indexed at the University of Cincinnati Library and made part of the OCLC bibliographic database.[15] For these reasons, it is my opinion that *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer*

[15] As shown on the first page of Exhibit 1040, the University of Cincinnati Library date-stamped the Conference Proceedings Book on October 17, 2000, shortly after it had been received, cataloged, and indexed by the University of Cincinnati Library.

Cloudflare, Inc. Exhibit 1003 Page 1343

Communication (Exhibit 1040), including the Savage paper (Exhibit 1005) contained therein, was published and publicly accessible at least by October 14, 2000.

46. Not only was the Savage paper (Exhibit 1005) accessible and available to others in the field as of October 2000, but individuals also obtained and cited this article in other works, which is further confirmation of its availability and accessibility.

47. **Park technical report.** For example, in their technical report titled *A Proactive Approach to Distributed DoS Attack Prevention using Route-Based Packet Filtering*, Kihong Park and Heejo Lee (hereafter "Park") cited the Savage paper (see Exhibit 1040, page 23, entry 23).

48. Exhibit 1040 is the technical report prepared by Park and Lee of the Network Systems Lab in the Department of Computer Science at Purdue University (West Lafayette, Indiana). Exhibit 1040 filed in this proceeding is a true and correct copy, which I obtained from the digital repository Purdue e-Pubs[16] written by faculty and graduate students at Purdue University.[17] Specifically, the text of Exhibit 1040 is complete; no pages are missing, and the text on each page appears to flow

[16] https://docs.lib.purdue.edu/

[17] https://www.cerias.purdue.edu/assets/pdf/bibtex_archive/bibtex_archive/2001-40.pdf

seamlessly from one page to the next; further, there are no visible alterations to the document. Further, Exhibit 1040 was found within the custody of the issuing higher education institution (i.e., the Purdue e-Pubs at Purdue University), a place where, if authentic, a copy of this technical report would likely be. Exhibit 1040 is a true and correct copy in a condition that creates no suspicion about its authenticity.

49. A technical report exhibits the following characteristics: the technical report is attributed to one or more creators or authors (in this case Kihong Park and Heejo Lee) who affix a date of creation (in this case December 3, 2000), recognize one or more outside funding agencies (in this case five National Science Foundation (NSF) grants, the Purdue Research Foundation, the Santa Fe Institute, Sprint, and Xerox, and the Center for Education and Research in Information Assurance and Security (CERIAS), a list of scholarly references (in this case 31 references), and publishes the document through a sponsoring institution or organization (in this case the Department of Computer Science at Purdue University).

50. Based on finding a digital copy of Exhibit 1040 on the Internet, it is my opinion that the technical report titled *A Proactive Approach to Distributed DoS Attack Prevention using Route-Based Packet Filtering* was available to the public by or shortly after December 3, 2000, as indicated on the first page of the document. It is also my opinion that there is no reason to doubt the authenticity of Exhibit 1040 given its appearance, contents, and the web address from which it was obtained.

Cloudflare, Inc. Exhibit 1003 Page 1345

51. **Gil doctoral dissertation.** The Savage paper was cited in a doctoral dissertation by Thomer M. Gil. In his dissertation titled *MULTOPS: A Data Structure for Denial-Of-Service Attack Detection* (hereafter "Gil") cited the Savage paper (see Exhibit 1041, page 48, entry 13).

52. Exhibit 1042 is a dissertation written by Thomer M. Gil as part of the degree requirements of the Division of Mathematics and Computer Science at Vrije Universitei (Amsterdam, Netherlands). Exhibit 1042 filed in this proceeding is a true and correct copy, which I obtained from the digital repository of thesis and dissertations at the Massachusetts Institute of Technology (MIT) (Cambridge, Massachusetts).[18] Specifically, the text of Exhibit 1042 is complete; no pages are missing, and the text on each page appears to flow seamlessly from one page to the next; further, there are no visible alterations to the document. Further, Exhibit 1042 was found within the custody of an issuing higher education institution (i.e., MIT at which two members of the dissertation committee were faculty members), a place where, if authentic, a copy of this dissertation would likely be. Exhibit 1042 is a true and correct copy in a condition that creates no suspicion about its authenticity.

53. A dissertation exhibits the following characteristics: the dissertation is attributed to a single creator or author (in this case Thomer M. Gil) who affixes a

[18] https://thomer.com/mit/thesis.pdf

Cloudflare, Inc. Exhibit 1003 Page 1346

date of creation (in this case December 2000), recognizes a committee of three or more faculty members with their affiliations, (in this case M. Frans Kaashoek, Associate Professor, MIT; Massimiliano Poletto, Research Staff, MIT; and, Andrew S. Tanenbaum, Professor, Vrije Universitei), a list of scholarly references (in this case 19 references), and publishes the document through the higher education institute conferring the doctoral degree (in this case the Division of Mathematics and Computer Science at Vrije Universitei).

54. Based on finding a digital copy of Exhibit 1042 on the Internet, it is my opinion that the dissertation titled *MULTOPS: A Data Structure for Denial-Of-Service Attack Detection* was available to the public by or shortly after December 2000, as indicated on the first page of the document. It is also my opinion that there is no reason to doubt the authenticity of Exhibit 1042 given its appearance, contents, and the web address from which it was obtained.

55. **Eronen research report.** The Savage paper was also cited in a research report titled *Denial of Service in Public Key Protocols* by Pasi Eronen (hereafter "Eronen") (see Exhibit 1043, page 13, entry 43).

56. Exhibit 1043 is the research report prepared by Eronen of the Telecommunications Software and Multimedia Laboratory at the Helsinki University of Technology (Helsinki, Finland). Exhibit 1043 filed in this proceeding

Cloudflare, Inc. Exhibit 1003 Page 1347

is a true and correct copy, which I obtained from the author's website.[19] Specifically, the text of Exhibit 1043 is complete; no pages are missing, and the text on each page appears to flow seamlessly from one page to the next; further, there are no visible alterations to the document. Further, Exhibit 1043 was found within the custody of the author, a place where, if authentic, a copy of this research report would likely be. Exhibit 1043 is a true and correct copy in a condition that creates no suspicion about its authenticity.

57. A research report exhibits the following characteristics: the research report is attributed to one or more creators or authors (in this case Pasi Eronen) who affix a date of creation (in this case 2000), includes a list of scholarly references (in this case 46 references), and publishes the document through a sponsoring institution or organization (in this case the Telecommunications Software and Multimedia Laboratory at the Helsinki University of Technology).

58. Based on finding a digital copy of Exhibit 1043 on the Internet, it is my opinion that the research report titled *Denial of Service in Public Key Protocols* was available to the public by or shortly after December 31, 2000, as indicated in the header of the document. It is also my opinion that there is no reason to doubt the authenticity of Exhibit 1043 given its appearance, contents, and the web address

[19] http://www.pasieronen.com/publications/netsec_2000.pdf

from which it was obtained.

VI. CONCLUSION

59. In signing this Declaration, I recognize that the Declaration will be filed as evidence in a case before Central Reexamination Unit of the United States Patent and Trademark Office. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

60. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: April 28, 2021

Respectfully submitted,



Sylvia D. Hall-Ellis, Ph.D.

Cloudflare, Inc. Exhibit 1003 Page 1349

ATTACHMENT A

Dr. Hall-Ellis *Curriculum Vitae*

CURRICULUM VITAE
SYLVIA D. HALL-ELLIS

EDUCATION

Ph.D., University of Pittsburgh, Pittsburgh, Pennsylvania, 1985
M.P.S., University of Denver, Denver, Colorado, 2014
Post Graduate Studies, University of Texas – San Antonio, Texas, 1975-1976
M.L.S., University of North Texas, Denton, Texas, 1972
B.A., Rockford University, Rockford, Illinois, 1971

PROFESSIONAL EXPERIENCE

1981- **Consultant for higher education, non-profit organizations, and corporations.**

2002- **Adjunct Professor, School of Information, San José State University, San José, California.** Serve as part-time faculty member teaching graduate students in technical services (cataloging, bibliographic control, classification), "core courses," and special topics.

2014-2016 **Director, Grants and Resource Development, Colorado Community College System.** Provided leadership and vision to foster the continued growth of rigorous scholarship, innovative projects, and creative work for statewide system, 13 campuses, and 50 teaching sites serving 155,000 students. Responsible for leadership and ensured efficient functioning of contract and grants in compliance with state & federal requirements and successful implementation and management. Served as a subject matter expert and liaison for college Grant Directors for all issues relating to grants and subcontracts.

2010-2014 **Senior Grant Administrator, Morgridge College of Education, University of Denver (Colorado).** Provided leadership and vision to foster continued growth of rigorous scholarship, innovative research, and creative work in the Morgridge College of Education. Ensure that contract and grants processes function effectively and efficiently for 60 faculty and researchers with a focus on the successful progression and efficient management of grants totaling $13M. Worked effectively and collegially with Department Chairs and Program Coordinators on operational grant-related management activities and with a broad range of internal and external constituencies. Supported the dissemination and promotion of faculty research and scholarship to outside constituents at conferences and through publications. Assisted Principal Investigators and grant project teams by coaching, mentoring, and financial management.

2011-2013 **Interim Director & Assistant Dean, Westminster Law Library, Sturm College of Law, University of Denver.** Planned, organized, and directed all administrative activities for the library serving students, faculty, and alumni; oversaw the employment, retention, promotion, transfer and termination of library personnel; represented the library at professional conferences and public meetings; created and promoted a climate and culture of acceptance for new programs and services, a positive high-quality image of the law library, and that reflect the organization's values, encourage excellent performance, and reward high productivity and innovation; provided leadership and set strategic direction of the organization; ensured that the library provided excellent customer service through solution-oriented staff response to patron needs and by responsiveness and continuous improvement of the organization; promoted, developed, and maintained positive working relationships with colleagues and customers including key stakeholders and groups, higher education institutions, the legal community, other regional libraries and districts statewide, and national library organizations.

2007-2014 **Associate Professor, Library & Information Science, Morgridge College of Education, University of Denver (Colorado).** Served in leadership role and worked collaboratively in program, college, campus and community environments. Advised and supervised students, taught core and specialized courses at the graduate level in an integrative, student-centered learning environment. Served on LIS, College, and University committees, and maintained

working relationships with colleagues in other academic units and information professionals in the Rocky Mountain region and beyond. Served on and chair doctoral student dissertation committees. Oversaw and facilitated the College and LIS graduate student association.

2002-2007 **Assistant Professor, Library & Information Science, College of Education, University of Denver (Colorado).** Served as tenure-track faculty member teaching graduate students in "core courses," resource description and access, service learning, and independent studies. Advised graduate students, participate on LIS and College committees, and serve on doctoral student dissertation committees. Oversaw and facilitated the LIS graduate student association and alumni association.

2000-2002 **Affiliate Faculty, Library & Information Science, College of Education, University of Denver (Colorado).** Served as part-time faculty member teaching graduate students in technical services (cataloging, bibliographic control, classification), "core courses," and special topics. Oversaw and facilitated the LIS graduate student association and alumni association.

2000-2001 **Special Assistant to the Secretary's Regional Representative, U.S. Department of Education, Region VIII, Denver, Colorado.** Served as the principal advisor and representative of the U. S. Secretary of Education's Regional Representative (SRR). Ensured the implementation of major goals of the SRR and the Secretary. Provided leadership on behalf of the SRR in contacts with high-level officials in Region VIII requiring sensitive policy interpretation in communication with senior Department officials to solve problems and resolve issues raised by State and local education officials. Served as the primary contact for School-to-Work/Career, Children's Health Insurance Program, and Safe and Drug-Free Schools. Delivered technical assistance to local education agencies and institutions of higher education in technology, professional development, and school construction.

1999-2000 **Catalog Librarian, Jefferson County Public Library, Lakewood, Colorado.** Performed original, copy cataloging and classification of library materials (English and Spanish) using standard library protocols; completed original descriptive cataloging and subject analysis; enhanced brief catalog and authority records in III.

1997-1999 **Development Officer, McREL International, Aurora, Colorado.** Served as senior member of corporate management team in strategic planning, development of proposals and contracts, implementation, and evaluation of new services, products, and programs for educational agencies. Provided creative leadership to corporate committees to solicit ideas, identify goals and objectives, plan, develop, present, and evaluate professional development opportunities.

1995-1997 **Education Specialist, Education Service Center, Region One, Edinburg, Texas.** Served as member of Administrative Cabinet team in strategic planning, development of proposals and contracts, implementation, and evaluation of telecommunications capabilities, services, products, and programs for 40 school districts serving 283,000 students in 7 counties. Provided creative leadership to regional and state committees to solicit ideas, identify strategic goals and objectives, plan, develop, present, and evaluate funding opportunities and professional development for 400 librarians.

1993-1996 **Assistant Professor of Library Science, Sam Houston State University, Huntsville, Texas.** Served a faculty member teaching 400 graduate students in technical services (cataloging, bibliographic control, classification), automation, and networking. Participated in distance education program and coordinated annual conference. Conducted university and Texas Library Association-funded field research focused on library collection development and academic achievement.

1992-1993	**Head Librarian, Rocky Mountain College of Art & Design, Denver, Colorado.** Responsible for the daily operation, selection and acquisition of materials, formulation of policies for library operations, media center, and photography/slides archives. Designed and implemented library automation and delivery of electronic resources to college community.
1981-1985	**Development Officer, PRLC, Inc., Pittsburgh, Pennsylvania.** Served as senior member of corporate management team in strategic planning, development of proposals and contracts, implementation, and evaluation of new services, products, and programs for 100 institutional member organizations. Coordinated the development of proposals and contracts totaling $4,000,000 annually. Provided creative leadership to corporate committees to solicit ideas, identify goals and objectives, plan, develop, present, and evaluate professional development opportunities.
1981	**Director of Library Development, Pennsylvania Department of Education, Harrisburg, Pennsylvania.** Responsible for statewide development, technical assistance, professional development, resource sharing, children's services, institutional library services, networking, and state aid program for all libraries throughout the Commonwealth. Functioned as liaison to Governor's Advisory Council, LSCA Advisory Council, District Administrators, private colleges, universities, consortia managers, and network directors. Supervised $14,000,000 formula-based state aid program and $3,000,000 grant awards to individual libraries, consortia, and networks.
1978-1981	**Assistant Director, Southern Tier Library System, Corning, New York.** Coordinated operation of system-wide programs (technical assistance, professional development, resource sharing, technical services, outreach) to 40 public libraries in 5 counties serving 500,000 residents. Solicited ideas, identified goals, sponsored, and evaluated professional development opportunities and technical assistance sessions.
1976-1978	**Division Librarian for Technical Services, Corpus Christi Public Libraries, Corpus Christi, Texas.** Provided leadership in acquisitions, cataloging, serials control, and processing for main library and 4 branches serving 250,000 residents. Participated as senior member of library management team. Compiled and prepared technical evaluations, reports, and statistical analyses of Division operations to measure the achievement and cost of annual goals, objectives, and staff performance.
1975-1976	**System Coordinator, San Antonio Major Resource Center, San Antonio, Texas.** Served as senior member of the management team for District X Office, charged to provide technical assistance, resource sharing, media services, and professional development to librarians and staff representing 30 public library jurisdictions in 21 counties serving 1,500,000 residents. Functioned as liaison to System Director, staff, and members of governing bodies with the System Board of Directors and the Texas State Library and Historical Commission. Prepared LSCA grant applications and monitored awards totaling $1,100,000 annually.
1973-1975	**Bilingual Branch Librarian, San Antonio Public Library, San Antonio, Texas.** Worked as librarian providing reference, information, and readers' advisory services in branch serving 50,000 Spanish-speaking residents in southwest San Antonio. Participated in collection development and resource acquisition activities, specializing in children's work, Spanish language resources, and multicultural studies.
1972-1973	**Librarian, Holding Institute, Laredo, Texas.** Worked as high school librarian serving 500 boarding students in Spanish-speaking environment of private school. Provided reference, research assistance, and library instruction to students and 35 faculty members.
1966-1971	**Rockford Public Library, Rockford, Illinois.** Worked in branches as part-time as a Library Assistant, Clerk, and Page in city library serving 150,000 residents.

PUBLICATIONS

Editor-reviewed Monographs (Completed and in Progress)

Hall-Ellis, Sylvia D., and Mary Beth Weber. *Contemporary Cataloging in an RDA Environment: A Handbook for Students and Practitioners.* Chicago, IL: American Library Association. Under contract & In development.

RDA Testing: Lessons Learned and Challenges Revealed. Sylvia D. Hall-Ellis and Robert O. Ellett, Jr., eds. Binghamton, N.Y.: Haworth, 2012. 128 p.

Hall-Ellis, Sylvia D., Stacey L. Bowers, Christopher D. Hudson, and M. Claire Williamson. *Librarian's Handbook for Seeking, Writing, and Managing Grants.* Santa Barbara, Calif.: Libraries Unlimited, 2011. 315 p.

Hall-Ellis, Sylvia D., with Ann Jerabek, and Merrie W. Valliant. *Contemporary Cataloging: A Handbook for Practitioners and Students.* Open access text. Athens, GA: University of Georgia System Regents, 2011. 767 p.

Grealy, Deborah S. and Sylvia D. Hall-Ellis. *From Research to Practice: The Scholarship of Teaching and Learning in LIS Education.* Westport, Conn.: Libraries Unlimited, 2009. 175 p.

Hall-Ellis, Sylvia D. with J. Ann Jerabek. *Grants for School Libraries.* Westport, Conn.: Libraries Unlimited, 2003. 197 p.

Hall-Ellis, Sylvia D., Doris Meyer, Frank W. Hoffmann, with J. Ann Jerabek. *Grant Writing for Small Libraries and School Library Media Centers.* Boulder, Colo.: NetLibrary, 2001. 173 p.

Hall-Ellis, Sylvia D., Doris Meyer, Frank W. Hoffmann, with J. Ann Jerabek. *Grant Writing for Small Libraries and School Library Media Centers.* Englewood, Colo: Libraries Unlimited, 2000. 173 p.

Editor-reviewed Chapters (Completed and In Progress)

Hall-Ellis, Sylvia D. "Prepared to Lead: Talent, Skills, and Competencies." In *Telling the Technical Services Story.* Chicago, IL: American Library Association, 2020. In development and review.

Hall-Ellis, Sylvia D. "Grant Writing and Sponsored Research Funding for Academic Librarians." In *The New Librarianship.* Vol. 4. Bradford Lee Eden, ed. New York: Scarecrow Press, 2015. (pp. 163-174)

Hall-Ellis, Sylvia D. "Organizing Information: Technical Services." In *Information Services Today: An Introduction.* Sandra Hirsch, ed. Lantham, Md.: Rowman and Littlefield, 2015. (pp. 139-148)

Hall-Ellis, Sylvia D. "Metadata, MARC, and More." In *Rethinking Technical Services, Considering Our Profession and Ourselves: What's the Future of Our Profession?* Mary Beth Weber, ed. Lantham, Md.: Rowman and Littlefield, 2015. (pp. 29-55)

Hall-Ellis, Sylvia D., ed. "Contingent Faculty: Non-Tenure Track Faculty Series." In the *Faculty Personnel Guidelines Relating to Appointment, Promotion, and Tenure.* November 2011. Denver, Colo.: University of Denver, 2011. 42 p.

Hall-Ellis, Sylvia D., ed. "Standard VII: Information Resources." In the *Sturm College of Law Self-Study Presentation for Accreditation by the American Bar Association.* Denver, Colo.: University of Denver, Sturm College of Law, 2011. 20 p.

Hall-Ellis, Sylvia D. "Applying for Grants from Foundations, Corporations, or Government." In *The Volunteers' Guide to Fundraising: Raise Money for Your School Team, Library or Community Group.* 1st ed. Ilona M. Bray, ed. Berkeley, Calif.: Nolo, 2011. (pp. 1-38 on accompanying disc)

Hall-Ellis, Sylvia D., ed. "Standard III: The Faculty." In the *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association.* Denver, Colo.: University of Denver, Morgridge College of Education, 2010. 22 p.

Hall-Ellis, Sylvia D. "Library and Information Science Programs and Education for Catalogers and Metadata Specialists: Challenges for the Twenty-first Century." In *Conversations with Catalogers in the Twenty-First Century.* Elaine R. Sanchez, ed. Santa Barbara, Calif.: ABC-Clio, 2010. (pp. 226-254)

Hall-Ellis, Sylvia D., ed. "Standard III: The Faculty." In the *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association.* Denver, Colo.: University of Denver, College of Education, 2003. 15 p.

Lesesne, Teri S. and Sylvia D. Hall-Ellis. "The Selection, Evaluation, and Integration of Culturally Authentic Texts: A Case for Making the Online Catalog Reflect Parallel Cultures." In *Literacy: Traditional, Cultural, Technological.* Pittsburgh, Pa.: International Association of School Librarianship, 1995. (pp. 110-113)

Lesesne, Teri S. and Sylvia D. Hall-Ellis. *The Selection, Evaluation, and Integration of Culturally Authentic Texts: A Case for Making the Online Catalog Reflect Parallel Cultures.* In *Conference Proceedings* of the 23rd Annual International Association of School Librarianship, Pittsburgh, Pennsylvania, July 17-22, 1994. ERIC Document ED374816. 17 p.

Hall-Ellis, Sylvia D. "Curriculum Folio for School Library Media Specialist Programs." In the *National Council for Accreditation of Teacher Education Self-Study for Sam Houston State University.* Huntsville, Tex.: Sam Houston State University, College of Education and Applied Science, 1994. 25 p.

Bruntjen, Scott and Sylvia D. Hall. "Attempting to Automate: Lessons Learned Over Five Years." In *Advances in Library Administration*. Volume 4. Weston, Conn.: JAI Press, 1985. (pp. 177-192)

Peer-Reviewed Journal Articles (Completed and In Progress)

Hall-Ellis-Sylvia D. "Job Design for Cataloging and Metadata Librarians." Submitted to *Journal of Library Administration.* In progress.

Hall-Ellis, Sylvia D. "The Relationship of Situational Leadership and the Dreyfus Model of Skill Acquisition for Supervisors in Cataloging and Metadata Services." To be submitted to *Library Resources and Technical Services.* In progress.

Hall-Ellis, Sylvia D. *The Relationship of Core Competencies on Learning Outcomes and Employers' Expectations for Catalog Librarians and Metadata Specialists.* To be submitted to the *Journal of Education for Library and Information Science.* In progress.

Hall-Ellis, Sylvia D. "Building Cataloger Competencies: The Dreyfus Model as a Prototype for the Education and Professional Development of Catalog Librarians and Metadata Specialists in Bibliographic Control." To be submitted to the *Cataloging & Classification Quarterly.* In progress.

Hall-Ellis, Sylvia D. "Descriptive Impressions of Entry-Level Cataloger and Metadata Specialists Positions, 2000-2016: Reflections and Trends." To be submitted to *Library Resources and Technical Services.* In progress.

Hall-Ellis, Sylvia D. "Competencies for Metadata and Cataloging Leaders: What Employers Expect as Reflected in Position Descriptions, 2000-2016." *Cataloging & Classification Quarterly.* In progress.

Hall-Ellis, Sylvia D. "Stackable Micro-credentials – A Framework for the Future." *The Bottom Line 29*, no. 4 (April 2016), http://www.emeraldinsight.com/doi/pdfplus/10.1108/BL-02-2016-0006

Hall-Ellis, Sylvia D. "Succession Planning and Staff Development – A Winning Combination." *The Bottom Line*, 28 no. 3 (May 2015): 95-98.

Hall-Ellis, Sylvia D. "Succession Planning and Staff Development – A Winning Combination." *The Informed Librarian Online* (September 2015), http://www.informedlibrarian.com/featuredArticle.cfm?FILE=succession_1509.pdf

Hall-Ellis, Sylvia D. "Nudges and Decision Making: A Winning Combination." *The Bottom Line 28*, no. 4 (July 2015): 133-136.

Hall-Ellis, Sylvia D. "Succession Planning and Staff Development – A Winning Combination." *The Bottom Line 28*, no. 3 (May 2015): 95-98.

Hall-Ellis, Sylvia D. "Metadata Competencies for Entry-Level Positions: What Employers Expect as Reflected in Position Descriptions, 2000-2013." *Journal of Library Metadata 15*, no. 2 (June 2015): 102-134.

Hall-Ellis, Sylvia D. "Onboarding to Improve Library Retention and Productivity." *The Bottom Line 27*, no. 4 (October 2014).

Hall-Ellis, Sylvia D. "Accept, Coach, and Inspire: A Formula for Success." *The Bottom Line 27*, no. 3 (July 2014): 103-106.

Hall-Ellis, Sylvia D. *Reward Systems for Staff in Higher Education Institutions: Case Studies of 15 Higher Education Institutions in Colorado.* Capstone Thesis. Denver, Colo.: University of Denver, University College, 2014.

Hall-Ellis, Sylvia D. "Reward Systems Promote High-Performance Work Teams Achieving Library Mission." *The Bottom Line 27*, no. 2 (April 2014): 66-69.

Hall-Ellis, Sylvia D. "Investment in Staff Development for Seeking External Grant Funds." *The Bottom Line 27*, no. 1 (January 2014): 22-25.

Hall-Ellis, Sylvia D. "Staff Investments for High Performance Teams." *The Bottom Line 26*, no. 4 (October 2013): 149-152.

Hall-Ellis, Sylvia D. and Deborah S. Grealy. "Service Learning in LIS Education: A Case for Holistic Assessment." *Catholic Library World* 82, no. 3 (March 2012): 178-187.

Hall-Ellis, Sylvia D., and Deborah S. Grealy. "The Dreyfus Model of Skill Acquisition: A Career Development Framework for Succession Planning and Management in Academic Libraries." *College and Research Libraries* (June 2012). http://crl.acrl.org/content/early/2012/06/19/crl12-349.full.pdf+html

Hall-Ellis, Sylvia D. "Cataloger Competencies…What Do Employers Require?" *Cataloging & Classification Quarterly* 46, no. 3 (2008): 305-330. http://ccq.haworthpress.com/doi:10.1080/01639370802034565

Hall-Ellis, Sylvia D. "Subject Access for Readers' Advisory Services: Their Impact on Contemporary Spanish Fiction in Selected Public Library Collections." *Public Library Quarterly* 27, no.1 (2007): 1-18. http://plq.haworthpress.com/doi:10.1080/01616840802122377

Hall-Ellis, Sylvia D. "Language Proficiencies among Catalogers and Technical Services Librarians." *Technical Services Quarterly* 25, no. 2 (2007) 31-47. http://tsq.haworthpress.com/doi:10.1300/J124v25n02_03

Hall-Ellis, Sylvia D. "Cataloging Electronic Resources & Metadata: Employers' Expectations as Reflected in *American Libraries and AutoCAT,* 2000-2005." *Journal of Education for Library and Information Science* 47, no. 1 (2006): 38-51.

Hall-Ellis, Sylvia D. "Descriptive Impressions of Managerial & Supervisory Cataloger Positions as Reflected in *American Libraries, AutoCAT,* and the *Colorado State Library Jobline,* 2000-2004: A Content Analysis of Education, Competencies, and Experience." *Cataloging & Classification Quarterly* 42, no. 1 (2006): 55-88.

Hall-Ellis, Sylvia D. "Descriptive Cataloging Proficiencies among Beginning Students: A Comparison among Traditional and Virtual Class Students." *Journal of Library & Information Services in Distance Education* 2, no. 2 (2005): 13-44.

Hall-Ellis, Sylvia D. "Descriptive Impressions of Entry-Level Cataloger Positions as Reflected in *American Libraries, AutoCAT,* and the *Colorado State Library Jobline,* 2000-2003." *Cataloging & Classification Quarterly* 40, no. 2 (2005): 33-72.

Hall-Ellis, Sylvia D. and Mary Ann Berry. "The Impact of Library Resource Centers on Academic Achievement in Selected Public Schools in South Texas: Preliminary Findings." *Texas Library Journal* 71, no. 2 (Summer 1995): 94-97.

Hall-Ellis, Sylvia D. "Visual Arts as Foundation for Successful Library Automation: The Rocky Mountain College of Art & Design Experience." *Art Reference Serials Quarterly* 4 (1994): 43-54.

Hall, Sylvia D. *High School Databases: An Examination of Bibliographic Data Elements*. Dissertation. Pittsburgh, Pa.: University of Pittsburgh, School of Library and Information Science, 1985.

Technical Reports

Hall-Ellis, Sylvia D. *Colorado Community College System Grant Handbook.* Denver, CO: Colorado Community College System, 2015.

.

Faculty Voices: Testimonials from the Frontiers of a Development Education Redesign. Elaine DeLott Baker, Marilyn Smith, and Sylvia D. Hall-Ellis, eds. Denver, CO: Colorado Community College System, 2014.

Lewis, Melanie, Jamie M. Van Leeuwen, and Sylvia D. Hall-Ellis. *Project Homeless Connect 7 Event Evaluation.* Denver, Colo.: University of Colorado Denver, 2009.

Hall-Ellis, Sylvia D. *Grant Writing Tips and Resources.* Denver, Colo.: University of Denver, 2008-2014. http://portfolio.du.edu/grants

Hall-Ellis, Sylvia D. and Duan Zhang. *Project Homeless Connect 6 Event Evaluation.* Denver, Colo.: Center for Community Engagement and Service Learning, 2008.

Hall-Ellis, Sylvia D. *Strategic Plan for Denver Medical Library System: Final Report.* Denver, Colo.: HealthONE and Denver Medical Library, 2008.

Hall-Ellis, Sylvia D. and Duan Zhang. *Project Homeless Connect 5 Event Evaluation.* Denver, Colo.: Center for Community Engagement and Service Learning, 2008.

Hall-Ellis, Sylvia D. *Strategic Plan for Denver Medical Library System: Phase Two Report.* Denver, Colo.: HealthONE and Denver Medical Library, 2007.

Hall-Ellis, Sylvia D. and Duan Zhang. *Project Homeless Connect 4 Event Evaluation.* Denver, Colo.: Center for Community Engagement and Service Learning, 2007.

Hall-Ellis, Sylvia D. *Strategic Plan for Denver Medical Library System: Phase One Report.* Denver, Colo.: HealthONE and Denver Medical Library, 2007.

Hall-Ellis, Sylvia D. *Denver Public Library: Opportunities for Change.* Denver, Colo.: University of Denver, Colorado Community-Based Research Network & Metropolitan Organization of People, 2005.

Hall-Ellis, Sylvia D. *Resources in Grant Writing for Colorado Schools and Libraries.* Sponsored by the Colorado Department of Education, Educational Telecommunications Unit, 1999. Revised September 2000.

Hall-Ellis, Sylvia D. *Lewis & Clark Fellows Concept Paper for the Bicentennial Commemoration.* Vancouver, Wash.: Lewis and Clark Bicentennial Council, 1998.

Hall-Ellis, Sylvia D. and Mary Ann Berry. *The Impact of Library Resource Centers on Academic Achievement in Public Schools in Texas: The Final Report.* Prepared under the sponsorship of the Research Enhancement Fund award from Sam Houston State University, the Texas Association of School Librarians, the Children's Services Round Table, and the Young Adults' Services Round Table. Huntsville, Tex.: Sam Houston State University, College of Education and Applied Science, 1997.

Hall-Ellis, Sylvia D. *The Impact of Library Resource Centers on Academic Achievement in Selected Public Schools in Texas.* Huntsville, Tex.: Sam Houston State University, College of Education and Applied Science, 1997.

Hall-Ellis, Sylvia D. *Planning for Educational Technology: the Distance Education Implementation Phase: An Investigation of Technology and Telecommunication Needs for Region One Education Service Center, South Texas Community College, TCI Cablevision of South Texas, Inc., Tech Prep of the Rio Grande Valley, Inc., Texas State Technical College, Centro Cultural en el Valle de Texas, McAllen Independent School District, Harlingen Consolidated Independent School District, Donna Independent School District, Edinburg Consolidated Independent School District, La Joya Independent School District, Mercedes Independent School District, Pharr-San Juan-Alamo Independent School District, Sharyland Independent School District, Easter Seal Society of the Rio Grande Valley, Inc., and McAllen Memorial Library.* Harlingen, Tex.: Tech Prep of the Rio Grande Valley, 1996.

Hall-Ellis, Sylvia D. "Barcodes for the Texas Library Connection." Prepared under the sponsorship of the Office of Technology Services, Texas Education Agency. Austin, Tex.: TENET News and Conferences, March 1995.

Hall-Ellis, Sylvia D. *[et. al.] Planning for Educational Technology: A Study of Future Technology Needs for the Donna Independent School District, Region One Education Region Service Center, Tech Prep of the Rio Grande Valley, Inc. South Texas Community College, and the Department of Library Science, Sam Houston State University.* Research team: Sylvia D. Hall-Ellis, Ruth Ann Riggins, Noe Torres, Patricia G. Bubb, and Michael D. Bell. Harlingen, Tex.: Tech Prep of the Rio Grande Valley, 1994. Funded through a grant award from the Texas Education Agency.

Hall-Ellis, Sylvia D. and Mary Ann Berry. *The Impact of Library Resource Centers on Academic Achievement in Selected Public Schools in South Texas.* Huntsville, Tex.: Sam Houston State University, College of Education and Applied Science, 1993. Funded through a grant award from the Sam Houston State University Research Enhancement Fund.

Hall-Ellis, Sylvia D., Alice Krueger, and Linda Brannan. *Thesaurus Development for National Curriculum Materials Database at the Eisenhower National Clearinghouse for Mathematics and Science Education, Columbus, Ohio.* Developed under the sponsorship of Mid-Continent Regional Educational Laboratory Corporate Library, Aurora, Colo. Washington, D.C.: National Science Foundation; Columbus, Ohio: ERIC Clearinghouse for Mathematics and Science Education, 1993.

Hall-Ellis, Sylvia D. and Clarence L. Coffindaffer. *Partners in Progress: The 1992 Montana Library Federation Study Final Report*. Helena, Mont: Montana State Library and The Blue Bear Group, Inc., 1992.

Hall-Ellis, Sylvia D. and Clarence L. Coffindaffer. *Montana Federation Study Identification Paper*. Denver, Colo.: The Blue Bear Group, Inc., 1992.

Bruntjen, Scott and Sylvia D. Hall. *Recommendations on the Development of System-Wide Efforts of the Northern Waters Library Service.* Ashland, Wisc.: Northern Waters Library Service and The Blue Bear Group, Inc., 1989.

Hall, Sylvia D. *The Iowa Union List of Serials as A Part of The Iowa Locator: Planning Document*. Des Moines, Ia.: State Library of Iowa, 1988.

Hall, Sylvia D. *Iowa Computer-Assisted Network (ICAN Users' Manual. IBM Compatible Edition: Version 7.0.* Des Moines, Ia.: State Library of Iowa, 1988.

Hall, Sylvia D. *Iowa Computer-Assisted Network (ICAN Users' Manual. IBM Compatible Edition: Version 5.0.* Des Moines, Ia.: State Library of Iowa, 1986.

Hall, Sylvia D. *Iowa Computer-Assisted Network (ICAN Users' Manual. Apple MacIntosh Edition: Version 5.0.* Des Moines, Ia.: State Library of Iowa, 1986.

Bruntjen, Scott and Sylvia D. Hall. *The District of Columbia Public Library Planning for Retrospective Conversion of Bibliographic Records and Automation Issues for the 1990's: Final Report*. Central City, Colo.: The Blue Bear Group, Inc., 1986.

Bruntjen, Scott and Sylvia D. Hall. *The Iowa Locator: Final Report.* Central City, Colo.: The Blue Bear Group, Inc., 1986.

Bruntjen, Scott and Sylvia D. Hall. *Discovering Resource Sharing Treasures: Accessing Pennsylvania's High School Library Collections*. Harrisburg, Pa.: State Library of Pennsylvania, Bureau of School Library Media Resources, 1985.

Hall, Sylvia D. *Retrospective Conversion and the Completeness of Cataloging in Pennsylvania High School Libraries.* Pittsburgh, Pa.: PRLC, Inc., 1985.

Bruntjen, Scott and Sylvia D. Hall. *The Iowa Locator.* Des Moines, Ia.: State Library of Iowa, 1985- http://www.statelibraryofiowa.org/ld/slfo/locator

Bruntjen, Scott and Sylvia D. Hall. *ICAN Analysis*. Des Moines, Ia.: State Library of Iowa, 1985.

Bruntjen, Scott and Sylvia D. Hall. *Zebras: Bar Codes, Standards, and Library Automation*. Central City, Colo.: The Blue Bear Group, Inc., 1985.

Hall, Sylvia D. and Scott Bruntjen. *Access Pennsylvania.* Harrisburg, Pa.: State Library of Pennsylvania, 1984- http://www.accesspa.org

Hall, Sylvia D. *Reconfiguration of the ITAC Cataloging Data Base to Machine Readable (MARC) Format*. Classified Document. Arlington, Va.: U. S. Army Intelligence and Threat Analysis Center, 1984.

Bruntjen, Scott and Sylvia D. Hall. *South Dakota Library Automation Plan: Final Report*. Pittsburgh, Pa.: PRLC, Inc., 1982. Reprinted by the South Dakota State Library, 1983.

Bruntjen, Scott and Sylvia D. Hall. *Altoona Automation Plan: Final Report prepared for the Altoona Area Public Library and the Altoona District Library Center*. Pittsburgh, Pa.: PRLC, 1983.

Bruntjen, Scott, Sylvia D. Hall and Paul J. Pugliese. *College Space Utilization Study*. Pittsburgh, Pa.: PRLC, Inc., 1983.

Hall, Sylvia D. *Pennsylvania Public Library Film Center Union Catalog Cataloging and Technical Services Procedures Manual*. Pittsburgh, Pa.: PRLC, Inc., 1982.

Hall, Sylvia D. *History of Library Development in Pennsylvania: A Review and Content Analysis of the Plans, 1955-1980.* Pittsburgh, Pa.: PRLC, Inc., 1982.

Hall, Sylvia D. *Newspapers on Microfilm: A Listing of Newspapers in Allegheny, Chemung, Schuyler,Steuben, and Yates Counties*. Ithaca, N.Y.: South Central Reference and Research Council, 1981.

Hall, Sylvia D. *OCLC/ILL Subsystem at Southern Tier Library System.* Albany, N.Y.: SUNY/OCLC Network, 1980.

Hall, Sylvia D. *Retrospective Conversion for Major Public Libraries in South Central New York State: A Technical Report*. 2 vols. Corning, N.Y.: Southern Tier Library System, 1979-1980.

Hall, Sylvia D. *The Librarian and the Mexican-American Patron in Texas.* Austin, Tex.: Texas State Library, 1978.

Other Publications

Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association.* St. Paul, MN: St. Catherine University, 2017. 255 p.

Hall-Ellis, Sylvia D., and Robert O. Ellett, Jr., guest editors. *RDA Testing: Lessons Learned and Challenges Revealed: A Special Issue of Cataloging & Classification Quarterly* 49, no. 7-8 (November 2011).

Hall-Ellis, Sylvia D. and Robert O. Ellett, Jr. *Fundamentals of Cataloging I: An Online ALCTS Course.* Chicago, Ill.: American Library Association, 2009.

Hall-Ellis, Sylvia D., Angela J. Kinney and Martha Ann Bace, editors. *Best Practices in Cataloging Education Identified at the 2007 ALCTS Pre-conference: A Special Issue of the Journal for Library and Information Science Education* 49, no. 2 (Spring 2008).

Hall-Ellis, Sylvia D. *Cataloging Education: A New Emphasis for the Library and Information Science Curriculum.* Presentation delivered to the ALCTS Pre-conference, 132nd Annual Conference, American Library Association, Washington, D.C., June 22, 2007. http://www.loc.gov/aba/professional/

Hall-Ellis, Sylvia D. "*Cataloger Competencies: Do the Employers Require What the Professors Teach?*" Presentation delivered to the ALCTS CCS Heads of Cataloging Discussion Group, 131st Annual Conference, American Library Association, New Orleans, Louisiana, June 26, 2006. http://dlist.sir.arizona.edu/

Hall-Ellis, Sylvia D., guest editor, with Deborah S. Grealy, Clara L. Sitter, and Tassos Petrou. *Library Students' Voices in Colorado: An Issue of Colorado Libraries* 31:1 (Spring 2005).

Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association.* Denver, Colo.: University of Denver, College of Education, 2003. 125 p. with CD-ROM.

Hall-Ellis, Sylvia D. and Deborah S. Grealy, guest editors. *Library Education in Colorado: An Issue of Colorado Libraries* 29, no. 2 (Summer 2003).

Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association.* Denver, Colo.: University of Denver, College of Education, 2003.

Hall-Ellis, Sylvia D. "Adaptive Furniture for the ADA-Qualified: An Investigative Study." 2001. Unpublished manuscript.

Hall-Ellis, Sylvia D. "Cataloging Electronic Serials: A Study in User Access and Article Delivery." 2001. Unpublished manuscript.

Hall-Ellis, Sylvia D. "E-books: Acquisition and Circulation Issues in Public Libraries." 2001. Unpublished manuscript.

Hall-Ellis, Sylvia D. "Authority Control and Its Importance in School Library Media Online Catalogs." 1999. Unpublished manuscript.

Hall, Sylvia D. "Cooperation Improving Service in San Antonio System." *Texas Libraries* XXXVII (Spring 1976): 27-29.

Hall, Sylvia D. "Summer is Sea Monstrously Fun!" *Texas Library Journal* L (December 1974): 224-226.

Hall, Sylvia D. "A Brief Annotated Bibliography of Non-book Materials on the Mexican American." *Texas Library Journal* XLVIII (October 1972): 213-232.

Book Reviews

ALA Book of Library Grant Money. Ann Kepler, ed. 8th ed. Chicago, IL: American Library Association, 2011. 978-0-8389-1058-0, $175.00 [*Colorado Libraries,* June 2012]

Developing a Compensation Plan for Your Library. Paula M. Singer and Laura L. Francisco, comps. 2nd ed. Chicago, IL: American Library Association, 2009. 978-0-8389-0985-0, $57.00 [*Colorado Libraries,* March 2012]

Frugal Librarian. Carol Smallwood, ed. Chicago, IL: American Library Association, 2011. 978-0-83891-075-7, $42.00 [*Colorado Libraries,* February 2012]

Dowlin, Ken. *Getting the Money: How to Succeed in Fundraising for Public and Nonprofit Libraries.* Westport, Conn.: Libraries Unlimited, 2008. 978-1-59158-597-X, $50.00 [*Colorado Libraries,* February 2009]

Lundahl, Mats. *Bebo de Cuba; Bebo Valdés y Su Mundo*, by Mats Lundahl. Traducción de Linda Oakeshott Dragó. Barcelona: RBA Libros, 2008. 978-849-86-7259-6, $33.99 [*Críticas*, November 2008]

Bonelli, Florencia. *Marlene*, by Florencia Bonelli. Buenos Aires: Aguilar, Altea, Taurus, Alfaguara, 2008. 978-987-04-0995-3, $24.95 [*Críticas*, October 2008]

Keret, Etgar. Traducción de Ana María Bejarano. *Pizzería Kamikaze y Otras Relatos.* México, D.F.: Editorial Sexto Piso, 2008. 968-5679-29-0, no price given [*Críticas*, September 2008]

Littauer, Marita, and Florence Littauer. *Enriquece tu Comunicación. Miami, Fla.: Editorial Unilit, 2008.* 978-0-7899-1521-4, $11.99 [*Críticas*, August 2008]

Martínez, Guillermo. *La Muerte Lenta de Luciana B.* New York: Rayo Planeta, 2008. 978-0-06-156551-9, $14.95 [*Críticas*, May 2008]

Matthews, Joseph R. *The Evaluation and Measurement of Library Services.* Westport, Conn.: Libraries Unlimited, 2007. 978-1-59158-532-9, $50.00 [*Colorado Libraries,* April 2008]

Gómez-Jurado, Juan. *A Masacre de Virginia Tech: Anatomía de una Mente Torturada.* Barcelona: Ediciones El Andén, 2007. 978-84-935789-4-7, €16.50 [*Críticas*, February, 2008]

Matthews, Joseph R. *Library Assessment in Higher Education.* Westport, Conn.: Libraries Unlimited, 2007. 978-1-59158-531-2, $45.00 [*Colorado Libraries,* February 2008]

Chacón, Inma. *La Princesa India: Cuando el Viento Azul.* México: Alfaguara, 2006. 970-770398-9, $19.95 [*Críticas*, June 2006]

Vargas Llosa, Mario. *Israel/Palestina: Paz o Guerra Santa.* fotogs. by Morgana Vargas Llosa. Spain/U.S.: Aguilar: Santillana, 2006. 84-03-09691-7, paper, $19.95 [*Críticas*, May 2006]

Findlay, Diane. *Digging into Dewey.* Fort Atkinson, Wisc.: Upstart Books, 2005. 978-1-932146-18-9, $16.95 [*Colorado Libraries,* May 2006]

Pequeñas Resistencias 4: Antología del Nuevo Cuento Norteamericano y Caribeño. Menéndez, Ronaldo, Ignacio Padilla & Enrique del Risco, eds. Madrid**:** Páginas de Espuma, 2005. 84-95642-59-X, paper, $36.95 [*Críticas*, April 2006]

Rueda, André. *Vengo a Salvar a España: Biografía de Un Franco Desconocido.* Madrid**:** Nowtilus (Investigación Abierta), 2005. photogs., bibliog. 84-9763-202-8, paper, $23.95 [*Críticas*, March 2006]

Becerra, Ángela. *El Penúltimo Sueño.* 1st ed. Barcelona: Planeta, 2005. 84-08-05795-2, $25.95. [*Críticas*, January 2006]

Aguilar Camín, Héctor. *La Conspiración de la Fortuna.* 1st ed. México: Planeta, 2005. 970-37-0368-2, $19.95 [*Críticas*, October 2005]

México en Sus Libros. Enrique Florescano with Pablo Mijangos, editors & compilers. 1st ed. México, D.F.: Taurus Aguilar, 2004. 968-19-0783-3, $16.95 [*Críticas*, September 2005]

Historia Económica de México. Coordinator, Enrique Semo. México: Universidad Nacional Autónoma de México, Editorial Oceano de México, 2004. 13 vols. 970-32-0805+3 (obra completa), $9.25 each [*Críticas*, August 2005]

Czarnowsky, Christyne A. and Michael H. Williams. *Managing Your Wired Workforce: A Practical Guide.* Denver, Colo.: Bradford Pub., 2003. 1-883726-79-4, $24.95 [*Colorado Libraries,* 2004]

Abraham Lincoln. Colección Grandes Biografías. Madrid: Edimat Libros, 2003. 84-8403-858-0, €4.95 [*Críticas*, 2004]

Alponte, Juan María. *Colón: el Hombre, el Navegante, la Leyenda.* México, D.F.: Aguilar, 2003. 968-19-1260-8, $16.95 [*Críticas*, 2004]

Olcese Salvatecci, Alfieri. *Cómo Estudiar con Éxito: Téchnicas yHhábitos Para Aprender Major.* México, D.F.: Alfaomega Grupo Editor, 2002. 970-15-0764-9 [*Críticas*, 2003]

Mujeres Como Islas: Antología de Narradoras Cubanas, Dominicanas, Puertoriqueñas. Edición: Olga Marta Pérez, Thelma Jiménez, Andrés Blanco D., eds. Santo Domingo, República Dominicana: Ediciones Ferilibro, 2002. 959-209-419-5 [*Críticas*, 2003]

Tibol, Raquel. *Los Murales de Diego Rivera: Universidad Autónoma Chapingo.* México, D.F.: Editorial RM, Universidad Autónoma Chapingo, 2002. 968-5208-08-5 [*Críticas*, 2003]

Lindsay-Poland, John. *Emperors in the Jungle: the Hidden History of the U.S. in Panama.* Durham, N.C.: Duke University Press, 2003. 0-8223-3098-9, $18.95 [*Críticas*, 2003]

Kaplan, Allison G. and Ann Marlowe Riedling. *Catalog It! A Guide to Cataloging School Library Materials.* Worthington, Ohio: Linworth Pub., 2002. 1-58683-014-7, $44.95 [*Colorado Libraries,* 2002]

Bardach, Ann Louise. *Cuba Confidential: Love and Vengeance in Miami and Havana.* New York: Random, 2002. 0-375-50489-3, $25.95 [*Library Journal*, 2002]

Swan, James. *Fundraising for Libraries: 25 Proven Ways to Get More Money for Your Library.* New York: Neal Schuman, 2002. 1-55570-433-6, $69.95 [*Colorado Libraries,* 2002]

Puíg de Lange, Victoria. *Sol Con Agua.* Nashville, Tenn.: Editorial Vistazo/Ediciones Reio Negro, 2002. 0-9724506-0-2, $17.75 [*Críticas*, 2002]

The Encyclopedia of Latin American Politics. Edited by Diana Kapiszewski; assistant editor, Alexander Kazan. Westport, Conn.: Oryx Press, 2002. 1-57356-306-4, $74.95 [*Library Journal*, 2002]

Vargas Lizano, Isabel. *Y Si Quieres Saber de Mi Pasado.* Con la colaboración de J.C. Vales. Madrid: Santillana Ediciones Generales, S.L., 2002. 84-03-09278-4, $18.95 [*Críticas*, 2002]

Ruy Sánchez, Alberto. *Los Jardines Secretos de Mogador: Voces de Tierra.* México, D.F.: Alfaguara, S.A., 2001. 968-19-0879-1, $16.95 [*Críticas*, 2002]

Chávez, Ricardo and Celso Santajuliana. *El Final de las Nubes.* Barcelona: RBA Libros, S.A., 2001. 84-7901-760-0, $11.95 [*Críticas*, 2002]

The Power of Language / El Poder de la Palabra: Selected Papers from the Second REFORMA National Conference. Edited by Lillian Castillo-Speed and the REFORMA National Conference Publications Committee. Englewood, Colo.: Libraries Unlimited, 2001. 1-563089459, $35.00 [*Colorado Libraries,* 2001]

Martínez, Rubén. *Crossing Over: A Mexican Family on the Migrant Trail.* Metropolitan: Holt, 2001. 0-8050-4908-8, $26.00 [*Library Journal*, 2001]

Guevara, Ernesto "Che." *The African Dream: the Diaries of the Revolutionary War in the Congo.* Translated from the Spanish by Patrick Camiller. With an introduction by Richard Gott and a foreword by Aleida Guevara March. New York: Grove Press, 2001. 0-8021-3834-9, $13.00 [*Library Journal*, 2001]

Cooper, Gail and Garry Cooper. *New Virtual Field Trips.* Englewood, Colo.: Libraries Unlimited, 2001. 1-56308-887-8, $27.50 [*Colorado Libraries,* 2001]

NATIONAL SERVICE & PROFESSIONAL AFFILIATIONS

American Library Association
- Association for Library Collections & Technical Services
  - Cataloging, Classification, and Metadata Section
    - ALCTS Editorial Board, Member, 2015-2019
    - Committee on Cataloging: Description and Access, Member, 2008-2012
    - Test Site for RDA, Manager, 2009-2011
    - Nominating Committee, Member, 2007-2008
    - Committee for Education and Training of Catalogers, Chair, 2006-2007, Member, 2003-2007
    - Competencies & Education for a Career in Cataloging Interest Group, Chair, 2010-2012; Member, 2010-
  - President's Program Committee, 2015-2016
  - Education Committee, Member, 2005-2007
  - Task Force on Competencies and Education for a Career in Cataloging, Chair, 2007-2009
  - 2007 Annual Meeting Pre-conference, "What They Don't Teach in Library School: Competencies, Education and Employer Expectations for a Career in Cataloging," Steering Committee Chair
  - Fundraising Committee, Member, 2004-2006
- Office of Diversity
  - Committee on Diversity, Chair, 2010-2011
  - Diversity Research Advisory Committee, Member, 2009-2010
  - Spectrum Scholar Mentor, 2009
  - Diversity Grants Review Committee, Member, 2007, 2010, 2012
- Office of Statistics and Research
  - Research and Statistics Committee, Member, 2005-2007; Intern, 2003-2005
  - Library Research Round Table
    - Membership Committee, 2019-2021
    - Board of Directors, Member At-Large, 2009-2012
    - Member, 2003-
- Office of Accreditation
  - Accreditation Panel Member (training completed 2004)
  - Accreditation Review Panel Member (training completed 2004)

REFORMA
- National Board of Directors, 2005-2013
- National Fundraising Chair, 2005-2013
- National Recruiting and Mentoring Committee, 2008-2010
- Colorado Chapter, Secretary, 2004-2005
- Colorado Chapter Liaison to National Board of Directors, 2004-2013

Online Audiovisual Catalogers Association (OLAC)
American Association of Law Libraries (AALL)
Colorado Association of Law Librarians (CoALL)

American Association of University Women
Association for Library and Information Science Education (ALISE)
- University of Denver ALISE Representative, 2003-2008; 2010-2011
- Membership Advisory Committee, 2007-2010
- Technical Services Education Special Interest Group, 2003-
- Garfield Doctoral Dissertation Award Reviewer, 2012
- Garfield Doctoral Dissertation Scholarship Reviewer, 2014

American Association of University Professors (AAUP), 2006-

American Grant Writers Association, 2014-
Grant Professionals Institute, 2016-
National Grants Management Association, 2014-
Grant Reviewer (National and Regional Team Leader), U.S. Department of Education, 1998-
Grant Reviewer, Broadband Technologies Opportunity Program, U.S. Department of Commerce, 2009
Grant Reviewer, Institute of Museum and Library Services, 1998-2000
Grant Reviewer, Colorado Department of Education, 2014-
Grant Reviewer, Colorado State Library, 1998-
Grant Reviewer, American Association of Community Colleges *Working Connections* program, 1999
Peer Reviewer, *Journal for Library and Information Science Education,* 2005-
Peer Reviewer, *Journal of Library Metadata,* 2009
Peer Reviewer, *International Journal of Library and Information Science,* 2011
Regular Columnist, *The Bottom Line*, 2013-2016
Book Reviewer, *Library Journal*, 2001
Book Reviewer, *Criticas*, 2002-2009
Book Reviewer, *Colorado Libraries,* 2000-2012

REGIONAL SERVICE & PROFESSIONAL AFFILIATIONS

Mountain Plains Library Association
Professional Development Grants Committee, Member, 2005-2006
Professional Development Policy and Guidelines Sub-committee Chair, 2006

Colorado Association of Libraries
"Student Voices" Column Editor, *Colorado Libraries,* 2005-2006
Conference Planning Committee, Member, 2002
Technical Services and Automation Division, Chair, 2002-2003; Member, 2000-
Academic Libraries Division, Peer Review Conference Papers Committee, Member, 2007-2009
Education Committee, Member, 1988-1993
Diverse Populations Committee, Member, 2009-2013

SERVICE TO THE UNIVERSITY OF DENVER

Center for Teaching and Learning Faculty Advisory Board, 2007-2008; 2010-2012
Center for Community Engagement and Service Learning Advisory Board, 2007-2012
Faculty Senate
Executive Committee, 2008-2013
Nominations, Rules & Credentials Committee, Chair, 2008-2013; Member 2007-2013
Appointment, Promotion, and Tenure Revision Committee, 2010-2012
Grievance Policy Committee Member, 2007-2010
Law Library Director Search Committee Chair, 2012-2013
PROF Grant Review, College Representative to the University Review Team, 2006, 2008
Project Homeless Connect Evaluations, Principal Investigator, 2006-2009
University Technology Council, 2007-2009
University of Denver Hyde Interviews for Incoming Freshmen, 2003-2012

SERVICE TO THE MORGRIDGE COLLEGE OF EDUCATION

Appointment, Tenure & Review Committee
Member, voting, 2003-2005; 2007- 2009
Chair, Clinical Faculty Promotion & Tenure Policy Subcommittee, 2008-2009
Member, Community Engagement Subcommittee, 2007-2009
Member, Tenure Review Panel, 2003
Advancement and Alumni Relations Committee, Chair, 2003-2007; Member, 2002-2007
College Building Committee Member, 2004-2010
College of Education Student Association, Faculty Advisor, 2004-2007; 2009-2010
Faculty Senator, 2007-2013
Research and Scholarship Committee, Chair, 2008-2009; Member, 2002-2003, 2008-2009

Research and Grant Mentoring Committee, Chair, 2009-2010
Research Task Force Member, 2010-2012
Search Committee Member, Assistant Professor for Curriculum & Instruction, 2010-2011
Workload Task Force Member, 2010-2011

SERVICE TO THE LIBRARY & INFORMATION SCIENCE PROGRAM
Library and Information Science Student and Alumni Association, Faculty Advisor, 2002-2008
ALA Student Chapter, Faculty Advisor, 2005-2008
Beta Phi Mu Phi Chapter, Faculty Advisor, 2004-2014
Steering Committee Member, Accreditation by the American Library Association, 2001-2004
Search Committee Member, Associate Professor for LIS, 2006
Search Committee Chair, Assistant Professors for LIS, 2003, 2005, 2006, 2007
Search Committee Ex-Officio Member, Director for LIS, 2004

SERVICE TO SAM HOUSTON STATE UNIVERSITY
Rio Roundup: South Texas Literature Conference, Conference Coordinator, 1993
External Relations, Fund Raising, Grants Committee, Chair, 1993-1995
Advisory Council Committee, Chair, 1993-1995
Students, Admissions, and Advisement Committee, Chair, 1994
Institutional Effectiveness Committee, Member, 1993-1995

SERVICE TO THE COLLEGE OF EDUCATION AND APPLIED SCIENCE
Faculty Affairs Committee, Member, 1993-1995
Curriculum Committee, Member, 1993-1995
Continuing Education Committee, Member, 1993-1995

SERVICE TO THE COMMUNITY
French Teacher & Student Exchange with Laon, France, Monarch High School, Boulder Valley School District (Boulder, Colorado), 2017-
Denver/Boulder Games 2022, Board of Directors, Secretary-Treasurer, 2015-2018
United Way Campaign Committee, College of Education, 2006
Arapahoe County, Election Judge for the County Clerk and Recorder, 2000-2005
Arthritis Foundation, Rocky Mountain Chapter, Certified Educator & Trainer, 1999-2008
Bonfils Blood Center, Silver Level Donor, 2000-2016
Denver Museum of Natural History, Docent, 1992
Tech Prep of the Lower Rio Grande Valley, Inc. (Harlingen, Texas) Board of Directors, 1995-1997 Executive Committee, 1996-1997; Chair, Development Committee, 1995-1997; Chair, Fiscal Agency Committee, 1995-1996; Chair, Colleges and Universities Committee, 1996-1997
Gilpin County (Colorado) Public Library Board of Trustees, 1986-1989; Vice President, 1987-1989
City of Central (Colorado) Economic Development Committee, 1987-1989
Columbine Family Health Centers, Inc. (Nederland and Black Hawk, Colorado) Board of Directors, 1988-1989

CERTIFICATION
Permanent Public Librarian Certificate - Pennsylvania, New York, Texas
Westlaw Expert Witness, 2008-
Certified Grant Writer®, 2016-

AWARDS AND HONORS
Advanced Practitioner for Service Learning and Community Engagement, University of Denver, 2011
Platinum Star Alumnae, College of Information, Library Science & Technologies, University of North Texas, 2009
Commendation for Integration of Technologies in Teaching & Learning Environment, University of Denver, 2006
Outstanding Adjunct Faculty Member Award, College of Education, University of Denver, 2002
Beta Phi Mu, Pi, University of Pittsburgh, 1985
Alpha Lambda Sigma, University of North Texas, 1972

Albert Nelson Marquis Lifetime Achievement Award, Who's Who in America
Who's Who in American Women
Who's Who of Women Executives
Dictionary of International Biography
Who's Who in the East
Who's Who in the South and Southwest
Who's Who in the World
Who's Who of Online Professionals
Who's Who in Library and Information Science
2,000 Notable Women
Who's Who of Emerging Leaders
Who's Who in Professional and Executive Women
Who's Who in American Education
International Who's Who of Professional and Business Women
International Leaders in Achievement
International Educator of the Year
Who's Who in Finance and Industry
Who's Who in Finance and Business

Invited International and National Conference Presentations

Driscoll, Margaret, Christy Confetti Higgins, Scott Brown, and Sylvia D. Hall-Ellis. *A View from Within: Open House Tour of Three Canvas Core Courses.* Presentation to be delivered at the School of Information, San José State University Professional Development Seminar, San José, California, October 17, 2019.

Hall-Ellis, Sylvia D. *Invest in Me -- I'm Your Future: Succession Planning for Libraries.* Keynote presentation delivered at the ALCTS President's Symposium, Boston, Mass., January 8, 2016.

Seidel, Kent E. and Sylvia D. Hall-Ellis. *Making Grants Work for You: Strategies for Doctoral Students and Early Career Scholars.* Presentation delivered at the University Council for Educational Administration, Early Career Scholars Session, Indianapolis, Ind., November 9, 2013.

Hall-Ellis, Sylvia D. *So You Want to be a Manager; Leadership Skills and Competencies for Technical Services Managers and Administrators.* Presentation delivered at the 138th Annual Conference, American Library Association, Chicago, Ill., June 29, 2013.

Seidel, Kent E., Karen S. Riley, Lyndsay Agans, Susan Korach, and Sylvia D. Hall-Ellis. *Making Grants Work for You (Instead of Just Working for Grants).* A panel discussion delivered at the University Council for Educational Administration, Early Career Scholars Session, Denver, Colo., November 17, 2012.

Hall-Ellis, Sylvia D. *After the Great TS Reorganization: The Westminster Law Library.* Presentation delivered at the 137th Annual Conference, American Library Association, Anaheim, Calif., June 23, 2012.

Hall-Ellis, Sylvia D. *Conversations with Catalogers in the 21st Century.* A panel discussion sponsored by the ALCTS Competencies for a Career in Cataloging Interest Group, delivered at the 137th Annual Conference, American Library Association, Anaheim, Calif., June 22, 2012.

Hall-Ellis, Sylvia D., moderator. *Mid-Career Leaders Program.* A panel discussion sponsored by the ALA Committee on Diversity delivered at the 136th Annual Conference, American Library Association, New Orleans, La., June 26, 2011.

Hall-Ellis, Sylvia D., moderator. *Diversity Town Hall.* A community conversation sponsored by the ALA Committee on Diversity delivered at the 136th Annual Conference, American Library Association, New Orleans, La., June 24, 2011.

LaBarre, Kathryn, Sylvia D. Hall-Ellis, Karen Anderson, Rick Hasenyager, Christopher Cronin, and Penny Baker. *Briefings from RDA Test Participants.* A panel discussion delivered at the Midwinter Conference, American Library Association, San Diego, Calif., January 7, 2011.

Miksa, Shawne, Marjorie Bloss, and Sylvia D. Hall-Ellis. *Educating the Next Generation of Catalogers: Teaching RDA.* A panel discussion delivered at the 97th Annual Conference, Association for Library and Information Science Education, San Diego, Calif., January 7, 2011.

Hall-Ellis, Sylvia D., Robert Maxwell, John Hostage, and George Prager. *RDA Panel: What Cataloging Managers Need to Know.* Presentation delivered at the 103rd Annual Conference, American Association of Law Librarians, Denver, Colo., July 12, 2010.

Hall-Ellis, Sylvia D. and Stacey L. Bowers. *Catalogers in the RDA Environment: Skill Sets, Expectations and Challenges.* Presentation delivered at the 103rd Annual Conference, American Association of Law Librarians, Denver, Colo., July 11, 2010.

Hall-Ellis, Sylvia D. *Comfortable in Your Cataloging and Metadata Specialist Skin? Or, So You Want to Hire a Cataloger.* Presentation delivered to the ALCTS Research Group at the 134th Annual Conference, American Library Association, Chicago, Ill., July 11, 2009.

Perez, Megan, Sylvia D. Hall-Ellis, and Denise Anthony. *From Novice to Expert: Collaboration for Succession Planning.* A "hot topic" presentation delivered at the 14th ACRL Conference, Seattle, Wash., March 13, 2009.

Hall-Ellis, Sylvia D. *Cataloging in the RDA Environment: Skill Sets, Expectations and Challenges.* Presentation delivered to the ALCTS Research and Publications Committee at the Midwinter Conference, American Library Association, Denver, Colo., January 24, 2009.

Hall-Ellis, Sylvia D. *LIS Cataloging Education for the 21st Century: Expectations and Challenges.* A panel discussion held at the 95th Annual Conference, Association for Library and Information Science Education, Denver, Colo., January 23, 2009.

Chu, Clara, Sylvia D. Hall-Ellis, and Mark Winston. *The Doctoral Degree & Building a Career.* A panel discussion delivered at the ALA Office of Diversity Spectrum Doctoral Fellows E.J. Josey Leadership Institute, Midwinter Conference, American Library Association, Denver, Colo., January 20, 2009.

Hall-Ellis, Sylvia D. and Robert O. Ellett, Jr. *Fundamentals of Cataloging Course: An Overview of the ALCTS Online Course.* Presentation delivered to the ALCTS Big Heads Group, 133rd Annual Conference, American Library Association, Anaheim, Calif., June 30, 2008.

Hall-Ellis, Sylvia D. *Employers' Expectations for Technical Services Librarians: What We Don't Know.* Presentation delivered to the ALCTS Research and Publications Committee Program, 133rd Annual Conference, American Library Association, Anaheim, Calif., June 28, 2008.

Hall-Ellis, Sylvia D., Virginia R. Maloney, and Mary Stansbury. *Institutional Responses to Engaged Scholarship: The Carnegie Foundation Engaged University Classification at Two Universities.* Presentation delivered to the 94th Annual Conference, Association for Library and Information Science Education, Philadelphia, Pa., January 11, 2008.

Hall-Ellis, Sylvia D. *Puzzles, Problems, and Predicaments.* Presentation delivered to the ALCTS Research Discussion Group, 132nd Annual Conference, American Library Association, Washington, D.C., June 23, 2007.

Hall-Ellis, Sylvia D. *Cataloging Education: A New Emphasis for the Library and Information Science Curriculum.* Presentation delivered to the ALCTS Pre-conference, 132nd Annual Conference, American Library Association, Washington, D.C., June 22, 2007. http://www.loc.gov/catdir/cpso/careercat.html

Ellett, Jr., Robert O. and Sylvia D. Hall-Ellis. *Copy Cataloging Done Smarter.* Presentation delivered to the International Conference on Interdisciplinary Information Sciences and Technologies (InSciT2006), October 25-28, 2006.

Hall-Ellis, Sylvia D. and Robert O. Ellett, Jr. *Cooperative Cataloging: Challenges and Opportunities for Defense Libraries.* Presentation delivered to the 1st Annual Conference of Defense Libraries, Spanish Ministry of Defense, Madrid, Spain, July 7, 2006.

Hall-Ellis, Sylvia D. *Cataloger Competencies: Do the Employers Require What the Professors Teach?* Presentation delivered to the ALCTS CCS Heads of Cataloging Discussion Group, 131st Annual Conference, American Library Association, New Orleans, La., June 26, 2006.

Grealy, Deborah S. and Sylvia D. Hall-Ellis. *From Research to Practice: The Scholarship of Teaching and Learning in LIS Education.* Presentation delivered to the at the 92nd Annual Conference, Association for Library and Information Science Education, San Antonio, Tex., January 18, 2006.

Hall-Ellis, Sylvia D. *Employers' Expectations for Entry-Level Catalogers: What Position Announcement Data Indicate.* Research paper delivered to the Technical Services Special Interests Group, 91st Annual Conference, Association for Library and Information Science Education, Boston, Mass., January 12, 2005. http://dlist.sir.arizona.edu/

Hall-Ellis, Sylvia D. *Common Errors in MARC Records Prepared by LIS Students: What Does It Mean?* Research paper delivered to the ALCTS CCS Cataloging Norms Discussion Group, Mid-Winter Conference, 90th American Library Association, San Diego, Calif., January 10, 2004.

Hall-Ellis, Sylvia D. *Visual Arts as Foundation for Successful Library Automation: The Rocky Mountain College of Art & Design Experience.* Paper delivered at the 6th Annual Conference of Higher Education, Charleston, S. C., 1993.

Hall, Sylvia D. *Design Elements for Bibliographic Databases: An Overview.* Paper delivered at the 14th Online National Conference, New York, 1983.

Invited Regional Conference Presentations

Hall-Ellis, Sylvia D., Hudson, Christopher D., Brittany Cronin, and Kathryn Michaels. "The Colorado Law Project: Meeting the Public's Need for Legal Information." Panel discussion delivered at the Mountain Plains Library Association Conference, Billings, Mont., April 9, 2011.

Grealy, Deborah S. and Sylvia D. Hall-Ellis. *Education for Information Professionals in New Mexico: Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the New Mexico Library Education Summit, Las Vegas, N.M., September 26, 2005.

Hall-Ellis, Sylvia D. *Public Library-School Library Partnerships.* Presentation delivered at the 3rd Annual Colorado Association of Libraries Conference with the Mountain Plains Library Association, Denver, Colo., October 22, 2004.

Hall-Ellis, Sylvia D. *Learn All You Can – Educational Partnership Opportunities for the Lewis and Clark Bicentennial Commemoration.* Paper delivered at the 3rd Annual Lewis and Clark Bicentennial Council National Planning Conference, Bismarck, N.D., April 26, 1998.

Invited State Conference Presentations

Grealy, Deborah S. and Sylvia D. Hall-Ellis. *Academic Library Leadership Changes: Using Succession Planning and Mentoring*. Presentation delivered at the Minnesota Library Education Conference, St. Cloud, Minn., October 10, 2013.

Hall-Ellis, Sylvia D., Merrie Valliant, and Melissa Powell. *RDA: What Is It and What Do You Need To Do With It At Your Library?* Presentation delivered at the Colorado Library Consortium Spring Conference, Ft. Morgan, Colo., April 26, 2013.

Hall-Ellis, Sylvia D. *Service Learning and the Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 4th LEADS Scholars Orientation, Denver, Colo., August 5, 2009.

Hall-Ellis, Sylvia D. *Service Learning: Enhancement to Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 3rd LEADS Scholars Orientation, Denver, Colo., June 2008.

Hall-Ellis, Sylvia D. *Law Librarianship: A Community Conversation.* Sponsored by the Colorado Association of Law Libraries. Paper delivered at the Colorado Supreme Court Library, Denver, Colo., May 14, 2008.

Hall-Ellis, Sylvia D. *Opportunities and Challenges in Law Librarianship: A Community Conversation.* Presentation delivered at the Sturm College of Law, Denver, Colo., November 7, 2007.

Hall-Ellis, Sylvia D. *Grant Writing Resources for Nursing Professionals.* Presentation delivered at the Presbyterian / St. Luke's Health One Medical Center 1st Annual Research Symposium, Denver, Colo., October 17, 2007.

Hall-Ellis, Sylvia D. *Project Homeless Connect 4 Event Evaluation – Insights and Lessons Learned.* Presentation delivered at the Homelessness Research Symposium: What is DU Doing about Homelessness in Denver, Denver, Colo., September 14, 2007.

Hall-Ellis, Sylvia D. *Service Learning and the Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 2nd LEADS Scholars Orientation, Denver, Colo., August 10, 2007.

Hall-Ellis, Sylvia D. *Education for Information Professionals in a Digital Environment: Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 15th Spring Mountains and Plains Parapros Conference, Denver, Colo., February 24, 2007.

Hall-Ellis, Sylvia D. *Public Library Service to Spanish-Speaking and Latino Residents in Denver: A Case Study.* Presentation delivered at the 4th Annual Colorado Association of Libraries Conference, Denver, Colo., November 10, 2005.

Hall-Ellis, Sylvia D. *Education for Information Professionals in a Digital Environment: Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 14th Annual Mountains and Plains Parapros Conference, Aurora, Colo., July 29, 2005.

Hall-Ellis, Sylvia D. *Educational Opportunities: Library & Information Science Graduate Education at the University of Denver.* Presentation delivered at the 13th Annual Mountains and Plains Parapros Conference, Centennial, Colo., August 6, 2004.

Hall-Ellis, Sylvia D. *Library Education & Training: Focus on the West: An LIS Faculty Member's Personal Response*. Presentation delivered at the 1st Annual Colorado Association of Libraries Conference, Keystone, Colo., October 18, 2002.

Hall-Ellis, Sylvia D. *Grant Writing for School Librarians.* Presentation delivered at the 2002 Annual Colorado Education Media Association Conference, Colorado Springs, Colo., February 15, 2002.

Hall-Ellis, Sylvia D. *Grants – Opportunities for the Future.* Paper delivered at the Southeast Regional Accountability Annual Conference, Lamar, Colo., November 12, 1998.

Hall-Ellis, Sylvia D. *The Texas Library Connection and Interlibrary Loan: An Experiment in Resource Sharing.* Paper delivered at the Texas Computer Educators' Association Annual Conference, Austin, Tex., February 6, 1997.

Hall-Ellis, Sylvia D. *Finding Grant Sources on the Internet: A Guide for Librarians.* Paper delivered at the 2nd Annual Institute for School Library Personnel, Pharr-San Juan-Alamo North High School, Pharr, Tex., July 29, 1996.

Hall-Ellis, Sylvia D. *Mathematical and Logical Thinking: A Critical Intelligence.* Paper delivered at the 3rd Annual Paraprofessional Conference at the University of Texas - Pan American, Edinburg, Tex., March 8, 1996.

Hall-Ellis, Sylvia D. *Cataloging Trends and Issues: Update Session.* Paper delivered at the 1st Annual Institute for School Library Personnel, South Texas Community College, McAllen, Tex., July 19, 1995.

Hall-Ellis, Sylvia D. *Grant Writing: Tips and Encouragement for School Librarians.* Paper delivered at the 1st Annual Institute for School Library Personnel, South Texas Community College, McAllen, Tex., July 18 and 19, 1995.

Hall-Ellis, Sylvia D. *How to Become an Expert Grant Writer.* Paper delivered at the 3rd Annual High School Principals' Academy, South Padre Island, Port Isabel, Tex., June 19, 1995.

Hall-Ellis, Sylvia D. *Texas Library Study: Results from Regions I and II.* Paper delivered at the 3rd Annual Technology Conference, Texas A&M University, College Station, Tex., November 18, 1994.

Hall-Ellis, Sylvia D. *Academic Achievement and Middle School Students.* Paper delivered at the 8th Annual Young Adult Conference, Sam Houston State University, Huntsville, Tex., November 5, 1994.

Hall-Ellis, Sylvia D. *Multimedia Resources for Library Leaders.* Paper delivered at the Institute for Librarians in A Multicultural Environment, Sam Houston State University, Huntsville, Tex., June 10, 1994.

Hall-Ellis, Sylvia D. *Finding the Resource: Empowering the User, or, the Case for Curriculum Based Subject Access to Learning Resource Center Collections.* Paper delivered at the 81st Annual Texas Library Association Conference, Corpus Christi, Tex., April 12-16, 1994.

Hall-Ellis, Sylvia D. and William H. Pichette. *Sam Houston State University Makes Use of OCLC/AMIGOS Collection Analysis CD.* Paper delivered at the 81st Annual Texas Library Association Conference, Corpus Christi, Tex., April 12-16, 1994.

Hall, Sylvia D. *Funding and Library Development in Pennsylvania: A Symbiotic Relationship.* Paper delivered at the Annual Graduate Student Colloquia, University of Pittsburgh, School of Library and Information Science, 1982.

Hall, Sylvia D. *Leadership for Public Library Trustees.* Paper delivered for the Trustees Division, Pennsylvania Library Association Annual Conference, Lancaster, Penn., 1981.

Seminar and Professional Development Presentations

Taylor, Meredith, and Sylvia D. Hall-Ellis. *Talent Management and Succession Planning.* ALCTS eForum, held March 22, 2017.

Hirsh, Sandra, Heather O'Brien, Michelle Holschuh Simmons, Michael Krasulski, and Sylvia D. Hall-Ellis. *Information Services Today: An Introduction. Part 3: Information Services: Roles in the Digital Age*. Rowan and Littlefield in partnership with Library Journal webinar, recorded February 5, 2015.

Hall-Ellis, Sylvia D. and Jennifer Sweda. *Copy Cataloging in an RDA Environment.* ALCTS eForum, held May 14 and 15, 2013.

Hall-Ellis, Sylvia D. *Law Librarianship: A Community Conversation.* Sponsored by the Colorado Association of Law Libraries, presented at the Colorado Supreme Court Library, Denver, Colo., May 14, 2008.

Hall-Ellis, Sylvia D. and Beatrice Z. Gerrish. *Reading and Libraries: Recent Research in Reading.* Presentation at the Ricks Center for Gifted Education, Denver, Colo., March 4, 2008.

Hall-Ellis, Sylvia D. *Cataloger Competencies: Do the Employers Require What the Professors Teach?* Presentation for the School of Library and Information Science, San José State University, February 12, 2008.

Hall-Ellis, Sylvia D. and Robert O. Ellett, Jr. *Cooperative Cataloging: Rules, Tools, and Conventions for Building a Multi-institutional Catalog.* Sponsored for the Spanish Ministry of Defense, Madrid, Spain, July 10, 2006.

Hall-Ellis, Sylvia D. *Cash for Kids: Grant Writing Opportunities for Youth Services Librarians.* Sponsored by the Colorado Young Adult Librarians; presented at Bemis Memorial Library, Littleton, Colo., November 13, 2002.

Hall-Ellis, Sylvia D. *MARC Records and Authority Control: Planning for Bibliographic Database Migration.* Sponsored and held at the Douglas County Public Library, Castle Rock, Colo., May 28, 2002.

Hall-Ellis, Sylvia D. *Grant Writing: A Refresher for Librarians.* Sponsored by Library and Information Science Program, College of Education, University of Denver; presented at University Center at Chaparral, August 12, 2000.

Hall-Ellis, Sylvia D. *Shaking the Money Tree – Basic Grant Writing for Colorado Educators.* Sponsored by the Office of Educational Telecommunications of the Colorado Department of Education.
Pikes Peak Community College, Colorado Springs, Colo., November 16, 1998.
Pueblo School District 60, Pueblo, Colo., November 12, 1998.
University of Northern Colorado, Greeley, Colo., November 10, 1998.
United Technology Educational Partnership, Grand Junction, Colo., November 9, 1998.

Hall-Ellis, Sylvia D. *Cataloging Multimedia, Kits, Globes and Map Materials in USMARC.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., April 24, 1997 and December 12, 1996.

Hall-Ellis, Sylvia D. *New Standards for School Library Media Centers.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., April 15, 1997.

Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future.* Sponsored and hosted by Pharr-San Juan-Alamo Independent School District, Pharr, Tex., April 4, 1997.

Hall-Ellis, Sylvia D. *Cataloging Sound Recordings and Audio Materials in USMARC.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., March 20, 1997 and November 21, 1996.

Hall-Ellis, Sylvia D. *Evaluating and Selecting CD-ROMS for School Library Media Collections.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., March 18, 1997.

Hall-Ellis, Sylvia D. *Cataloging Audiovisual Materials in USMARC.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., February 27, 1997, November 7, 1996 and October 24, 1996.

Hall-Ellis, Sylvia D. *Introduction to Dialog: Basic Searching Strategies.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., February 14, 1997.

Hall-Ellis, Sylvia D. *Cataloging Books in USMARC.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., January 30, 1997, October 10, 1996, September 26, 1996 and August 8, 1996.

Hall-Ellis, Sylvia D. *Advanced Internet Searching Techniques for Librarians.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., January 21, 1997.

Hall-Ellis, Sylvia D. *Texas Library Connection Full-Text Searching.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., December 13, 1996 and October 25, 1996.

Hall-Ellis, Sylvia D. *Developing Evaluation Strategies for Grants and Proposals.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., December 10, 1996.

Hall-Ellis, Sylvia D. *Developing Needs Assessment for Grants and Proposals.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., November 12, 1996.

Hall-Ellis, Sylvia D. *Texas Library Connection Union Catalog Searching.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., November 8, 1996 and September 6, 1996.

Hall-Ellis, Sylvia D. *Grant Resources on the Internet.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., October 23, 1996.

Hall-Ellis, Sylvia D. *Preparing a Response to the Telecommunications Infrastructure Fund Board: Needs Assessment, Professional Development Framework, and Evaluation Strategies.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., October 16, 1996.

Hall-Ellis, Sylvia D. *Texas Library Connection: Building the Districtwide Bibliographic Database.* Sponsored and hosted by Mercedes Independent School District, Mercedes, Tex., September 13, 1996 and August 8, 1996. Sponsored and hosted by Los Fresnos Consolidated Independent School District, Los Fresnos, Tex., August 12, 1996.

Hall-Ellis, Sylvia D. *School-to-Work and Special Education: An Inclusive Partnership for Success.* Sponsored and hosted by the Office of Special Education, Region One Education Service Center, Edinburg, Tex., September 11, 1996.

Hall-Ellis, Sylvia D. *Telecommunications Infrastructure Fund Board: An Overview of Funding for Secondary Schools.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., August 31, 1996.

Hall-Ellis, Sylvia D. *Cataloging Books, Multimedia, and Realia in USMARC.* Sponsored and hosted by Edinburg Consolidated Independent School District, Edinburg, Tex., August 6, 1996.

Hall-Ellis, Sylvia D. *Internet Resources for Grant Writers.* Sponsored and hosted by Region One Education Service Center, Edinburg, Tex., April 20, 1996.

Hall-Ellis, Sylvia D. *Enhanced Grant Writing Skills for Mathematics Educators: Writing Skills for Campus Teams.* Sponsored and hosted by the Office of General Education, Region One Education Service Center, Edinburg, Tex., March 22, 1996.

Hall-Ellis, Sylvia D. *Internet Resources for Grant Writers.* Sponsored and hosted by the College of Education and Applied Science, Sam Houston State University, Huntsville, Tex., March 9, 1996.

Hall-Ellis, Sylvia D. *Grant Writing for Mathematics Educators: A Development Process for Campus Teams.* Sponsored and hosted by the Office of General Education, Region One Education Service Center, Edinburg, Tex., March 4, 1996.

Hall-Ellis, Sylvia D. *Shaking the Money Tree: Preparing Successful Technology Grant Applications.* Sponsored and hosted by the Office of Technology and Media Services, Region One Education Service Center, Edinburg, Tex., February 29, 1996 and April 29, 1996.

Hall-Ellis, Sylvia D. *District-wide Technology Planning: Technical Assistance for the Texas Education Agency Initiative.* Sponsored and hosted by the Office of Technology and Media, Region One Education Service Center, Edinburg, Tex., February 2, 1996.

Hall-Ellis, Sylvia D. *Enhanced Grant Writing Skills.* Sponsored and hosted by Pharr-San Juan-Alamo Independent School District, Pharr, Tex., January 27, 1996.

Hall-Ellis, Sylvia D. *United States Copyright Act of 1976, Video Transmissions, Computer Software and the Internet.* Sponsored and hosted by the Office of Technology and Media, Region One Education Service Center, Edinburg, Tex., January 24, 1996.

Hall-Ellis, Sylvia D. *The Grant Writing Development Process.* Sponsored and hosted by Pharr-San Juan-Alamo Independent School District, Pharr, Tex., January 20, 1996.

Hall-Ellis, Sylvia D. *MARC Cataloging of Materials for Library Media Centers.* Sponsored and hosted by the Office of Media Services, Region One Education Service Center for the Donna Independent School District, Donna, Tex., August 17, 1995.

Hall-Ellis, Sylvia D. *Classification.* Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., May 10, 1995.

Hall-Ellis, Sylvia D. *Library Technical Services.* Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., April 26, 1995.

Hall-Ellis, Sylvia D. *Grant Writing: An Introduction for Public School Administrators.* Sponsored and hosted by the Office of Administrative Services, Region One Education Service Center, Edinburg, Tex., April 19, 1995.

Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future.* Sponsored and hosted by the United Independent School District, Laredo, Tex., March 31, 1995.

Hall-Ellis, Sylvia D. *Using USMARC.* Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., March 29, 1995.

Hall-Ellis, Sylvia D. *Library Media Center Policies and Guidelines: How to Prepare for School Board Adoption.* Sponsored and hosted by the Office of Technology and Media Services, La Joya Independent School District, La Joya, Tex., March 10, 1995.

Hall-Ellis, Sylvia D. *Developing District-wide Policies and Guidelines for Library Media Centers.* Sponsored and hosted by the Office of Technology and Media Services, La Joya Independent School District, La Joya, Tex., January 10, 1995.

Hall-Ellis, Sylvia D. *MARC Cataloging of Audiovisual Materials for Library Media Centers.* Sponsored and hosted by the Office of Library Media and Technology Services, Cypress-Fairbanks Independent School District, Houston, Tex., December 1, 1994.

Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future.* Sponsored and hosted by United Independent School District, Laredo, Tex., October 14, 1994.

Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future.* Sponsored and hosted by the Laredo Independent School District, Laredo, Tex., October 7, 1994.

Hall-Ellis, Sylvia D. *MARC Cataloging for Library Media Centers.* Sponsored by the Office of Library Media and Technology Services, Cypress-Fairbanks Independent School District, Houston, Tex., October 5, 1994.

Hall-Ellis, Sylvia D. *Jump to the Head of the Class: Undergraduate Library Resources Available at Sam Houston State University.* Sponsored by the Office of the Associate Vice President for Student Services, Sam Houston State University, Huntsville, Tex., October 4, 1994.

Hall-Ellis, Sylvia D. *CD-ROMs - 1994's Newest and the Best for Secondary Level Media Centers.* Sponsored and hosted by Clear Lake Independent School District, Houston, Tex., June 3, 1994.

Hall-Ellis, Sylvia D. *Introduction to Classification.* Sponsored by AJ Seminars, Rockville, Maryland; presented at Holiday Inn, Market Center, Dallas, Tex., May 18, 1994.

Hall-Ellis, Sylvia D. *Basic Descriptive Cataloging.* Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., May 4, 1994.

Hall-Ellis, Sylvia D. *Shaking the Money Tree - Part II: Writing Successful Grant Applications.* Sponsored by Donna Independent School District, Donna, Texas, and Region One Education Service Center Edinburg, Texas; presented at South Texas Community College Library, McAllen, Tex., April 29, 1994.

Hall-Ellis, Sylvia D. *Using MARC.* Sponsored by AJ Seminars, Rockville, Maryland; presented at Holiday Inn, Market Center, Dallas, Tex., April 6, 1994.

Hall-Ellis, Sylvia D. *Automated Authority Control.* Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., March 23, 1994.

Hall-Ellis, Sylvia D. *Automating the District School Library Media Centers: Choices and Opportunities.* Sponsored and hosted by the Office of Technology and Library Media Services, Fort Bend Independent School District, Sugar Land, Tex., March 4, 1994.

Hall-Ellis, Sylvia D. *Shaking the Money Tree - Part I: Preparing Successful Grant Applications.* Sponsored and hosted by the Office of Library Media Services, Donna Independent School District, Donna, Tex., February 4, 1994.

GRADUATE COURSES TAUGHT

San José State University, School of Library and Information Science
INFO 249 – Advanced Cataloging and Classification (Fall 2015, 2016; Summer 2016)
INFO 287 – Special Topics in Cataloging and Classification (Spring 2017, 2018, 2019, 2020; Fall 2020)
LIBR 248 – Beginning Cataloging and Classification (Summer 2002, 2003, 2004, 2005, 2006)
LIBR 249 – Advanced Cataloging and Classification (Summer 2003; Fall 2014; Summer 2015)

University of Denver, University College
Grant Writing Certificate Seminar for Non-Profit Organizations (Fall 2020)

University of Denver, Morgridge College of Education
HED 5991 – Grant Writing in Higher Education (Spring 2011)
LIS 4010 – Organization of Information (Fall 2002, 2003, 2004, 2005, 2006, 2009; Winter 2005; Spring 2004, 2005, 2006)
LIS 4020 – Professional Principles and Ethics (Summer 2000)
LIS 4040 – Management of Libraries and Information Centers (Fall 2010; Spring 2003, 2009 (DS); Winter 2005)
LIS 4070 – Cataloging and Classification (Winter 2008, 2009, 2010, 2011, 2012; Fall 2009 (DS))
LIS 4321 – Collection Management (Spring 2005)
LIS 4326 – LIS Research (Winter 2009 (DS); Spring 2009 (DS))
LIS 4350 – Adult Materials and Services (Summer 2006, 2009)
LIS 4379 – Social Sciences Resources (Spring 2009)
LIS 4400 – Cataloging and Classification (Spring 2000, 2001; Winter 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009; Summer 2005, 2006 – course renumbered LIS 4070, September 2007)
LIS 4401 – Descriptive Cataloging (Winter 2001; Spring 2002, 2003, 2006, 2007, 2008, 2009; Summer 2005)
LIS 4402 – Subject Cataloging (Spring 2001, 2009; Summer 2005, 2006, 2007, 2008)
LIS 4403 – Classification Schemes (Fall 2007, 2008, 2009 (DS))
LIS 4405 – Authority Control (Winter 2009)
LIS 4510 – Materials and Services for Children (Winter 2004; Summer 2005)
LIS 4620 – Grant Writing and Fundraising (Summer 2000, 2002, 2004, 2010; Winter 2006; Fall 2006, 2007, 2008, 2009)
LIS 4700 – Seminar in Technical Services (Fall 2001)
LIS 4700 – Seminar in Public Libraries (Summer 2002)
LIS 4804 – Management of Electronic Records (Spring 2004; Fall 2005)
LIS 4902 – Capstone Projects (Winter 2003, Spring 2008, 2009, 2011)
LIS 4910 – Independent Study (every quarter Winter 2000 through Spring 2011)
LIS 4920 – Service Learning (every quarter Summer 2004 through Spring 2011)
RMS 4954 – Grant Writing (Summer 2013, 2014)
RMS 4959 – Content Analysis Methodology (Spring 2015)
QRM 4978 – Grant Writing (Summer 2011, 2012)

Rutgers University, School of Communication and Information
SC&I 522 – Cataloging and Classification (Summer 2013)

University of Arizona, College of Behavioral and Social Sciences
LIS 602 – Cataloging and Classification (Summer 1995)
LIS 612 – Advanced Online Search and Retrieval (Summer 1995)

Sam Houston State University, College of Education & Applied Science
LIS 532 - Cataloging and Classification (Fall 1993, 1994, 1995)
LIS 563 - Advanced Cataloging and Classification (Summer 1994)
LIS 567 - Research Methods (Spring 1994, Spring 1995)
LIS 591 - Educational Technology (Spring 1994, 1995)
LIS 596 - Networking and Computer Technologies in Education (Fall 1993, Summer 1994, Fall 1995)

Dissertations at the University of Denver

Bowers, Stacey L. *Library Anxiety of Law School Students: A Study Utilizing the Multidimensional Library Anxiety Scale.* Chair, May 5, 2010.

Fattor, Melissa M. *Student Engagement Differences by Ethnicity and Scale for Ninth Grade Students.* Chair, November 1, 2010.

Fulton, Roseanne. *A Case Study of Culturally Responsive Teaching in Middle School Mathematics.* Kent Seidel, Chair, June 18, 2009, Outside Chair.

Grealy, Deborah S. *Tribes and Territories in Library and Information Studies Education.* Bruce Uhrmacher, Chair, June 10, 2008, Committee Member.

McCord, J. Michael. *Developing a Standard of Care for Educational Malpractice.* Chair, April 15, 2011.

Priebe, Sarah J. *Distinguishing Effects of Domain and General Knowledge on Passage Fluency and Comprehension.* Jan Keenan, Chair. July 21, 2011, Outside Chair.

Snyder-Mondragon, Sandra M. *Institutional Factors that Impact the Retention of Graduate Students of Color in Schools of Library and Information Science: A Metaregression of Accredited Library School Statistics on Student Retention and Graduation Rates.* Kathy Green, Chair, July 24, 2009, Committee Member.

Taylor, Karen Pickles. *Effective Teaching.* Elinor Katz, Chair, July 15, 2009, Outside Chair.

Thompson, Jennifer. *Distinguishing a Western Women's College: A History of the Curriculum and Student Experience at Colorado Women's College.* Edith W. King, Chair, July 16, 2010, Committee Member.

Walker, Emelda. *Influence of Organizational Factors on Job Satisfaction of Disability Service Providers at Postsecondary Institutions.* Chair, April 29, 2010.

Dissertation Proposals at the University of Denver

Bowers, Stacey L. *Library Anxiety of Law School Students: A Study Utilizing the Multidimensional Library Anxiety Scale.* Chair, November 5, 2009.

McCord, J. Michael. *Developing a Standard of Care for Educational Malpractice.* Edith W. King, Chair. April 23, 2010.

Thompson, Jennifer. *History of Colorado Women's College.* Edith W. King, Chair, October 30, 2008. Committee Member.

Walker, Emelda. *Influence of Organizational Factors on Job Satisfaction of Disability Service Providers at Postsecondary Institutions.* Chair, July 21, 2009.

Dissertations at Other Institutions

Rodríguez-Mori, Howard. *The Information Behavior of Puerto Rican Immigrants to Central Florida, 2003-2009: Grounded Analysis of Six Case Studies Use of Social Networks during the Migration Process.* Kathleen Burnett, Chair, April 10, 2009, Florida State University, Committee Member.

Schwartz, Brian *More than a Look-up Skill: Medical Information Literacy Education in Osteopathic Medical Schools.* D. Mirah Dow, Chair, July 18, 2017. Emporia State University, Committee Member.

Snow, Karen. *A Study of the Perception of Cataloging Quality among Catalogers.* Shawne D. Miksa, Chair, August 1, 2011. University of North Texas, Committee Member.

Dissertation Proposals at Other Institutions

Schwartz, Brian *More than a Look-up Skill: Medical Information Literacy Education in Osteopathic Medical Schools.* D. Mirah Dow, Chair, December 4, 2015. Emporia State University, Committee Member.

Snow, Karen. *A Study of the Perception of Cataloging Quality among Catalogers.* Shawne D. Miksa, Chair, May 11, 2010. University of North Texas, Committee Member.

Master's Thesis at the University of Denver

Hemingson, Jeff. *Recital Paper.* Lamont School of Music, February 2010, Outside Chair.

Capstones at the University of Denver

Anthony, Alisa. *Correlation of Library and Information Science Program outcomes and Vacant Position Qualifications Listed on the Colorado State Library Jobline by Employers During the Period September 1, 2000 through August 31, 2002.* Chair, Winter Quarter, 2003.

Borden, Donna M. *Improving Emergency Communications Systems: Is a Radio Communications Network the Answer?* Fall Quarter, 2014.

Bowden, Heather L.M. *Exploring Biological Models for Long-term Data Preservation.* Chair, Spring Quarter, 2008.

Casenada, Cassandra Y. *Challenges to the Recruitment, Education, and Retention of Librarians of Color*. Chair, Spring Quarter, 2008.

Chang, Jennifer C. *Legal Research Practice and Preference: A Law Firm Perspective.* Chair, Spring Quarter 2011.

Ellis, Megan S. Fitzgerald. *Design of a Public Library Adult Volunteer Recruitment Program and Training Curriculum.* Chair, Winter Quarter, 2003.

Kircher, Kathy. *Development of a Library Pathfinder for Exobiology and Posting it on the Internet.* Chair, Winter Quarter, 2003.

Melhado, Loretta. *Design of St. John's Episcopal Church Library.* Chair, Spring Quarter, 2008.

Radcliff, Kathy. *Original Cataloging of Archival materials in the HERS Collection in Penrose Library.* Chair, Winter Quarter, 2003.

Sass, Carol Ann. *ACT Periodical Index: Access to Catholic Thinking Periodical Index: Web Index to Select Catholic Periodicals.* Spring Quarter 2000.

Stone, Sergio D. *Conducting Community Analysis for the Bemis Public Library (Littleton, Colorado) Using 2002 U.S. Census Data and the Online Outcome-Based Evaluation Toolkit.* Chair, Winter Quarter, 2003.

Tureson, Tamara. *Design and Use of an Information Audit Tool for Use in a Law Library.* Chair, Winter Quarter, 2003.

Tweed, Beth. *Evaluation of Email Reference Service in a Consumer Health Library Environment.* Chair, Winter Quarter, 2003.

GRANTS AND CONTRACTS

Student Learning in Agriculture STEM through Teacher Professional Development. Research team: Stanton Gartin (PI) and Cyndi Hofmeister, Northeast Junior College; Jeff Cash, Cheryl Sánchez, Anne-Marie Crampton, Lamar Community College; Suzanna Spears, Morgan Community College; Jack Wiley, Kerry Gabrielson, Trinidad State Junior College; Michael J. Miller, Colorado State University; Michael Womochil (Co-PI), Casey Sacks, and Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Agriculture, Agriculture and Food Research Initiative, Food, Agriculture, Natural Resources and Human Sciences Education and Literacy Initiative, $404,460 (2017-2020)

Leading and Achieving: the Colorado Agriculture Regional Consortium. Research team: Jack Wiley (PI), Kerry Gabrielson, Trinidad State Junior College; Jeff Cash (PI), Cheryl Sánchez, Anne-Marie Crampton, Lamar Community College; Suzanna Spears, Morgan Community College; Cyndi Hofmeister, Northeast Junior College; Michael Womochil (Co-PI), Casey Sacks, and Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Agriculture, Hispanic Serving Institutions Education Grants Program, $504,414 (2017-2020)

Cyber Prep Program Planning Grant. Research team: Debbie Sagen (PI), Brenda Lauer, Pikes Peak Community College; Gretchen Martin, Koiosa Insights; and, Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Commerce, Regional Alliances and Multi-stakeholder Partnerships to Stimulate (RAMPS) Cybersecurity Education and Workforce Development, $199,681 (2016-2018)

Career and Technical Education in Colorado: Pathways to Education and Employment. Research team: Heather McKay (PI), Rutgers University; Sarah Heath, Casey Sacks, Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Education, Institute of Education Sciences, $1,400,000 (2017-2021)

Pre-Alliance Planning Grant Colorado Community College Alliance. Research team: Victor Vialpondo (PI) and Janel Highfill, Community College of Aurora; Heidi Loshbaugh (Co-PI), Community College of Denver; Rick Reeves, Bill McGreevy, Liz Cox, and Kristin Aslin, Red Rocks Community College; Cathy Pellish, Front Range Community College; Samuel DeVries, Arapahoe Community College; Sylvia D. Hall-Ellis, Colorado Community College System. National Science Foundation, Louis Stokes Alliance for Minority Participation (LSAMP), $86,817 (2016-2017)

Towards Scalable Differentiated Instruction Using Technology-enabled, Competency-based, Dynamic Scaffolding. Research team: Karen Wilcox (PI) and Vijay Kumar, Center for Computational Engineering, Massachusetts Institute of Technology; Flora McMartin, Broad-based Knowledge; Quinsigamond Community College; and, Casey Sacks, and Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Education, First in the World Developmental Grant, $2,891,882 (2015-2019)

Colorado Strategic Partnerships Emergency Grant. Research team: Elise Lowe-Vaughn (PI), Amy Hodson, Celia Hardin, Barbara McBride, James Newby, Christopher Dewhurst, Nina Holland, Kate Anderson, MaryAnn Roe, Chrystalynn Chrystalynn, Elaine Edon, Mona Barnes, Tom Morgan, Rob Hanni, and Marie Valenzuela, Colorado Department of Labor and Employment; Steve Anton, Joelle Brouner, and, Katie Griego, Colorado Department of Human Services; Rebecca Holmes, Judith Martinez, and Jennifer Jirous, Colorado Department of Education; Emily Templin Lesh, Colorado Workforce Development Council; Cory Everett, Colorado Department of Regulatory Agencies; and, Casey Sacks and Sylvia D. Hall-Ellis, Colorado Community College System. U.S. Department of Labor, Strategic Partnerships Emergency Grant, $5,000,000 (2015-2017)

CAEL Jump Start Program: Competency-Based Education. Research team: Jerry Migler (PI), Casey Sacks, Debra Cohn, Thomas Hartman, and Sylvia D. Hall-Ellis, Colorado Community College System; Matt Jamison, Front Range Community College; Mike Coste, Red Rocks Community College; Amanda Corum, Pueblo Community College; Janet Colvin, Pikes Peak Community College; and, MaryAnn Matheny, Community College of Denver. Council for Adult and Experiential Learning, $15,000 (2015)

Summit on the Redesign of Developmental Education. Research team: Jerry Migler (PI), Casey Sacks, and Sylvia D. Hall-Ellis, Colorado Community College System; Chip Nava, Pueblo Community College; Kim Moultney,

Arapahoe Community College; Debbie Ulibarri, Trinidad State Junior College; and, Kris Bernard, Front Range Community College. American Association of Community Colleges, $15,000 (2015-2016)

Equity in Excellence at Colorado Community Colleges. Research team: Keith Howard (PI) and Sylvia D. Hall-Ellis, Colorado Community College System; Estela Mata Bensimon, Center for Urban Education, University of Southern California; and Kerry Gabrielson, Trinidad State Junior College. Colorado Department of Higher Education, Colorado Opportunity Scholarship Program, $150,000 (2015-2016)

MBA High School of Business in Colorado. Research team: Laurie Urich (PI) and Sylvia D. Hall-Ellis, Colorado Community College System; Rudolph Sumpter and Beatrice Gerrish, Boulder Valley Schools; Keith Curry Lance, RSL Research Group. Colorado Department of Higher Education, Colorado Opportunity Scholarship Program, $501,295 (2015-2016)

Fullbridge Program in Colorado. Research team: Keith Howard (PI) and Sylvia D. Hall-Ellis, Colorado Community College System; Suzanna Spears, Fort Morgan Community College; and, Cheryl Sánchez and Anne-Marie Compton, Lamar Community College. Colorado Department of Higher Education, Colorado Opportunity Scholarship Program, $300,000 (2015-2017)

Colorado Community College System Alternative Credit Program. Research team: Jerry Migler (PI), Keith Howard, and Sylvia D. Hall-Ellis, Colorado Community College System. American Council on Education, $13,000 (2015)

Internationalization in Higher Education. Researcher: Samuel D. Museus. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $8,000 (2014-2015)

International Perspectives on Bilingual Education. Researcher: Sharolyn Pollard-Durodola. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $3,700 (2014-2015)

Cultivating Culturally Relevant and Responsive Curriculum and Pedagogy in College. Researcher: Samuel Museus. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, iRise Grant, $5,000 (2014-2015)

Project EMERGE (Educational Model for Evaluation and Replicability in Gifted Environments). Research team: Norma Hafenstein (PI) and Bruce Uhrmacher. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Lynde and Harry Bradley Foundation, $235,000 (2014-2015)

Collecting Asian American Refugee Stories. Researcher: Samuel Museus. Technical reviewer for MCE: Sylvia D. Hall-Ellis. American Educational Research Association, Research Grant, $5,000 (2014-2015)

Developing Expertise in Teaching K-5 Mathematics. Research team: Julie Sarama (PI) and Douglas H. Clements (Co-PI), in partnership with the School of Education at the University of Michigan. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, $130,344 (2013-2015)

Investigation of the Long-Term Outcomes for Special Education Students. Researcher: Antonio Olmos-Gallos. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Jefferson County School District, Office of Assessment, $40,000 (2014)

Refugee Community Collaboration. Researcher: Vicki Tomlin (PI) in partnership with the African Community Center and Jewish Family Service. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, $14,854 (2014-2015)

Access to Mathematics for All. Research team: Richard S. Kitchen (PI), Nicole M. Russell (Co-PI), and Terrence Blackman (Co-PI), Curriculum Studies and Teaching Program, Morgridge College of Education; Álvaro Árias (Co-PI, Department of Mathematics); and, James Gray (Co-PI), Department of Mathematics, Community

College of Aurora. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, The Robert Noyce Scholarship Program, Capacity Building Project, $349,926 (2014-2016)

Cognitive Test Battery for Intellectual Disabilities. Research team: Karen Riley (PI), Lyndsay Agans, Jessica Lerner, and Karin Dittrick-Nathan in partnership with David Hessl (PI), The MIND Institute at the University of California – Davis, and Elizabeth Berry-Kravis, Rush University Medical Center. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Institutes of Health, Outcome Measures for Use in Treatment Trials for Individuals with Intellectual and Developmental Disabilities (R01), $2,499,996 ($588,672 at DU) (2014-2019)

Broadening Participation in Engineering among Women and Latino/as: A Longitudinal, Multi-Site Study. Researcher: Patton Garriott (PI) in partnership with the University of North Dakota and the University of Missouri. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, HER Core Research, $677,390 ($69,992 at DU) (2014-2019)

Designing a Teacher Evaluation System to Improve Teacher Effectiveness for Culturally and Linguistically Diverse Learners. Research team: María del Carmen Sálazar (PI), Jessica Lerner, and Kathy Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Professional Research Opportunities for Faculty, $29,988 (2014-2016)

Developing a College-Going Culture in Latina/O Families: Exploring the Influence of Funds of Knowledge on Family Outreach Programs. Researcher: Judy Marquez Kiyama. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Professional Research Opportunities for Faculty, $18,720 (2014-2016)

Assessment of Quality of Life in Neutral Implantation Surgery for the Treatment of Parkinson's Disease. Researcher: Cynthia McRae. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $3,967 (2014-2015)

Pura Vida: Cloud Forest, Curriculum and Cross-Cultural Study. Research team: Norma Hafenstein (PI) and Bruce Uhrmacher (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $8,000 (2014-2015)

Online Course Development for Curriculum and Instruction. Researcher: Ruth Chao. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $3,000 (2014)

Online Course Development for Curriculum and Instruction. Researcher: María del Carmen Sálazar. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $3,000 (2014)

Online Course Development for Curriculum and Instruction. Researcher: Jessica Lerner. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $3,000 (2014)

Online Course Development for Curriculum and Instruction. Researcher: Duan Zhang. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $3,000 (2014)

The Mathematics Education of African Americans, 1866 – 1954. Researcher: Nicole M. Russell. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Academy of Education, Spencer Foundation Postdoctoral Fellowship, $55,000 (2014-2016)

Early Childhood Care and Education Study. Research team: Carrie Germeroth (PI), Melissa Mincic, and Douglas H. Clements. Technical reviewer for MCE: Sylvia D. Hall-Ellis. State of North Dakota, Department of Public Instruction, $73,500 (2013-2014)

Graduate Level Specialty in Addiction Counselor Training with Emphasis on Integration of Native American Specific Content. Research team: Ruth Chao (PI) and Michael J. Faragher (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. The Galena Foundation, $289,732 (2013-2016)

The Collecting Asian American and Pacific Islander Refugee Stories (CARS) Project. Researcher: Samuel Museus. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, $15,000 (2014)

Parents in Transition: A Multiple Case Study of Parent and Family Orientation Programs. Researcher: Judy Marquéz Kiyama. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Faculty Research Fund, $1,500 (2014-2015)

The Sistah Network: Black Women Graduate Students Supporting and Retaining Each Other. Researcher; Nicole M. Russell. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Faculty Research Fund, $2,708 (2014-2015)

Evaluation of the Northeast Denver Babies Ready for College Program. Research team: Carrie Germeroth (PI), Melissa Mincic, and Douglas H. Clements. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Mile High Montessori, $9,922 (2013-2014)

Developing Teaching Expertise in K-5 Mathematics. Research team: Julie Sarama (PI) and Douglas Clements (University of Denver), in partnership with Timothy Boerst (PI), Meghan Shaughnessy, Deborah Ball, Hyman Bass (School of Education at the University of Michigan). Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, $449,827 ($130,344 at the University of Denver (2013-2015)

Early Learning Care and Education Study Program Grant for the State of North Dakota. Research team: Carrie Germeroth (PI), Melissa Mincic, and Sheridan Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. State of North Dakota Department of Public Instruction, $73,500 (2013-2014)

Local Professional Learning Community for School Leaders. Research team: Susan Korach (PI), Kristina Hesbol, and Rebecca McClure. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Education Development Center, $10,950 (2013)

Healthy Eaters, Lifelong Movers 2: Implementing Evidence-Based School Environment, Policy, and Curricular Changes to Increase Opportunities for Healthy Eating and Physical Activity in Low Income, Rural Colorado. Research team: Elaine Berlansky, University of Colorado at Denver (PI), Nicholas Cutforth, University of Denver (Co-PI), and Allison Reeds. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Health Foundation, $3,103,108 (2013-2017)

Hughes Rare Books Library Room Renovation in the Sturm College of Law. Project team: Patti H. Marks, Sylvia D. Hall-Ellis, and Leigh Elliott. Mabel T. Hughes Charitable Trust, $34,000 (2013-2014)

The Promise Center Partnership with the Marsico Institute for Early Learning and Literacy and the City and County of Denver. Research team: Karen Riley (PI), Douglas H. Clements (Co-PI), and Sheridan Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. The Piton Foundation, $223,468 (2013-2014)

Center of Excellence for Problem Gambling. Research team: Ruth Chao (PI) and J. Mike Faragher. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Department of Behavioral Health, $68,021 (2012-2013)

United Way Implementation and Validation Review. Research team: Douglas H. Clements (PI), Amanda Moreno, and Sheridan Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Mile High United Way, $19,737 (2013)

Math/Science Partnership in Rural Districts. Research team: Kristen Bunn (PI, Eagle County Schools), Paul Michalec (Co-PI), and Alegra Reiber. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Department of Education, Colorado's Mathematics and Science Partnership Program, $750,000 (2013-2014)

Mathematics and Science Education of African Americans. Research team: Nicole Russell (PI), Sylvia D. Hall-Ellis, Steve Fisher (Penrose Library). Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Professional Research Opportunities for Faculty, $29,994 (2013-2015)

Online Course Development for Curriculum and Instruction. Researcher: Nicole Russell. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $3,000 (2013)

Discourse and Opportunity: Undocumented Students and Higher Education Policy. Researcher: Ryan Gildersleeve. Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Academy of Education, Spencer Foundation Postdoctoral Fellowship, $55,000 (2012-2014)

An Anthropological Study of the Latino Graduation Ceremony. Researcher: Ryan Gildersleeve. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Spencer Foundation, $39,900 (2012-2013)

Fragile-X and Pharmaceutical Company: Clinical Trial of AFQ056. University of California – Davis Children's Hospital MIND Institute (Sacramento), Children's Hospital Denver, and Rush Children's Hospital at Rush University Medical Center (Chicago). Karen Riley, PI. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Novartis Pharmaceuticals Corporation, $394,206 (2012-2014)

Project Engage, Phase 2, a DAPRA Grant in partnership with Total Immersion Systems, Inc., and Texas A&M University. Research team: Karen Riley (PI) and Lyndsay Agans (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. U.S. Department of Defense, $60,000 (2011-2013)

International School Psychology Practicum Exchange. Researcher: Gloria L. Miller. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, $14,951 (2012-2013)

Online Course Development for Curriculum and Instruction. Research team: Bruce Uhrmacher (PI) and Norma Hafenstein (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Teaching and Learning, $18,193 (2012-2014)

Writers in the Schools. Research team: Karen Riley (PI) and Amanda Moreno (Co-PI).Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Humanities, $4,965 (2012-2013)

Learning Ecosystem Validation Grant. Research team: Karen Riley (PI), Lyndsay Agans, Kent Seidel, and Shimelis Assefa. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Bill & Melinda Gates Foundation, $315,000 (2012-2013)

Project Words of Oral Reading and Language Development (WORLD). Research team: Jorge E. Gonzalez (PI), Texas A&M University; Laura Saenz (Co-PI), University of Texas – Pan American; and, Sharolyn Pollard-Durodola (Co-PI), University of Denver. Technical reviewer for MCE: Sylvia D. Hall-Ellis. U.S. Department of Education, Institute of Education Sciences, $53,354 (2012-2015); award $640,718, transfer from Texas A&M University

Increasing the Efficacy of an Early Mathematics Curriculum with Scaffolding Designed to Promote Self-Regulation. Research team: Douglas H. Clements (PI) and Julie Sarama (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. U.S. Department of Education, Institute of Education Sciences, $1,445,315 (2008-2014); award $4,541,975, transfer from University at Buffalo, The State University of New York

Using Rule Space and Poset-based Adaptive Testing Methodologies to Identify Ability Patterns in Early Mathematics and Create a Comprehensive Mathematics Ability Test. Research team: Douglas H. Clements (PI) and Julie Sarama (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, $323,791 (2010-2014); award $1,194,944, transfer from University at Buffalo, The State University of New York

Comprehensive Postdoctoral Training in Scientific Education Research. Researcher: Julie Sarama (PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. U.S. Department of Education, Institute of Education Sciences, $133,458 (2010-2014); award $613,353, transfer from University at Buffalo, The State University of New York

Longitudinal Study of a Successful Scaling Up Project: Extending TRIAD. Research team: Douglas H. Clements (PI), Julie Sarama (Co-PI), and Abt Associates (Carolyn Layzer, Fatih Unlu, Laurie Bozzi, Lily Fesler, Alina Martinez, Cristofer Price, James van Orden). Technical reviewer for MCE: Sylvia D. Hall-Ellis. Institute of Education Sciences, $384,940 (2011-2015; award $1,250,286, transfer from University at Buffalo, The State University of New York

Early Childhood Education in the Context of Mathematics, Science, and Literacy. Research team: Julie Sarama (PI), Douglas H. Clements (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. National Science Foundation, $990,020 (2010-2014; award $2,285,228, transfer from University at Buffalo, The State University of New York

Early Childhood Clearinghouse Information Center Redesign. Research team: Karen Riley (PI) and Amanda Moreno (Co-PI).Technical reviewer for MCE: Sylvia D. Hall-Ellis. Office of the Lt. Governor of Colorado, $12,000 (2012-2013)

Morgridge Rural Educational Leadership Initiative. Research team: Lyndsay Agans (PI), Linda Brookhart (Co-PI), Susan Korach, and Rebecca McClure. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Morgridge Family Foundation, $100,000 (2012-2014)

Center of Excellence for Problem Gambling. Research team: Patrick Sherry (PI) and J. Mike Faragher. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Department of Behavioral Health, $55,000 (2012-2013)

Intermodal Transportation Institute Research Initiatives. Researcher: Patrick Sherry. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University Transportation Centers Program, Research and Innovative Technology Administration, U.S. Department of Transportation. $600,000 (2012-2014)

Mile High United Way Social Innovation Fund Early Literacy Initiative. Gloria L. Miller, PI, Amanda Moreano, Kim Hartnett-Edwards, and Sheridan Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Mile High United Way, $89,992 (2012-2013)

International School Psychology Practicum Exchange. Researcher: Gloria L. Miller. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, $1,000 (2012-2013)

Refugee Student Art Outreach. Researcher: Karin Dittrick-Nathan. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Faculty Research Fund, $3,000 (2012-2013)

Resistance, Resilience, and Reciprocity: Centering the Voices of Black Doctoral Women with Faculty Aspirations. Researcher: Nicole M. Russell. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Faculty Research Fund, $2,965 (2012-2013)

Assessing Learning through Student Notebooks. Research team: Keith Miller, Nancy Sasaki, and Kathy Green. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Professional Research Opportunities for Faculty, $30,000 (2012-2014)

International School Psychology Practicum Exchange. Researcher: Gloria L. Miller. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $2,000 (2012-2013)

ELO in Colorado. Research team: Cynthia Hazel and Duan Zhang. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Legacy Foundation, $150,000 (2012-2013)

Creating Online LIS Courses. Research team: Mary C. Stansbury (PI), Shimelis Assefa, Denise Anthony, Xiao Hu, and Krystyna Matusiak. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Teaching and Learning, $15,000 (2011-2013)

Maritime Piracy Seminar. Researcher: Ved Nanda, The Nanda Center, Sturm College of Law. Technical reviewer for Sturm College of Law: Sylvia D. Hall-Ellis. Arsenault Family Foundation, $15,000 (2012-2013).

Early Childhood Librarianship: An Interdisciplinary, Experiential Learning MLIS. Researcher: Mary C. Stansbury (PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. Laura Bush's 21st Century Librarians Program, Institute for Museums and Library Services, $249,066 (2012-2014)

Learning Ecosystem Planning Grant. Research team: Karen Riley (PI), Lyndsay Agans, Kent Seidel, and Shimelis Assefa. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Bill & Melinda Gates Foundation, $281,217 (2011-2012)

Advanced Service Learning Practitioner Faculty Grant, Sylvia D. Hall-Ellis, University of Denver, Center for Community Engagement and Service Learning, $400 (2011)

Evaluating and Enhancing the EspeciallyMe Program. Research team: Lori D. Patton and Nicole Russell. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, Public Good Fellowship Grant, $24,780 (2012)

Choosing Excellence: Let Every Child Bloom. Research team: Shimelis Assefa (PI) and Mary C. Stansbury (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Community Engagement and Service Learning, Public Good Grant, $7,657 (2012)

K-8 STEM Content Specific Professional Development to Improve Elementary Student Achievement. Research team: Kent Seidel (PI), Nicole Russell (Co-PI), Kimberly Hartnett-Edwards, Paul Michalec, Jeff Farmer, Keith Miller, Alegra Reiber, and Nancy Sasaki. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Department of Education, Improving Teacher Quality Grant, (Title II ESEA), $307,299 (2011-2012)

Building a Better Principal Pipeline to Boost Student Achievement. Researcher: Susan Korach. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Denver Public Schools; a subcontract from the Wallace Foundation Grant, $170,000 (2011-2017)

Center of Excellence for Problem Gambling. Research team: J. Mike Faragher (Co-PI) and Bobbie Vollmer (PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Department of Behavioral Health, $79,999 (2011-2012)

An Exploration of Novice Teachers' Core Competencies: Impacts on Student Achievement, and Effectiveness of Preparation. Research team: Kent Seidel (PI), Kathy Green (Co-PI), Kimberly Hartnett-Edwards, and Duan Zhang. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Institute of Education Sciences, Effective Teachers and Effective Teaching, $990,987 (2012-2015)

Project Engage, a DAPRA Grant in partnership with Total Immersion Systems, Inc., and Texas A&M University. Research team: Karen Riley (PI) and Lyndsay Agans (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. U.S. Department of Defense, $49,964 (2011-2014)

User-centered Evaluation of Music Search Engines. Researcher: Xiao Hu. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Faculty Research Fund, $2,931 (2011-2012)

Educational Practicum in Vietnam and China to Promote the Inclusion of Young Children with Disabilities. Researcher: Gloria L. Miller. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Office of Internationalization, $5,000 (2011-2012)

Evaluation of Colorado's Enhancing Quality in Infant-Toddler (EQIT) Initiative. Research team: Virginia R. Maloney and Amanda Moreno. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Buell Foundation, $395,884 (2011-2013)

Creating Engaging Environments to Teach Pre-Algebra Mathematics to Elementary Students. Research team: Álvaro Arias, Mario López, María del Carmen Salazar, and Lyndsay Agans. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver Interdisciplinary Grant, $60,000 (2011-2012)

Morgridge Education Technology Accessible (META) Resource Project. Research team: Lyndsay Agans (PI) and Shimelis Assefa (Co-PI). Technical reviewer for MCE: Sylvia D. Hall-Ellis. Morgridge Family Foundation, $36,000 (2011)

MCE Connect: A 21st Century Framework for Faculty Development. Research team: Bruce Uhrmacher (PI), Shimelis Assefa (Co-PI), Lyndsay Agans (Co-PI), Kimberly Hartnett-Edwards, Norma Hafenstein, Xiao Hu, Paul Michalec (Co-PI), María del Carmen Salazar (Co-PI), and Sandra Snyder-Mondragon. Technical reviewer for MCE: Sylvia D. Hall-Ellis. University of Denver, Center for Teaching and Learning, $22,355 (2011-2012)

Parent Education and Parent Leadership and Advocacy. Research team: Virginia Maloney (PI) and Amanda Moreno, Marsico Institute for Early Learning and Literacy. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Health Foundation, $27,358 (2010-2011)

Intentional School Culture. Researcher: Cynthia Hazel. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Denver Public Schools, $38,040 (2010-2011)

Healthy Eaters, Lifelong Movers: Implementing Evidence-Based School Environment, Policy, and Curricular Changes to Increase Opportunities for Healthy Eating and Physical Activity in Low Income, Rural Colorado. Research team: Elaine Berlansky, University of Colorado at Denver (PI), Nicholas Cutforth, University of Denver (Co-PI), and Allison Reeds. Technical reviewer for MCE: Sylvia D. Hall-Ellis. Colorado Health Foundation, $1,683,277 (2011-2014)

Second Life Learning Community. Research team: Don McCubbrey (PI), Sylvia D. Hall-Ellis (Co-PI), Walter LaMendola, and Paul Novak. University of Denver, Center for Teaching and Learning, $13,000 (2010-2011)

Reintroducing the Value of Law Librarians to Academic and Public Librarians in Colorado through the Identification and Use of Emerging Technologies. Research team: Sylvia D. Hall-Ellis (Co-PI), Stacey L. Bowers (PI), and Christopher Hudson, in partnership with Denver Public Library, Arapahoe Library District, and the Colorado State Supreme Court Library. University of Denver, Center for Community Engagement and Service Learning, $5,773 (2010-11)

Reintroducing the Value of Law Librarians to Public Librarians through the Identification and Use of Emerging Technologies and Resources. Research team: Stacey Bowers (PI) and Sylvia D. Hall-Ellis (Co-PI). American Association of Libraries, Wolters Kluwer Law & Business Grant Program, $2,725 (2010-2011)

Faculty Service Learning Pod. Research team: María del Carmen Salazar and Nicholas Cutforth (Curriculum and Instruction Program), Frank Tuitt (Higher Education Program), Cynthia Hazel and Gloria Miller (Child, Family, and School Psychology Program), and Sylvia D. Hall-Ellis (Library and Information Science Program). University of Denver, Center for Community Engagement and Service Learning, $8,000 (2010-2011)

Lincoln Collaborative. Research team: Antonio Esquibel (Principal, Lincoln High School), María del Carmen Salazar (Curriculum and Instruction Program), and. Sylvia D. Hall-Ellis. School Improvement Grant, Denver Public Schools, $375,000 (2010-2012)

Future LEADers of America: Leaders III. Research team: Denver Public Library (Kristen Svendson, PI), the University of Denver (Sylvia D. Hall-Ellis), the Colorado Chapter of REFORMA (Orlando Archibeque), and the Colorado Association of Libraries (Martin Garnar). Laura Bush's Recruiting Librarians for the 21st Century Program, Institute for Museums and Library Services, $988,366 (2009-2012)

Teaching for Success in the Library Environment: LIS 4030 in Library 2.0. Research team: Deborah S. Grealy (PI) and Sylvia D. Hall-Ellis (Co-PI). University of Denver, Center for Teaching and Learning, $ 9,350 (2009-10)

Writing Group Faculty Grant, Sylvia D. Hall-Ellis, University of Denver, Center for Community Engagement and Service Learning, $750 (2008-2009)

Collaborative Learning Faculty Grant, Sylvia D. Hall-Ellis, University of Denver, Center for Teaching and Learning, $1,500 (2008)

Connecting Information Literacy to Learning. Research team: Lori Micho, Merrie Valliant, Amanda Samland, and Sylvia D. Hall-Ellis. Colorado State Library, LSTA Discretionary Grant Program, $19,548 (2008-2009)

Law Librarianship Fellows Program. Research team: Sylvia D. Hall-Ellis (PI), Stacey Bowers (Co-PI), and Christopher Hudson, Westminster Law Library. Laura Bush's 21st Century Librarians Program, Institute for Museums and Library Services, $999,370 (2008-2012)

Grant Writing in a Cooperative Learning Environment. Sylvia D. Hall-Ellis. University of Denver, Center for Teaching and Learning, $2,000 (2008-2009)

Project Homeless Connect 6 Event Evaluation. Research team: Sylvia D. Hall-Ellis and Duan Zhang. University of Denver Center for Community Engagement and Service Learning, $9,785 (2008)

Project Ecuador: International Learning Service Libraries – A Faculty Development Experience in Ecuador. Sylvia D. Hall-Ellis, Office of Internationalization, University of Denver, $600 (2007)

Project Ecuador: International Learning Service Libraries – A Faculty Development Experience in Ecuador. Sylvia D. Hall-Ellis, International Service Learning Office, University of Denver, $400 (2007)

Project Homeless Connect 5 Event Evaluation. Research team: Sylvia D. Hall-Ellis and Duan Zhang. University of Denver Center for Community Engagement and Service Learning, $1,000 (2007)

Project Homeless Connect 4 Event Evaluation. Research team: Sylvia D. Hall-Ellis and Duan Zhang. University of Denver Center for Community Engagement and Service Learning, $4,595 (2007)

Future LEADers of America: Leaders II. Research team: Denver Public Library (Kristen Svendson, PI), the University of Denver (Sylvia D. Hall-Ellis), and the Colorado Chapter of REFORMA (Orlando Archibeque). Laura Bush's Recruiting Librarians for the 21st Century Program, Institute for Museums and Library Services, $988,518 (2007-2010)

Strategic Planning Assistance for the Denver Medical Library. Sylvia D. Hall-Ellis. Denver Medical Library, Inc., $30,000 (2006-2007)

Destiny Software for JMAC Student Lab. Sylvia D. Hall-Ellis, gift from the Sagebrush Corporation, Minneapolis, Minnesota, $10,000 (2006)

Libraries: Tools for Education and Development Worldwide: A Faculty Development Experience in France. Sylvia D. Hall-Ellis, Office of Internationalization, University of Denver, $800 (2005)

Denver Public Library: Opportunities for Change. Sylvia D. Hall-Ellis, Colorado Community Based Research Network, $2,000 (2005)

Future LEADers of America. Research team: Denver Public Library (Letty Icolari and Steve Taylor, PIs), Emporia State University (Jim Agee), and the University of Denver (Sylvia D. Hall-Ellis). Recruiting Librarians for the 21st Century Program, Institute for Museums and Library Services, $670,315 (2005-2008)

American Association of University Professors Summer Institute Professional Development Grant. Sylvia D. Hall-Ellis, American Association of University Professors, $300 (2005)

Beta Phi Mu Alumni Tea. Sylvia D. Hall-Ellis, gift from the Office of Alumni and Parent and Relations, University of Denver, $2,000 (2005)

Spectrum Software for LIS Student Lab. Sylvia D. Hall-Ellis, gift from the Sagebrush Corporation, Minneapolis, Minnesota, $5,000 (2005)

Developing Research Capacity in Community Organizations and Residents through Training and Technical Assistance. Nicholas J. Cutforth (PI) and Sylvia D. Hall-Ellis. Piton Foundation, $25,000 (2005-2006)

Libraries: Tools for Education and Development Worldwide: A Faculty Development Experience in Argentina. Sylvia D. Hall-Ellis, Office of Internationalization, University of Denver, $500 (2004)

Developing Research Capacity in Community Organizations and Residents through Training and Technical Assistance. Research team: Nicholas J. Cutforth (PI), Gary Lichtenstein and Sylvia D. Hall-Ellis, Piton Foundation, $25,000 (2003-2004)

Increasing Spanish-Speaking and Hispanic Diversity among Library and Information Science Students at the University of Denver: Development of a Student Recruitment Model. Researcher: Sylvia D. Hall-Ellis (PI), Office of Multicultural Excellence, University of Denver, $2,750 (2003-2004)

Collection Development Enrichment to Support the Cataloging & Classification Specialization and School Librarianship within the Library & Information Science Program at the University of Denver. Research team: Sylvia D. Hall-Ellis (PI) and Deborah S. Grealy. Women's Library Association, University of Denver, $4,000 (2003)

Collection Development Enrichment to Support the Library & Information Science Program at the University of Denver. Research team: Deborah S. Grealy (PI) and Sylvia D. Hall-Ellis. Women's Library Association, University of Denver, $4,000 (2001)

Wireless LAN for Library Education. Research team: Deborah S. Grealy (PI) and Sylvia Hall-Ellis. University of Denver, Center for Teaching and Learning, $24,000 (2001-2002)

Upward Bound – A Program for South Texas Youth. Research team: Monte Churchill, Executive Director of Community Relations; Gary R. Sauceda, Outreach Coordinator; Raymond Hernandez, (PI) Associate Dean for Student Success; and Sylvia D. Hall-Ellis, South Texas Community College (McAllen, Texas). U.S. Department of Education, Office of Postsecondary Education, $2,672,003 (1999-2003)

Write Now! Improving Elementary Students' Writing Skills. Research team: Caryl G. Thomason, Assistant Superintendent (PI); Karen Tankersley, Principal; and, Hal Anderson, Director of Technology, Cheyenne Mountain School District 12 (Colorado Springs, Colorado); and Sylvia D. Hall-Ellis. Colorado Department of Education, Educational Telecommunications Unit, $189,453 (1999-2001)

Operation Quick Start -- Distance Learning in Rural Colorado. Research team: Randal Weigum, Technology Coordinator (PI); Bonnie Barns, Director of Federal Programs and Staff Development; and, Adam "Joe" Raskop, Executive Director, Southeastern Board of Cooperative Educational Services (Lamar, Colorado); and Sylvia D. Hall-Ellis. Colorado Department of Education, Educational Telecommunications Unit, $400,000 (1999-2001)

Advanced Technological Training for Information Professionals for the 21st Century. Research team: Mario Reyna, Division of Business Director (PI) and Sylvia D. Hall-Ellis, South Texas Community College (McAllen, Texas). Phi Delta Kappa and the National Science Foundation, $250,000 (1999-2001)

Working Connections – Training Information Technologies Professionals for the 21st Century. Research team: Mario Reyna, Division of Business Director (PI) and Sylvia D. Hall-Ellis, South Texas Community College

(McAllen, Texas). Microsoft Corporation and the American Association of Community Colleges, $1,147,775 (1999-2001)

The ROAD (Research Oriented Amplification of Development to Literacy) Program. Research team: Deborah J. Leong (PI), Metropolitan State University; Elena Bodrova, Metropolitan State University; Dmitri Semenov, Robert J. Marzano, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). Hewlett-Packard Foundation, $25,000 (1998-1999)

Press to Literacy. Research team: Deborah J. Leong (PI), Metropolitan State University; Elena Bodrova, Metropolitan State University; Dmitri Semenov, Robert J. Marzano, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). The Denver Post and the Robert S. McCormick Foundation, $49,818 (1998-1999)

International Telementor Center. Research team: David Neils, David B. Frost (PI), and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). Hewlett-Packard Philanthropy, $100,000 (1998-1999)

America Reads: Providing Tutor Training for the America Reads Challenge. Research team: Louis F. Cicchinelli (PI) and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). U.S. Department of Education, Office of Educational Research and Improvement, $306,000 (1998-1999)

Comprehensive School Reform. Research team: Louis F. Cicchinelli (PI), J. Timothy Waters, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). U.S. Department of Education, Office of Educational Research and Improvement, $285,000 (1998-1999)

Sustainable Energy Education (SEE) Program for Grades 4-8: Preparing Today's Youth for Lifelong Learning and Responsible Actions in Energy Conservation. Research team: Mary Gromko, Colorado Department of Education; Gene McCarthy, Rocky Flats Field Office, U.S. Department of Energy; Gina Kissell, National Renewable Energy Laboratory; and Barbara L. McCombs (PI), Janet L. Bishop, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). State of Colorado Governor's Office for Energy Conservation, $499,936 (1997-1999)

Review of Kansas Curriculum Standards in Mathematics and Language Arts. Research team: Robert J. Marzano (PI), John S. Kendall, David B. Frost, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). Submitted to the Kansas State Department of Education, $25,903 (1998)

International Telementor Center. Research team: David Neils, John Kuglin, Chris Rapp, David B. Frost (PI), and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). Hewlett-Packard Company, $54,752 (1998)

Genesis Mission: Education Public Outreach. Developed in partnership by the Jet Propulsion Laboratory, California Institute of Technology, Los Alamos National Laboratory, Lockheed Martin Astronautics, and the Mid-continent Regional Educational Laboratory (Aurora, Colorado). Research team: John T. Sutton (PI), Alice Krueger, Martha Henry, Greg Rawls, Shae Isaacs, Jeff Johnson, Deb Jordan, Arlene Mitchell, David B. Frost, Jana Caldwell, J. Timothy Waters, and Sylvia D. Hall-Ellis. National Aeronautics and Space Administration, $139,700,000; subcontract award, $4,750,000 (1997-2007)

North Dakota Mathematics Assessment Project. Developed in partnership by the North Dakota Department of Education and Mid-continent Regional Educational Laboratory (Mid-continent Regional Educational Laboratory). Research team: Ann Clapper, Clarence Bina, Greg Gallagher, North Dakota Department of Education; Don Burger (PI), Hillary Michaels, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). U.S. Office of Education, $1,618,214; subcontract award, $389,076 (1997-2001)

Pacific Resources for Education and Learning Distance Education: Project Evaluation. Developed in partnership by the Pacific Educational Community, the Pacific Resources for Education and Learning (PREL), and Mid-continent Regional Educational Laboratory (McREL). Research team: John W. Kofel (PI), Executive Director, PREL; J. Timothy Waters, Executive Director; Joan Buttram, Robert Keller, and Sylvia D. Hall-Ellis, Mid-continent Regional Educational Laboratory (Aurora, Colorado). U.S. Office of Education, Star Schools Program, $10,000,000; subcontract award, $500,000 (1997-2002)

Identification of Bilingual Gifted and Talented Children: A Comprehensive School Grants for Bilingual Education. Developed in partnership by Hidalgo (Texas) ISD, Los Fresnos (Texas) CISD, Progreso (Texas) ISD, University of Texas – Pan American (Edinburg, Texas), and Education Service Center, Region One (Edinburg, Texas). Research team: Hilda Medrano, Dean, College of Education, University of Texas – Pan American; Linda Phemister (PI), Janie Navarro, and Sylvia D. Hall-Ellis, Education Service Center, Region One. U.S. Department of Education, Office of Bilingual Education and Minority Languages Affairs, $1,670,633 (1997-2002)

Academics 2000: First Things First -- The Texas Goals 2000. Developed for Jim Hogg (Hebbronville, Texas) County ISD, Mirando City (Texas) ISD, and San Isidro (Texas) ISD. Research team: Hilda Medrano, Dean, College of Education, University of Texas - Pan American (Edinburg, Texas); Angie Lehmann, Amy Mares, Ellen Gonzalez (PI), and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg, Texas). Texas Education Agency, $339,987 (1997-2000)

Southwestern Bell's Learning Communities Initiative. Developed for Tech Prep of the Rio Grande Valley, Inc. (Harlingen, Texas); the Center for Professional Teacher Development, University of Texas – Brownsville; Teach for America – Rio Grande Valley (McAllen, Texas); and Education Service Center, Region One (Edinburg, Texas). Research team: Patricia G. Bubb, Executive Director (PI), Tech Prep of the Rio Grande Valley; Martin Winchester, Executive Director, Teach for America – Rio Grande Valley; Aileen Johnson, Director, School of Education, University of Texas – Brownsville; and Sylvia D. Hall-Ellis. Southwestern Bell Foundation, $100,000 (1997-1998)

Texas School to Work: Regional Implementation. Developed for Tech Prep of the Rio Grande Valley, Inc. (Harlingen), South Texas Community College (McAllen), Texas State Technical College (Harlingen), Texas Southmost College (Brownsville), Empowerment Zone of the Rio Grande Valley (Mercedes), Project VIDA (Weslaco), Youth Fair Chance (McAllen), and Education Service Center, Region One (Edinburg). Research team: Patricia G. Bubb (PI), Tech Prep of the Rio Grande Valley; Stephen Vassberg, Texas State Technical College; Ellen Trevino, Youth Fair Chance; Wanda Garza, Project VIDA; Leonardo Olivares, University of Texas – Pan American; Michael Bell, South Texas Community College; and Sylvia D. Hall-Ellis. Texas Workforce Commission, $4,250,000 (1997-2002).

Comprehensive Bilingual Education Grant for Hidalgo ISD and Roma ISD. Developed for Hidalgo (Texas) ISD, Roma (Texas) ISD, College of Education, University of Texas – Pan American (Edinburg, Texas), and Education Service Center, Region One (Edinburg, Texas). Research team: Tomas Thomas (PI), Director, Office of Bilingual Education and Sylvia D. Hall-Ellis. U.S. Department of Education, Office of Bilingual Education and Minority Languages Affairs, $1,531,361 (1997-2002)

Lopez High School, Porter High School, Rivera High School, Central Middle School, and Perkins Middle School, Brownsville (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $1,473,611 (1997-1998)

Donna High School, Todd Middle School, and Solis Middle School, Donna (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $704,052 (1997-1998)

Memorial Middle School and Nellie Schunior Middle School, La Joya Independent School District (Texas). Telecommunications Infrastructure Fund Board, $293,641 (1997-1998)

Martin High School, Christen Middle School, Cigarroa Middle School, Lamar Middle School, Laredo (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $1,200,000 (1997-1998).

Lasara Middle School, Lasara (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $246,210 (1997-1998)

Los Fresnos High School and Resaca Middle School, Los Fresnos (Texas) Consolidated Independent School District. Telecommunications Infrastructure Fund Board, $500,000 (1997-1998)

Lyford Junior High School, Lyford (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $354,997 (1997-1998)

Travis Middle School, McAllen (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $269,139 (1997-1998)

Mission High School, Mission (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $297,010 (1997-1998)

Progreso High School, Progreso (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $300,000 (1997-1998)

San Perlita High School, San Perlita (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $293,490 (1997-1998)

Myra Green Junior High School, Raymondville (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $225,000 (1997-1998)

Rio Grande City High School, Ringgold Middle School, Gruella Middle School, Rio Grande City (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $809,934 (1997-1998)

Rio Hondo Junior High School, Rio Hondo (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $297,024 (1997-1998)

Roma High School, Roma (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $333,178 (1997-1998)

San Benito High School, Miller Jordan Junior High School, San Benito (Texas) Consolidated Independent School District. Telecommunications Infrastructure Fund Board, $547,000 (1997-1998)

Santa Maria High School and Santa Maria Middle School, Santa Maria (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $400,000 (1997-1998)

Salvador High School, United Independent School District (Laredo, Texas). Telecommunications Infrastructure Fund Board, $250,879 (1997-1998)

Weslaco High School, Cabaza Middle School, Cuellar Middle School, Weslaco (Texas) Independent School District. Telecommunications Infrastructure Fund Board, $874,242 (1997-1998)

Principals' Assessment and Training Center. Developed for Education Service Center, Region One (Edinburg, Texas). Research team: Roberto Zamora, Leonel Barrera, William H. Parry, and Sylvia D. Hall-Ellis. Texas Principals' Leadership Initiative, Texas Association of Secondary School Principals, Texas Association of Elementary School Principals, and the Sid Richardson Foundation. $461,439 (1996-1998)

The Texas Library Connection -- Integrating and Sharing Resources. Developed for Hidalgo County Library System (McAllen), Cameron County Library System (Brownsville), South Texas Community College (McAllen), the

University of Texas - Pan American (Edinburg), University of Texas – Brownsville, and Education Service Center, Region One (Edinburg). Research team: William R. McGee, Coordinator, Hidalgo County Library System; Joe Garcia, Director, Cameron County Library System; Michael D. Bell, Library Director, South Texas Community College; Eleanor Folger Foster, University Library, University of Texas - Pan American; Jaime Vela (PI), Director of Instructional Technology and Media Services, Ron Pontius, Instructional Technology, and Fabiola Fuentes, Media Services, Education Service Center, Region One; and Sylvia D. Hall-Ellis. Office of Library Media Services, Technology Services, Texas Education Agency. $8,000 (1996-1997)

Developing Leadership Communities for Improving Algebra I for All Students. Developed for a partnership of the Region One Statewide Systemic Initiative Team. Research team: Noel Villarreal (PI), Eduardo Cancino, and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Statewide Systemic Initiative for Reform in Mathematics, Science and Technology Education to The Charles A. Dana Center for Mathematics and Science Education, The University of Texas at Austin. $25,000 (1996-1997)

Academics 2000: First Things First -- The Texas Goals 2000. Developed for Edinburg (Texas) CISD. Research team: Hilda Medrano, Dean, College of Education, University of Texas - Pan American (Edinburg); Helen Jones, Director (PI), Gifted and Talented Education, Edinburg CISD; and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Education Agency, $750,000 (1996-2001)

Texas Teachers Empowered for Achievement in Mathematics (TEXTEAM) Institute for Algebra I. Developed for Hidalgo ISD, Jim Hogg (Hebbronville) County ISD, La Joya ISD, La Villa ISD, Lyford ISD, Mission CISD, Pharr-San Juan-Alamo ISD, Rio Hondo ISD, San Isidro ISD, San Perlita ISD, Santa Maria ISD, Santa Rosa, Sharyland ISD, Valley View ISD, Webb (Laredo) CISD, Weslaco ISD, and Zapata County (Zapata) ISD. Research team: Noel Villarreal (PI), Chuck McInteer, Ellen M. Gonzalez, Education Service Center, Region One (Edinburg); and Sylvia D. Hall-Ellis. Charles A. Dana Center for Mathematics and Science Education, University of Texas - Austin, $15,975 (1996)

Community Learning Center for La Villa, Texas. Developed for Edcouch-Elsa Independent School District. Research team: Noe Gonzalez (PI), Assistant Superintendent, Edcouch-Elsa Independent School District, and Sylvia D. Hall-Ellis. Delta Region Subzone, Rio Grande Valley (Texas) Rural Empowerment Zone, $325,000 (1996-1997)

Monte Alto (Texas) Community Learning Center. Developed for Monte Alto Independent School District. Research team: Homero A. Diaz (PI), Superintendent, Monte Alto Independent School District, and Sylvia D. Hall-Ellis. Delta Region Subzone, Rio Grande Valley (Texas) Rural Empowerment Zone, $300,000 (1996-1997)

Community Learning Center for La Villa, Texas. Developed for La Villa Independent School District. Research team: Sam Gonzalez (PI), Assistant Superintendent, La Villa Independent School District; and Sylvia D. Hall-Ellis. Delta Region Subzone, Rio Grande Valley (Texas) Rural Empowerment Zone, $325,000 (1996-1997)

Building Project for Taylor Elementary. Developed for Mercedes (Texas) Independent School District. Research team: Mrs. Denise Rivera, Librarian, and Eduardo Infante, Principal, Taylor Elementary School; Ismael S. Cantu (PI), Federal Programs Director, Mercedes Independent School District; and Sylvia D. Hall-Ellis. Delta Region Subzone, Rio Grande Valley (Texas) Rural Empowerment Zone, $375,000 (1996-1997)

Border Education Network (BEN) - Distance Education through Cable Television. Developed for Edinburg (Texas) Consolidated Independent School District. Research team: Noe Torres (PI), Library Media Specialist, Magdalena Rosas, Library Media and Technology Coordinator, Edinburg Consolidated Independent School District; and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Edinburg (Texas) Consolidated School District, $250,000 (1996-1997)

Multiservice One-Stop Open Enrollment Charter School. Developed for the Information Referral Resource Assistance, Inc. (McAllen, Texas). Research team: Pablo Perez and Aguie Pena (PI), Executive Director, Information Referral Resource Assistance, Inc.; Roberto Zamora and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Education Agency, $3,066,000 (1996-2001)

Project OK: A Community Youth Opportunities Grant for Summer, 1996. Developed for McAllen (Texas) Independent School District, St. Joseph the Worker Catholic Church, and the Office of Adult Education, Education Service Center, Region One. Research team: Father Bart Flatt, St. Joseph the Worker Catholic Church; Noe Calvillo (PI), Office of Adult Education; Maria Louisa Garcia, McAllen Independent School District; and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Protective and Regulatory Agency through the Office of the Mayor, City of McAllen, $350,000 (1996-2001).

Early Childhood: a Time of Discovery. Developed for a partnership of Lasara ISD, Rio Hondo ISD, San Perlita ISD, the School of Education, University of Texas - Brownsville, and Education Service Center, Region One (Edinburg). Research team: Hugo Rodriguez, Dean, School of Education, University of Texas - Pan American; Leonel Barrera and Jack Damron, Field Service Agents, Ellen M. Gonzalez (PI), Administrator for Student Instructional Services, Ruth Solis, Education Specialist in Special Education, and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Education Agency, $722,090 (1996-2001)

Innovative Gifted and Talented Programs for Early Childhood and Elementary Education Students. Developed for a partnership of Edinburg (Texas) CISD, the School of Education, University of Texas - Pan American (Edinburg), and Education Service Center, Region One (Edinburg). Research team: Hilda Medrano, Dean, School of Education, University of Texas - Pan American; Helen de la Garza (PI), Director of Elementary Curriculum and Instruction, Edinburg CISD; and Sylvia D. Hall-Ellis. Texas Education Agency, $750,000 (1996-2001)

Reading Recovery in Early Elementary Grades. Developed for a partnership of La Villa (Texas) ISD the School of Education, University of Texas - Pan American (Edinburg), and Education Service Center, Region One (Edinburg). Research team: Hilda Medrano, Dean, School of Education, University of Texas - Pan American; Marcario Salinas (PI), Supervisor of Elementary Curriculum and Instruction, La Villa ISD; and Sylvia D. Hall-Ellis, Education Service Center, Region One. Texas Education Agency, $750,000 (1996-2001)

Connected Mathematics Project for Middle and Junior High Students. Developed for a partnership of Michigan State University, The Charles A. Dana Center for Mathematics and Science Education, The University of Texas at Austin and Region One Statewide Systemic Initiative Team. Research team: Jack Damron, Chuck McInteer, Noel Villarreal (PI), and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). National Science Foundation through Michigan State University to the Charles A. Dana Center at the University of Texas - Austin, $539,610 (1996-1999)

Sharing Resources: Testing the Interlibrary Loan Potential of the Texas Library Connection -- Integrating Media Resources. Developed under the sponsorship of the Hidalgo County Library System (McAllen) and Education Service Center, Region One (Edinburg). Research team: William H. McGee, Hidalgo County System Coordinator; Fabiola Fuentes, Library Media; Ronald Pontius (PI), Instructional Technology, Education Service Center, Region One; and Sylvia D. Hall-Ellis. Library Media Services Program, Office of Technology Services, Texas Education Agency, $25,000 (1996-1997)

State and Federal Adult Education JOBS Program. Developed for Education Service Center, Region One (Edinburg). Research team: Noe Calvillo (PI), Director, Adult Education Program, and Sylvia D. Hall-Ellis. Adult Education Program, Texas Education Agency, $532,846 (1995-1996)

Creating Safe and Drug-Free Schools and Communities. Developed for a partnership of the Region One Consortium for Safe and Drug-Free Education Environments. Research team: Clara Contreras (PI), Health Specialist, and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Education Agency, $150,000 (1995-1996)

Developing Leadership Communities for Improving Mathematics Performance for All Students on Title I Campuses. Developed for a partnership of the Region One Statewide Systemic Initiative Team. Research team: Jack Damron, Chuck McInteer, Noel Villarreal (PI), and Sylvia D. Hall-Ellis, Education Service Center, Region One (Edinburg). Texas Statewide Systemic Initiative for Reform in Mathematics, Science and Technology Education to The Charles A. Dana Center for Mathematics and Science Education, The University of Texas at Austin, $50,669 (1995-1996)

The Impact of Library Resource Centers on Academic Achievement in Selected Public Schools in South Texas. Sylvia D. Hall-Ellis (PI). Developed under the sponsorship of the Department of Library Science, College of Education and Applied Science, Sam Houston State University (Huntsville). Texas Association of School Librarians, Children's Services Round Table, and Young Adults Round Table (Austin, Texas), $2,500 (1995)

School Library Media Specialists Fellowship Program. Sylvia D. Hall-Ellis (PI). Developed under the sponsorship of the Department of Library Science, College of Education and Applied Science, Sam Houston State University (Huntsville, Texas). U.S. Department of Education, HEA Title II-B, Library Education and Human Resource Development Program, $44,000 (1995-1996)

Planning for Educational Technology. Research team: Ruth Ann Riggins (PI), Director of Library Media and Technology Services, Donna (Texas) Independent School District; Noe Torres, Education Service Center, Region One; Patricia G. Bubb, Tech Prep of the Rio Grande Calley, Inc; Michael D. Bell, South Texas Community College; and Sylvia D. Hall-Ellis, Sam Houston State University. Developed under the sponsorship of Donna (Texas) Independent School District, Education Service Center, Region One (Edinburg), Tech Prep of the Rio Grande Valley, Inc. (Harlingen), South Texas Community College (McAllen), and the Department of Library Science, Sam Houston State University (Huntsville). Funded through Infusion of Educational Technology Planning Grant Program (H.B. 18: Models). Texas Education Agency, $18,000 (1993-1994)

The Impact of Library Resource Centers on Academic Achievement in Selected Public Schools in South Texas. Sylvia D. Hall-Ellis (PI) and Mary Ann Berry. Developed under the sponsorship of the Department of Library Science, College of Education and Applied Science, Sam Houston State University (Huntsville). Sam Houston State University Research Enhancement Fund, $7,500 (1993-1994)

Electronic Mail Resource Sharing System: Management and Operation of the Iowa Computer-Assisted Network. State Library of Iowa, $294,000 (1986-1993)

The Iowa Locator: A CD-ROM Resource Sharing Tool. State Library of Iowa, $567,000 (1986-1992)

Statewide Database Development: An OCLC Tape Analysis To Determine Feasibility. State Library of South Dakota, $15,000 (1987)

The Iowa Locator: A Feasibility Study. State Library of Iowa, $50,000 (1986)

Electronic Mail Resource Sharing System: A Feasibility Study for Libraries in the State of Iowa. State Library of Iowa, $5,000 (1985)

Access Pennsylvania: A Feasibility Study. State Library of Pennsylvania, $50,000 (1984)

Automating Library Administrative and Management Tasks: Procurement and Distribution of Microcomputer Systems for District Library Centers in the Commonwealth of Pennsylvania. State Library of Pennsylvania, $495,000 (1982)

Sharing Serial Titles Resources: Procurement and Distribution of Microfiche Readers for 550 Libraries in the Commonwealth of Pennsylvania. State Library of Pennsylvania, $500,000 (1982)

Electronic Mail System: A Pilot Project for Libraries throughout the Commonwealth of Pennsylvania. State Library of Pennsylvania, $125,000 (1982)

Literacy Program for Adults in Rural Upstate New York State: Program Outreach and Evaluation - Phase 3. Appalachian Regional Commission, $50,000 (1981)

Database Building: A Cooperative Project of the Finger Lakes Library System (Ithaca), Four County Library System (Binghamton), and the Southern Tier Library System (Corning). New York State Library, $450,000 (1980-1985)

Faces of the Southern Tier: A Professional Photographer-in-Residence. New York State National Endowment for the Humanities, $25,000 (1980)

Literacy Program for Adults in Rural Upstate New York State: Program Implementation - Phase 2. Appalachian Regional Commission, $50,000 (1980)

Small Business Resources Center: A Pilot Project for Rural Public Libraries. New York State Library, $50,000 (1980)

Media Programming: A Professional Development Program for Public Librarians in Upstate New York: A Program in Allegheny, Chemung, Schuyler, Steuben, and Yates Counties. New York State Library, $35,000 (1979)

Information and Library Resources for Inmates and Prisoners in Selected Upstate New York Facilities: A Cooperative Program in Allegheny, Chemung, Schuyler, Steuben, and Yates Counties. New York State Library, $25,000 (1979)

Library Resources for Homebound Adults in Upstate New York: An Outreach Program in Allegheny, Chemung, Schuyler, Steuben, and Yates Counties. New York State Library, $40,000 (1979)

Literacy Program for Adults in Rural Upstate New York State: Program Initiation and Establishment - Phase 1. Funded by the Appalachian Regional Commission. Awarded to the Corning (New York) Public Library, $50,000 (1979)

Information Reference Services to Homebound Adults. New York State Library, $31,000 (1978)

Books-By-Mail Services to Homebound Adults. New York State Library, $31,000 (1978)

System Headquarters Services and Programs for Public Libraries in District 10. Funded through Library Services and Construction Act Title I. Texas State Library and Historical Commission, $880,000 (1976)

County Library Development Grant for Atascosita County. Funded through Library Services and Construction Act Title I. Texas State Library and Historical Commission, $15,000 (1976)

Spanish Language Materials Collection Development Program. Texas State Library and Historical Commission, $40,000 (1975)

Establishment of System Headquarters for Alamo Regional Library System (District 10) Headquartered at the San Antonio (Texas) Public Library. Texas State Library and Historical Commission, $800,000 (1975)

SELECTED CONSULTANTSHIPS

Project Strategic Planning & Funding Proposal Development

Walden University, 2015.

Pikes Peak Library District, 2014.

Douglas County Public Libraries, 2011-2012.

Denver School for Science and Technology, 2008-2011.

Johnson & Wales University, Denver Campus, 2007-2009.

Challenges, Choices, and Images K-12 Charter School, Denver Public Schools, 2007.

Denver Medical Library, Inc., 2006-2010.

Bemis Public Library (Littleton), Ergonomic Design and Facilities Enhancement Consultant, 2003.

Curtis Park Community Center (Denver), Community Technology Center Evaluation and Proposal Development Consultant, 2003.

Colorado Community Based Research Network (Denver), Funding Research Associate, 2002-

Our Lady of the Rosary Academy (Edgewater), Learning Resource Center Development Project Consultant, 2001-

University of Southern Colorado (Pueblo). Technology Integration and Curriculum Enhancement into Higher Education Learning Environment. Proposal Development Consultant, 2000.

Jefferson County Library System (Lakewood), 1999-2000.

Southeastern BOCES (Lamar, CO). Distance Learning Curriculum Content Development Project. Proposal Development Consultant, 1998-2004.

Cheyenne Mountain School District 12 (Colorado Springs, CO). Technology Integration into Elementary Writing Curriculum Project. Proposal Development Consultant, 1998.

Telecommunications Infrastructure Fund Board Round #2 Application, 1997. Technical assistance to the following: La Villa (Texas) Independent School District; Mirando City (Texas) Independent School District; Monte Alto (Texas) Independent School District; San Isidro (Texas) Independent School District.

Telecommunications Infrastructure Fund Board Round #1, 1996. Technical assistance to the following: Brownsville (Texas) Independent School District; Donna (Texas) Independent School District; Edinburg (Texas) Consolidated Independent School District; Harlingen (Texas) Consolidated Independent School District; Jim Hogg County Independent School District (Hebbronville, Texas); La Feria (Texas) Independent School District; La Joya (Texas) Independent School District; La Villa (Texas) Independent School District; Los Fresnos (Texas) Consolidated Independent School District; Progreso (Texas) Independent School District; Raymondville (Texas) Independent School District; Rio Hondo (Texas) Independent School District; San Benito (Texas) Independent School District; Sharyland Independent School District (Mission, Texas); Weslaco (Texas) Independent School District; Zapata (Texas) County Independent School District.

South Texas Community College (McAllen). Strategy Development to Meet the Technology Instructional Needs for Vocational, School-To-Work and Academic Programs, 1994-2000. Development Consultant.

Donna (Texas) Independent School District. Planning for the Infusion of Technology in Middle Schools Serving High At-Risk Students, 1994-1995. Research Associate.

Pettus (Texas) Independent School District. Planning for District-wide Automation in School Library Media Centers: Preparation of Data, and CD-ROM Hardware/Software Configuration Evaluation, 1994. Principal Investigator.

Fort Bend Independent School District (Sugar Land, TX). Planning for District-wide Automation in School Library Media Centers, Preparation of Data, and CD-ROM Hardware/Software Configuration Evaluation, 1994-1995. Principal Investigator.

Northern Waters Library Services (Ashland, WI). Professional Consulting Services to Make Recommendations on the Development of System-Wide Efforts of the Northern Waters Library Service, 1988-1989. Research Associate.

State Library of Iowa (Des Moines). The Iowa Locator Compact Disc to Support Multi-Type Libraries Resource Sharing including Database Design, Building, Production, and Distribution, 1986-1991. Project Director.

State Library of Iowa (Des Moines). Iowa Computer-Assisted Network Development, Enhancement, and Operation, 1985-1992. Technical Director.

State Library of South Dakota (Pierre). Statistical Sampling and Analysis of Multi-Institutional OCLC Archive Tape for Statewide Online Database Building, 1985-1986. Research Associate.

District of Columbia Public Library (Washington, DC). Planning Document for the Retrospective Conversion of Bibliographic Records and Automation Issues in the 1990's, 1984-1985. Research Associate.

State Library of Pennsylvania (Harrisburg). Database Development for High School Libraries, 1984-1985. Research Associate.

MINITEX (Minneapolis, MN). Workshops for Retrospective Conversion, Bar Coding, Library Statistics, and OCLC Serials Format, 1985. Presenter and Research Associate.

ABC Film Consortium (Altoona, Bellefonte, Johnstown, PA). Database Building, System Design, and Implementation of Online Film Booking System, 1983-1985. Database Manager.

Library Automation -- System Design & Implementation

Douglas County Library District (Castle Rock, CO). Library Automation Technical and System Performance Specifications, 2001.

Pharr-San Juan-Alamo Independent School District (Pharr, TX). Library Automation Technical and System Performance Specifications, 1997. Project Development Consultant.

Harlingen (Texas) Independent School District. Planning for District-wide Automation in Junior High School Library Media Centers: Preparation of Data, and CD-ROM Hardware/Software Configuration Evaluation, 1996. Principal Investigator.

Sharyland (Texas) Independent School District. Design and Construction of New Elementary School Library and Technology Resources Center, 1995-1996. Principal Investigator.

South Texas Independent School District (Mercedes, TX). Design and Construction of a New High School Library and Technology Resources Center, 1995-1996. Principal Investigator.

Citizens' Library (Washington, PA). Online Community Resources Files: Design and Implementation, 1984. Research Associate.

Altoona (Pennsylvania) Hospital. Integrated Online System Upgrade Study, 1984. Research Associate.

State Library of South Dakota (Pierre). Automation Plan and State Database Development, 1982-1983. Principal Investigator.

State Library of Pennsylvania (Harrisburg). COM Production and Preparation of the Technical Specifications Document, Procurement, and Distribution of Microcomputers in the Commonwealth of Pennsylvania, 1983-1984. Research Associate.

Altoona (Pennsylvania) Area Public Library. Integrated System (Circulation/Online Catalog) Study, 1982-1983. Principal Investigator.

COPSCAULD (Council of Pennsylvania State College and University Library Directors, Edinboro, PA). Design of Online Media Catalog, 1982. Principal Investigator.

Erie (Pennsylvania) County Library System. Operations Research and Design for Automated Circulation, Hardware Upgrade and Re-Retrospective Conversion, 1982. Research Associate.

Michigan Library Consortium (Lansing, MI). Tape Management System Design and Implementation, 1982-1984. Research Associate.

State Library of Pennsylvania (Harrisburg). Preparation of the Technical Specifications Document, Procurement and Distribution of Microfiche Readers in the Commonwealth of Pennsylvania, 1982. Research Associate.

Community College of Allegheny County, Allegheny Campus (Pittsburgh, PA). Integrated Systems Study, 1981. Principal Investigator.

Community College of Allegheny County, South Campus (West Mifflin, PA). Retrospective Conversion Training Program for Handicapped Students, 1981. Technical Director.

Southern Tier Library System (Corning, NY). OCLC Acquisitions Subsystem Evaluation, 1981. Principal Investigator.

South Central Reference and Research Council (Ithaca, NY). South Central Regional Delivery System, 1980. Principal Investigator.

Southern Tier Library System (Corning, NY). Newspapers on Microfilm in the Chemung-Southern Tier Library System, 1979-1981. Principal Investigator.

Cataloging & Bibliographic Database Building

Challenges, Choices, and Images K-12 Charter School, Denver Public Schools, 2007.

Denver Medical Library, Inc., 2006-2010.

Ricks Gifted and Talented School Library, 2005-

American Humane Society (Englewood), 2004-2006.

Fisher Early Learning Center Library, 2004-

Colorado Community Based Research Network (Denver), 2004-

Boulder Valley School District (Boulder), 2002-2003.

Douglas Public Library District (Castle Rock), 2001-2002.

Bibliographic Center for Research, Inc. (Denver). Original Cataloging, 2001-2003 & 1992-1993. Professional Cataloger. Libraries include Arapahoe Library District (Centennial); Clarke College (Clarke, IA); Kaiser Permanente Center for Health Research (Portland, KS); Douglas County Library System (Carson City, NV); Grinnell College (Grinnell, IA); Kansas State University (Manhattan, KS); Kansas State Library (Topeka, KS); McPherson College (McPherson, KS); Newman University (Wichita, KS); Montana Technical University (Helena, MT); Montana State Library (Helena, MT); University of Texas Southwest Medical Center Library (Dallas, TX); Community College of Southern Nevada (Las Vegas, NV); Pikes Peak Library District (Colorado Springs); Mesa Community College (Grand Junction); Coe College (Cedar Rapids, IA); Briar Cliff College (Sioux City, IA); Tri-Care Health Systems (Aurora); The Penrose-St. Francis Healthcare System (Colorado Springs); Westminster Public Library (Westminster); Lutheran Medical Center Medical Library (Wheat Ridge); American Heritage Center, University of Wyoming (Laramie, WY); National Wildlife Research Center Library (Lakewood); University of Colorado Health Sciences Center, Denison Memorial Library (Denver); Water Resources Library, Denver District Office, Bureau of Land Management (Lakewood); Mesa State Community College Learning Resources Center (Pueblo); Front Range Community College Learning Resources Center (Westminster).

Colorado Historical Society Library (Denver). Original Cataloging of Serials for Multi-Institutional Union Listing Project, 1993. Technical Services Associate. Institutional partners: Denver Art Museum, Denver Museum of Natural History Library, and Denver Botanical Gardens Library.

Arapahoe Library District (Centennial). Original Cataloging of Major Media, 1991.

Original Cataloging Projects: Davis and Elkins College (Elkins, WV), 1989-1990; Waldorf College (Forest City, IA), 1989-1990; Kennametal Corporation (Latrobe, PA), 1984-1985; Pennsylvania Public Libraries Film Center (Harrisburg), 1982-1985.

Retrospective Conversion Projects: North Central Regional Library Service (Mason City, IA), 1987-1988; Southeastern Library Services (Davenport, IA), 1986-1987; Virginia Theological Seminary (Alexandria), 1984-1985; California University of Pennsylvania (California, PA), 1984-1985; Kennametal Corporation (Latrobe, PA), 1984-1985;West Virginia University (Morgantown), 1983-1985; Altoona Area Public Library (Altoona, PA), 1983-1984; George Washington University (Washington, DC), 1983-1984; Indiana University of Pennsylvania (Indiana, PA), 1983-1985; Health Education Center (Pittsburgh), 1983-1984; Central Pennsylvania District Library Center (Bellefonte, PA), 1982-1985; Duquesne University Law School Library (Pittsburgh), 1982-1983; Allegheny County Law Library (Pittsburgh), 1982-1984;Calgon Technical Information Center (Pittsburgh), 1982; Dow Corning Corporation, Technical Information Center (Midland, MI), 1982-1983; Tri-System Public Library Retrospective Conversion Project (Binghamton, Corning, and Ithaca, NY), 1978-1981, Technical Project Director.

Updated March 2020

EXHIBIT 1005

Stefan Savage et al., "Practical Network Support for IP Traceback," ACM SIGCOMM 2000: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, August 2000, pp. 295-306 ("*Savage paper*")

Practical Network Support for IP Traceback

Stefan Savage, David Wetherall, Anna Karlin and Tom Anderson

Department of Computer Science and Engineering
University of Washington
Seattle, WA, USA

Abstract

This paper describes a technique for tracing anonymous packet flooding attacks in the Internet back towards their source. This work is motivated by the increased frequency and sophistication of denial-of-service attacks and by the difficulty in tracing packets with incorrect, or "spoofed", source addresses. In this paper we describe a general purpose traceback mechanism based on probabilistic packet marking in the network. Our approach allows a victim to identify the network path(s) traversed by attack traffic without requiring interactive operational support from Internet Service Providers (ISPs). Moreover, this traceback can be performed "post-mortem" – after an attack has completed. We present an implementation of this technology that is incrementally deployable, (mostly) backwards compatible and can be efficiently implemented using conventional technology.

1. INTRODUCTION

Denial-of-service attacks consume the resources of a remote host or network, thereby denying or degrading service to legitimate users. Such attacks are among the hardest security problems to address because they are simple to implement, difficult to prevent, and very difficult to trace. In the last several years, Internet denial-of-service attacks have increased in frequency, severity and sophistication. Howard reports that between the years of 1989 and 1995, the number of such attacks reported to the Computer Emergency Response Team (CERT) increased by 50 percent per year [25]. More recently, a 1999 CSI/FBI survey reports that 32 percent of respondents detected denial-of-service attacks directed against their sites [16]. Even more worrying, recent reports indicate that attackers have developed tools to coordinate distributed attacks from many separate sites [14].

Unfortunately, mechanisms for dealing with denial-of-service have not advanced at the same pace. Most work in this area has focused on *tolerating* attacks by mitigating their effects on the victim [38, 2, 26, 29, 9]. This approach can provide an effective stop-gap measure, but does not eliminate the problem nor does it discourage attackers. The other option, and the focus of this paper, is to trace attacks back towards their origin – ideally stopping an attacker at the source.

A perfect solution to this problem is complicated by the potential use of indirection to "launder" the true causal origin of an attack. For example, an attack may consist of packets sent from many different slave machines, themselves under the control of a remote master machine. Such indirection may be achieved either explicitly (by compromising the individual slave hosts directly) or implicitly (by sending false requests to the slaves on behalf of the victim – a so-called *reflector*). More challenging still, the true origin and identity of the attacker can be similarly concealed through chains of false computer accounts, call forwarding, and so forth. Consequently, we regard a complete solution – particularly one able to address the forensic needs of law enforcement – as an open problem.

Instead, we address the more limited operational goal of simply identifying the machines that *directly* generate attack traffic and the network path this traffic subsequently follows. We call this the *traceback problem* and it is motivated by the operational need to control and contain attacks. In this setting, even incomplete or approximate information is valuable because the efficacy of measures such as packet filtering improve as they are applied further from the victim and closer to the source.

However, even for our restricted problem, determining the source generating attack traffic is surprisingly difficult due to the stateless nature of Internet routing. Attackers routinely disguise their location using incorrect, or "spoofed", IP source addresses. As these packets traverse the Internet their true origin is lost and a victim is left with little useful information. While there are several ad hoc traceback techniques in use, they all have significant drawbacks that limit their practical utility in the current Internet.

In this paper we present a new approach to the traceback problem that addresses the needs of both victims and network operators. Our solution is to probabilistically mark packets with partial path information as they arrive at routers. This approach exploits the observation that attacks generally comprise large numbers of packets. While each marked packet represents only a "sample" of the path it has traversed, by combining a modest number of such packets a victim can reconstruct the entire path. This allows victims to locate the approximate source of attack traffic without requiring the assistance of outside network operators. Moreover, this determination can be made even after an attack has completed. Both facets of our solution represent substantial improvements over existing capabilities for dealing with flooding-style denial-of-service attacks.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM'00, Stockholm, Sweden.
Copyright 2000 ACM 1-58113-224-7/00/0008...$5.00.

A key practical deployment issue with any modification of Internet routers is to ensure that the mechanisms are efficiently implementable, may be incrementally deployed, and are backwards compatible with the existing infrastructure. We describe a traceback algorithm that adds little or no overhead to the router's critical forwarding path and may be incrementally deployed to allow traceback within the subset of routers supporting our scheme. Further, we demonstrate that we can encode the necessary path information in a way that peacefully co-exists with existing routers, host systems and more than 99% of today's traffic.

The rest of this paper is organized as follows: In Section 2, we describe related work concerning IP spoofing and solutions to the traceback problem. Section 3 outlines our basic approach and section 4 characterizes several abstract algorithms for implementing it. In Section 5 we detail a concrete encoding strategy for our algorithm that can be implemented within the current Internet environment. We also present experimental results demonstrating the effectiveness of our solution. In section 6 we discuss the main limitations and weaknesses of our proposal and potential extensions to address some of them. Finally, we summarize our findings in Section 7.

2. RELATED WORK

It has been long understood that the IP protocol permits anonymous attacks. In his 1985 paper on TCP/IP weaknesses, Morris writes:

> "The weakness in this scheme [the Internet Protocol] is that the source host itself fills in the IP source host id, and there is no provision in ... TCP/IP to discover the true origin of a packet." [31]

In addition to denial-of-service attacks, IP spoofing can be used in conjunction with other vulnerabilities to implement anonymous one-way TCP channels and covert port scanning [31, 3, 24, 44].

There have been several efforts to reduce the anonymity afforded by IP spoofing. Table 1 provides a subjective characterization of each of these approaches in terms of management cost, additional network load, overhead on the router, the ability to trace multiple simultaneous attacks, the ability trace attacks after they have completed, and whether they are preventative or reactive. We also characterize our proposed traceback scheme according to the same criteria. In the remainder of this section we describe each previous approach in more detail.

2.1 Ingress filtering

One way to address the problem of anonymous attacks is to eliminate the ability to forge source addresses. One such approach, frequently called *ingress filtering*, is to configure routers to block packets that arrive with illegitimate source addresses [20]. This requires a router with sufficient power to examine the source address of every packet and sufficient knowledge to distinguish between legitimate and illegitimate addresses. Consequently, ingress filtering is most feasible in customer networks or at the border of Internet Service Providers (ISP) where address ownership is relatively unambiguous and traffic load is low. As traffic is aggregated from multiple ISPs into transit networks, there is no longer enough information to unambiguously determine if a packet arriving on a particular interface has a "legal" source address. Moreover, on many deployed router architectures the overhead of ingress filter becomes prohibitive on high-speed links.

The principal problem with ingress filtering is that its effectiveness depends on widespread, if not universal, deployment. Unfortunately, a significant fraction of ISPs, perhaps the majority, do not implement this service – either because they are uninformed or have been discouraged by the administrative burden[1], potential router overhead and complications with existing services that depend on source address spoofing (e.g. some versions of Mobile IP [33] and some hybrid satellite communications architectures). A secondary problem is that even if ingress filtering were universally deployed at the customer-to-ISP level, attackers could still forge addresses from the hundreds or thousands of hosts within a valid customer network [14].

It is clear that wider use of ingress filtering would dramatically improve the Internet's robustness to denial-of-service attacks. At the same time it is prudent to assume that such a system will never be fullproof – and therefore traceback technologies will continue to be important.

2.2 Link testing

Most existing traceback techniques start from the router closest to the victim and interactively test its upstream links until they determine which one is used to carry the attacker's traffic. Ideally, this procedure is repeated recursively on the upstream router until the source is reached. This technique assumes that an attack remains active until the completion of a trace and is therefore inappropriate for attacks that are detected after the fact, attacks that occur intermittently, or attacks that modulate their behavior in response to a traceback (it is prudent to assume the attacker is fully informed). Below we describe two varieties of link testing schemes, *input debugging* and *controlled flooding*.

2.2.1 Input debugging

Many routers include a feature called *input debugging*, that allows an operator to filter particular packets on some egress port and determine which ingress port they arrived on. This capability is used to implement a trace as follows: First, the victim must recognize that it is being attacked and develop an *attack signature* that describes a common feature contained in all the attack packets. The victim communicates this signature to a network operator, frequently via telephone, who then installs a corresponding input debugging filter on the victim's upstream egress port. This filter reveals the associated input port, and hence which upstream router originated the traffic. The process is then repeated recursively on the upstream router, until the originating site is reached or the trace leaves the ISP's border (and hence its administrative control over the routers). In the later case, the upstream ISP must be contacted and the procedure repeats itself. While such tracing is frequently performed manually, several ISPs have developed tools to automatically trace attacks across their own networks [41].

The most obvious problem with the input debugging approach, even with automated tools, is its considerable management overhead. Communicating and coordinating with network operators at multiple ISPs requires the time, attention and commitment of both the victim and the remote personnel – many of whom have no direct economic incentive to provide aid. If the appropriate network

[1] Some modern routers ease the administrative burden of ingress filtering by providing functionality to automatically check source addresses against the destination-based routing tables (e.g. `ip verify unicast reverse-path` on Cisco's IOS). This approach is only valid if the route to and from the customer is symmetric – generally at the border of single-homed stub networks.

	Management overhead	Network overhead	Router overhead	Distributed capability	Post-mortem capability	Preventative/ reactive
Ingress filtering	Moderate	Low	Moderate	N/A	N/A	Preventative
Link testing						
Input debugging	High	Low	High	Good	Poor	Reactive
Controlled flooding	Low	High	Low	Poor	Poor	Reactive
Logging	High	Low	High	Excellent	Excellent	Reactive
ICMP Traceback	Low	Low	Low	Good	Excellent	Reactive
Marking	Low	Low	Low	Good	Excellent	Reactive

Table 1: Qualitative comparison of existing schemes for combating anonymous attacks and the probabilistic marking approach we propose.

operators are not available, if they are unwilling to assist, or if they do not have the appropriate technical skills and capabilities, then a traceback may be slow or impossible to complete [21].

2.2.2 *Controlled flooding*

Burch and Cheswick have developed a link testing traceback technique that does not require any support from network operators [6]. We call this technique *controlled flooding* because it tests links by flooding them with large bursts of traffic and observing how this perturbs traffic from the attacker. Using a pre-generated "map" of Internet topology, the victim coerces selected hosts along the upstream route into iteratively flooding each incoming link on the router closest to the victim. Since router buffers are shared, packets traveling across the loaded link – including any sent by the attacker – have an increased probability of being dropped. By observing changes in the rate of packets received from the attacker, the victim can therefore infer which link they arrived from. As with other link testing schemes, the basic procedure is then applied recursively on the next upstream router until the source is reached.

While the scheme is both ingenious and pragmatic, it has several drawbacks and limitations. Most problematic among these is that controlled flooding is itself a denial-of-service attack – exploiting vulnerabilities in unsuspecting hosts to achieve its ends. This drawback alone makes it unsuitable for routine use. Also, controlled flooding requires the victim to have a good topological map of large sections of the Internet in addition to an associated list of "willing" flooding hosts. As Burch and Cheswick note, controlled flooding is also poorly suited for tracing distributed denial-of-service attacks because the link-testing mechanism is inherently noisy and it can be difficult to discern the set of paths being exploited when multiple upstream links are contributing to the attack. Finally, like all link-testing schemes, controlled flooding is only effective at tracing an on-going attack and cannot be used "post-mortem".

2.3 Logging

An approach suggested in [37] and [41] is to log packets at key routers and then use data mining techniques to determine the path that the packets traversed. This scheme has the useful property that it can trace an attack long after the attack has completed. However, it also has obvious drawbacks, including potentially enormous resource requirements (possibly addressed by sampling) and a large scale inter-provider database integration problem. We are unaware of any commercial organizations using a fully operational traceback approach based on logging[2].

[2]Historically, the T3-NFSNET *did* log network-to-network traffic statistics and these were used on at least one occasion to trace IP spoofing attacks to an upstream provider [43].

2.4 ICMP Traceback

Since the first writing of this paper, a new traceback proposal bas emerged based on the use of explicit router-generated ICMP traceback messages [4]. The principle idea in this scheme is for every router to sample, with low probability (e.g., 1/20,000), one of the packets it is forwarding and copy the contents into a special ICMP traceback message including information about the adjacent routers along the path to the destination. During a flooding-style attack, the victim host can then use these messages to reconstruct a path back to the attacker. This scheme has many benefits compared to previous work and is in many ways similar to the packet marking approach we have taken. However, there are several disadvantages in the current design that complicate its use. Among these: ICMP traffic is increasingly differentiated and may be filtered or rate limited differently from normal traffic, the ICMP Traceback message relies on an input debugging capability (i.e. the ability to associate a packet with the input port and/or MAC address on which it arrived) that is not available in some router architectures, if only some of the routers participate it seems difficult to positively "connect" traceback messages from participating routers separated by a non-participating router, and finally, it requires a key distribution infrastructure to deal with the problem of attackers sending false ICMP Traceback messages. That said, we believe that the scheme is promising and that hybrid approaches combining it with some of the algorithms we propose are likely to be quite effective.

3. OVERVIEW

Burch and Cheswick mention the possibility of tracing flooding attacks by "marking" packets, either probabilistically or deterministically, with the addresses of the routers they traverse [6]. The victim uses the information in the marked packets to trace an attack back to its source. This approach has not been previously explored in any depth, but has many potential advantages. It does not require interactive cooperation with ISPs and therefore avoids the high management overhead of input debugging. Unlike controlled flooding, it does not require significant additional network traffic and can potentially be used to track multiple attacks. Moreover, like logging, packet marking can be used to trace attacks "post-mortem" – long after the attack has stopped. Finally, we have found that marking algorithms can be implemented without incurring any significant overhead on network routers. The remainder of this paper focuses on fully exploring and characterizing this approach.

3.1 Definitions

Figure 1 depicts the network as seen from a victim V. For the purposes of this paper, V may be a single host under attack, or a network border device such as a firewall or intrusion detection system that represents many such hosts. Every potential *attack origin* A_i is a leaf in a tree rooted at V and every router R_i is an internal

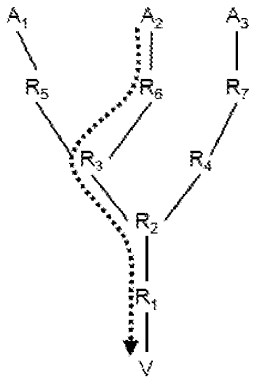


Figure 1: Network as seen from the victim of an attack, V. Routers are represented by R_i, and potential attackers by A_i. The dotted line represents a particular *attack path* between an attacker and the victim.

node along a path between some A_i and V. The *attack path* from A_i is the unique ordered list of routers between A_i and V. For instance, if an attack originates from A_3 then to reach V it must first traverse the path R_6, R_3, R_2, and R_1 – as shown by the dotted line in Figure 1.

The *exact traceback* problem is to determine the attack path and the associated attack origin for each attacker. However, solving this problem is complicated by several practical limitations. The exact attack origin may never be revealed (even MAC source addresses may be spoofed) and a wily attacker may send false signals to "invent" additional routers in the traceback path. We address these issues in section 6, but for now we restrict our discussion to solving a more limited problem. We define the *approximate traceback* problem as finding a candidate attack path for each attacker that contains the true attack path as a suffix. We call this the *valid suffix* of the candidate path. For example, $(R_6, R_5, R_3, R_2, R_1)$ is a valid approximate solution to Figure 1 because it contains the true attack path as a suffix. We say a solution to this problem is *robust* if an attacker cannot prevent the victim from discovering candidate paths containing the valid suffix.

All marking algorithms have two components: a *marking procedure* executed by routers in the network and a *path reconstruction procedure* implemented by the victim. A router "marks" one or more packets by augmenting them with additional information about the path they are traveling. The victim attempts to reconstruct the attack path using only the information in these marked packets. The *convergence time* of an algorithm is the number of packets that the victim must observe to reconstruct the attack path.

3.2 Basic assumptions

The design space of possible marking algorithms is large, and to place our work in context we identify the assumptions that motivate and constrain our design:

- an attacker may generate any packet,
- multiple attackers may conspire,
- attackers may be aware they are being traced,
- packets may be lost or reordered,
- attackers send numerous packets,
- the route between attacker and victim is fairly stable,
- routers are both CPU and memory limited, and
- routers are not widely compromised.

The first four assumptions represent conservative assessments of the abilities of the modern attackers and limitations of the network. Designing a traceback system for the Internet environment is extremely challenging because there is very little that can be trusted. In particular, the attacker's ability to create arbitrary packets significantly constrains potential solutions. When a router receives a packet, it has no way to tell whether that packet has been marked by an upstream router or if the attacker simply has forged this information. In fact, the only invariant that we can depend on is that a packet from the attacker must traverse all of the routers between it and the victim.

The remaining assumptions reflect the basis for our design and deserve additional discussion. First, denial-of-service attacks are only effective so long as they occupy the resources of the victim. Consequently, most attacks are comprised of thousands or millions of packets. Our approach relies on this property because we mark each packet with only a small piece of path state and the victim must observe many such packets to reconstruct the complete path back the the attacker. If many attacks emerge that require only a single packet to disable a host (e.g. ping-of-death [11]), then this assumption may not hold (although we note that even these attacks require multiple packets to *keep* a machine down).

Second, measurement evidence suggests that while Internet routes do change, it is extremely rare for packets to follow many different paths over the short time-scales of a traceback operation (seconds in our system) [32]. This assumption greatly simplifies the role of the victim, since it can therefore limit its consideration to a single primary path for each attacker. If the Internet evolves to allow significant degrees of multi-path routing then this assumption may not hold.

Third, while there have been considerable improvements in router implementation technology, link speeds have also increased dramatically. Consequently, we assert that any viable implementation must have low per-packet overhead and must not require per-flow state. Significantly simpler schemes than ours can be implemented if we assume that routers are not resource constrained.

Finally, since a compromised router can effectively eliminate any information provided by upstream routers, it is effectively indistinguishable from an attacker. In such circumstances, the security violation at the router must be addressed first, before any further traceback is attempted. In normal circumstances, we believe this is an acceptable design point. However, if non-malicious, but information hiding, routing infrastructures become popular, such as described in [22, 35], then this issue may need to be revisited.

4. BASIC MARKING ALGORITHMS

In this section we describe a series of marking algorithms – starting from the most simple and advancing in complexity. Each algorithm attempts to solve the approximate traceback problem in a manner consistent with our assumptions.

Marking procedure at router R:
for each packet w, append R to w

Path reconstruction procedure at victim v:
for any packet w from attacker
extract path $(R_i..R_j)$ from the suffix of w

Figure 2: Node append algorithm.

4.1 Node append

The simplest marking algorithm – conceptually similar to the IP Record Route option [34] – is to append each node's address to the end of the packet as it travels through the network from attacker to victim (see Figure 2). Consequently, every packet received by the victim arrives with a complete ordered list of the routers it traversed – a built-in attack path.

The node append algorithm is both robust and extremely quick to converge (a single packet), however it has several serious limitations. Principal among these is the infeasibly high router overhead incurred by appending data to packets in flight. Moreover, since the length of the path is not known *a priori*, it is impossible to ensure that there is sufficient unused space in the packet for the complete list. This can lead to unnecessary fragmentation and bad interactions with services such as MTU discovery [30]. This problem cannot be solved by reserving "enough" space, as the attacker can completely fill any such space with false, or misleading, path information.

4.2 Node sampling

To reduce both the router overhead and the per-packet space requirement, we can sample the path one node at a time instead of recording the entire path. A single static "node" field is reserved in the packet header – large enough to hold a single router address (i.e. 32 bits for IPv4). Upon receiving a packet, each router chooses to write its address in the node field with some probability p. After enough packets have been sent, the victim will have received at least one sample for every router in the attack path. As stated in section 3, we assume that the attacker sends enough packets and the route is stable enough that this sampling can converge.

Although it might seem impossible to reconstruct an ordered path given only an unordered collection of node samples, it turns out that with a sufficient number of trials, the order can be deduced from the relative number of samples per node. Since routers are arranged serially, the probability that a packet will be marked by a router and then left unmolested by all downstream routers is a strictly decreasing function of the distance to the victim. If we constrain p to be identical at each router, then the probability of receiving a marked packet from a router d hops away is $p(1-p)^{d-1}$. Since this function is monotonic in the distance from the victim, ranking each router by the number of samples it contributes will tend to produce the accurate attack path. The full algorithm is shown in Figure 3.

Putting aside for the moment the difficulty in changing the IP header to add a 32-bit node field, this algorithm is efficient to implement because it only requires the addition of a write and checksum update to the forwarding path. Current high-speed routers already must perform these operations efficiently to update the *time-to-live* field on each hop. Moreover, if $p > 0.5$ then this algorithm is robust against a single attacker because there is no way for an attacker to insert a "false" router into the path's valid suffix by contributing

Marking procedure at router R:
for each packet w
let x be a random number from [0..1)
if $x < p$ then,
write R into w.node

Path reconstruction procedure at victim v:
let $NodeTbl$ be a table of tuples (node,count)
for each packet w from attacker
z := lookup w.node in $NodeTbl$
if z != NIL then
increment z.count
else
insert tuple (w node,1) in $NodeTbl$
sort $NodeTbl$ by count
extract path $(R_i..R_j)$ from ordered node fields in $NodeTbl$

Figure 3: Node sampling algorithm.

more samples than a downstream router, nor to reorder valid routers in the path by contributing more samples than the difference between any two downstream routers.

However, there are also two serious limitations. First, inferring the total router order from the distribution of samples is a slow process. Routers far away from the victim contribute relatively few samples (especially since p must be large) and random variability can easily lead to misordering unless a very large number of samples are observed. For instance, if $d = 15$ and $p = 0.51$, the receiver must receive more than 42,000 packets on average before it receives a *single* sample from the furthest router. To guarantee that the order is correct with 95% certainty requires more than seven times that number.

Second, if there are multiple attackers then multiple routers may exist at the same distance – and hence be sampled with the sample probability. Therefore, this technique is not robust against multiple attackers.

4.3 Edge sampling

A straightforward solution to these problems is to explicitly encode *edges* in the attack path rather than simply individual nodes. To do this, we would need to reserve *two* static address-sized fields, *start* and *end*, in each packet to represent the routers at each end of a link, as well as an additional small field to represent the distance of an edge sample from the victim.

When a router decides to mark a packet, it writes its own address into the start field and writes a zero into the distance field. Otherwise, if the distance field is already zero this indicates that the packet was marked by the previous router. In this case, the router writes its own address into the end field – thereby representing the edge between itself and the previous router. Finally, if the router doesn't mark the packet then it always increments the distance field. This somewhat baroque signaling mechanism allows edge sampling to be incrementally deployed – edges are constructed only between participating routers.

The mandatory increment is *critical* to minimize spoofing by an attacker. When the packet arrives at the victim its distance field represents the number of hops traversed since the edge it contains

Marking procedure at router R:

```
for each packet w
    let x be a random number from [0..1)
    if x < p then
        write R into w.start and 0 into w.distance
    else
        if w.distance = 0 then
            write R into w.end
        increment w.distance
```

Path reconstruction procedure at victim v:

```
let G be a tree with root v
let edges in G be tuples (start,end,distance)
for each packet w from attacker
    if w.distance = 0 then
        insert edge (w.start,v,0) into G
    else
        insert edge (w.start,w.end,w.distance) into G
remove any edge (x,y,d) with d ≠ distance from x to v in G
extract path (R_i..R_j) by enumerating acyclic paths in G
```

Figure 4: Edge sampling algorithm.

was sampled.[3] Any packets written by the attacker will necessarily have a distance greater or equal to the length of the true attack path. Therefore, a single attacker is unable to forge any edge between themselves and the victim (for a distributed attack, of course, this applies only to the closest attacker) and the victim does not have to worry about "chaff" while reconstructing the valid suffix of the attack path. Consequently, since we no longer use the sampling rank approach to distinguish "false" samples, we are free to use arbitrary values for the marking probability p.

The victim uses the edges sampled in these packets to create a graph (much as in Figure 1) leading back to the source, or sources, of attack. The full algorithm is described in Figure 4. Because the probability of receiving a sample is geometrically smaller the further away it is from the victim, the time for this algorithm to converge is dominated by the the time to receive a sample from the furthest router, $\frac{1}{p(1-p)^{d-1}}$ in expectation, for a router d hops away. However, there is a small probability that we will receive a sample from the furthest router, but not from some nearer router. We can bound this effect to a factor of $ln(d)$ by the following argument: We conservatively assume that samples from all of the d routers appear with the same likelihood as the furthest router. Since these probabilities are disjoint, the probability that a given packet will deliver a sample from some router is at least $dp(1-p)^{d-1}$. Finally, as per the well-known *coupon collector* problem, the number of trials required to select one of each of d equi-probable items is $d(ln(d) + O(1))$ [19]. Therefore, the number of packets, X, required for the victim to reconstruct a path of length d has the following bounded expectation:

$$E(X) < \frac{ln(d)}{p(1-p)^{d-1}}$$

For example, if $p = \frac{1}{10}$, and the attack path has a length of 10, then a victim can typically reconstruct this path after receiving 75 packets from the attacker. While this choice of $p = \frac{1}{d}$ is optimal, the convergence time is not overly sensitive to this parameter for the path lengths that occur in the Internet. So long as $p \leq \frac{1}{d}$, the results are generally within a small constant of optimal. In the rest of this paper we will use $p = \frac{1}{25}$ since few paths exceed this length [7, 42, 17]. For comparison, the previous example converges with only 108 packets using $p = \frac{1}{25}$.

This same algorithm can efficiently discern multiple attacks because attackers from different sources produce disjoint edges in the tree structure used during reconstruction. The number of packets needed to reconstruct each path is independent, so the number of packets needed to reconstruct all paths is a linear function of the number of attackers. Finally, edge sampling is also robust. That is it is impossible for any edge closer than the closest attacker to be spoofed, due to the robust distance determination. Conversely, in a distributed attack this also means that it is impossible to trust the contents of any edge *further* away than the closest attacker. As with the ICMP Traceback approach [4], an additional mechanism incorporating a shared secret is required to completely address the problem of attackers spoofing edges.

Of course, a significant practical limitation of this approach is that it requires additional space in the IP packet header and therefore is not backwards compatible. In the next section we discuss a modified version of edge-sampling that addresses this problem, albeit at some cost in performance and a reduction in robustness during large distributed attacks.

5. ENCODING ISSUES

The edge sampling algorithm requires 72 bits of space in every IP packet (two 32-bit IP addresses and 8 bits for distance to represent the theoretical maximum number of hops allowed using IP). It would be possible to directly encode these values into an MPLS label stack [36], to enable traceback within a single homogeneous ISP network. However, our focus is on a heterogeneous environment based purely on IP datagrams. One obvious approach is to store the edge sample data in an IP option, but this is a poor choice for many of the same reasons that the node append algorithm is infeasible – appending additional data to a packet in flight is expensive and there may not be sufficient space to append this data. We could also send this data out-of-band – in a separate packet – but this would add both router and network overhead plus the complexity of a new and incompatible protocol.

Instead, we have developed a modified version of edge sampling that dramatically reduces the space requirement in return for a modest increase in convergence time and a reduction in robustness to multiple attackers. Following an analysis of our algorithm we explore the practical implementation issues and discuss one concrete encoding of this scheme based on overloading the 16-bit IP *identification* field used for fragmentation. Any solution involving such overloading necessarily requires compromises and we stress that our solution reflects only one design point among many potential implementation tradeoffs for this class of algorithm and *does not* necessarily reflect an optimal balance among them.

5.1 Compressed edge fragment sampling

We use three techniques to reduce per-packet storage requirements while preserving robustness. First, we encode each edge in half the space by representing it as the *exclusive-or* (XOR) of the two IP addresses making up the edge, as depicted in Figure 5. When

[3] It is important that distance field is updated using a saturating addition. If the distance field were allowed to wrap, then the attacker could spoof edges close to the victim by sending packets with a distance value close to the maximum.

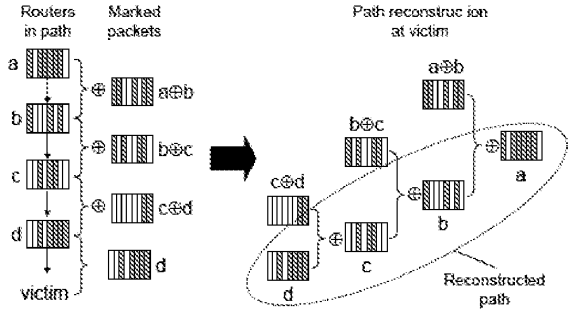


Figure 5: Edge data can be communicated in half the space by sending the XOR of the two nodes (i.e. router IP addresses) making up an edge, rather than sending each node separately. Over time the victim receives the messages d, $c \oplus d$, $b \oplus c$, and $a \oplus b$. By XORing these messages together the original path can be reconstructed.

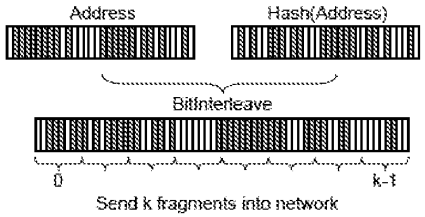


Figure 6: Each router calculates a uniform hash of its IP address once, at startup, using a well-known function. This hash is interleaved with the original IP address (the original address on odd bits, the hash on even bits). The resulting quantity is then broken into k fragments, which the router selects among randomly when marking a packet. Although it is not shown, each of these fragments is further labeled with its offset. The next downstream router uses this offset to select the appropriate fragment to XOR – thereby encoding part of an edge.

some router decides to mark a packet it writes its address, a, into the packet. The following router, b, notices that the distance field is 0 and (assuming it does not mark the packet itself) reads a from the packet, XORs this value with its own address and writes the resulting value, $a \oplus b$, into the packet. We call the resulting value the *edge-id* for the edge between a and b. The edge-ids in the packets received by the victim always contain the XOR of two adjacent routers, except for samples from routers one hop away from the victim, which arrive unmodified. Since $b \oplus a \oplus b = a$, marked packets from the final router can be used to decode the previous edge id, and so on, hop-by-hop until we reach the first router.

Our second modification further reduces our per-packet space requirements by subdividing each edge-id into some number, k, of smaller non-overlapping fragments. When a router decides to mark a packet, it selects one of these fragments at random and stores it in the packet. We use a few additional bits ($log_2 k$) to store the offset of this fragment within the original address – this is necessary to ensure that both fragments making up an edge-id are taken from the same offset. If enough packets are sent by the attacker, the victim will eventually receive all fragments from all edge-ids.

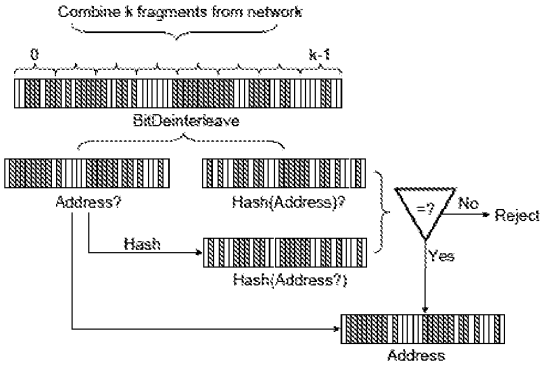


Figure 7: When reconstructing a candidate edge, the victim combines k fragments to produce a bit string. By deinterleaving this string, the address portion and the hash portion are extracted. We recalculate the hash over this address portion using the same hash function used by the router. If the resulting hash is the same as the hash portion extracted, then the address is accepted as valid. This procedure protects against accidentally combining fragments of different edges.

Finally, unlike full IP addresses, edge-id fragments are not unique and multiple fragments from different edge-ids may have the same value. If there are multiple attackers, a victim may receive multiple edge fragments with the same offset and distance. To reduce the probability that we accidentally reconstruct a "false" edge-id by combining fragments from different paths, we add a simple error detection code to our algorithm. We *increase* the size of each router address, and hence each edge-id, by bit-interleaving its IP address with a random hash of itself (depicted in Figure 6). As described earlier, this value is split into fragments, each fragment is selected randomly and stored with an offset, and downstream routers use XOR to combine fragments at the same offset to make up edge-id fragments. The victim constructs *candidate edge-ids* by combining all combinations of fragments at each distance with disjoint offset values. As shown in Figure 7, a candidate edge-id is only accepted if the hash portion matches the data portion for each of its two nodes. As we increase the size of the hash, the probability of a collision is reduced. We provide a describe the full procedure in Figure 8.

The expected number of packets for this algorithm to converge is similar to the edge sampling approach, except now we need k fragments for each edge-id, rather than just one, a total of kd fragments. If we again assume conservatively that each of these fragments is delivered equi-probably with probability $p(1-p)^{d-1}$, the expected number of packets required for path reconstruction is bounded by:

$$E(X) < \frac{k \cdot ln(kd)}{p(1-p)^{d-1}}$$

For example, if there are 8 fragments per edge-id, an attacker is 10 hops away, and $p = \frac{1}{25}$, then a victim can reconstruct the full path after receiving slightly less than 1,300 packets on average. Using techniques similar to those used to show sharp concentration results for the coupon collectors problem, we can further show that the approximate the number of packets required to ensure that a path

Marking procedure at router R:

let R' = BitInterleave(R, Hash(R))
let k be the number of non-overlapping fragments in R'
for each packet w
  let x be a random number from [0..1)
  if $x < p$ then
    let o be a random integer from [0..$k-1$]
    let f be the fragment of R' at offset o
    write f into w.frag
    write 0 into w.distance
    write o into w.offset
  else
    if w.distance = 0 then
      let f be the fragment of R' at offset w.offset
      write $f \oplus w$.frag into w.frag
    increment w.distance

Path reconstruction procedure at victim v:

let $FragTbl$ be a table of tuples (frag,offset,distance)
let G be a tree with root v
let edges in G be tuples (start,end,distance)
let $maxd := 0$
let $last := v$
for each packet w from attacker
  $FragTbl$.Insert(w.frag,w.offset,w.distance)
  if w.distance > $maxd$ then
    $maxd := w$.distance
for $d := 0$ to $maxd$
  for all ordered combinations of fragments at distance d
    construct edge z
    if $d \neq 0$ then
      $z := z \oplus last$
    if Hash(EvenBits(z)) = OddBits(z) then
      insert edge (z,EvenBits(z),d) into G
      $last$:= EvenBits(z);
remove any edge (x,y,d) with $d \neq$ distance from x to v in G
extract path (R_0..R_j) by enumerating acyclic paths in G

Figure 8: Compressed edge fragment sampling algorithm.

can be reconstructed with probability $1 - \frac{1}{e}$ is:

$$\frac{k \cdot ln(kdc)}{p(1-p)^{d-1}}$$

packets. To completely reconstruct the previous path with 95% certainty should require no more than 2150 packets. Many denial-of-service attacks send this many packets in a few seconds.

Finally, we explore the robustness of this algorithm with respect to multiple attackers. For a random hash of length h, the probability of accepting an arbitrarily constructed candidate edge-id is $\frac{1}{2^h}$. In the event that there are m attackers, then at any particular distance d, in the worst case there may be up to m distinct routers.[4] Consequently the probability that any edge-id at distance d is accepted incorrectly is at most:

$$1 - (1 - \frac{1}{2^h})^{m^k}$$

[4] In practice, the number of distinct routers is likely to be smaller for the portion of the path closest to the receiver, since many attackers will still share significant portions of their attack path with one another.

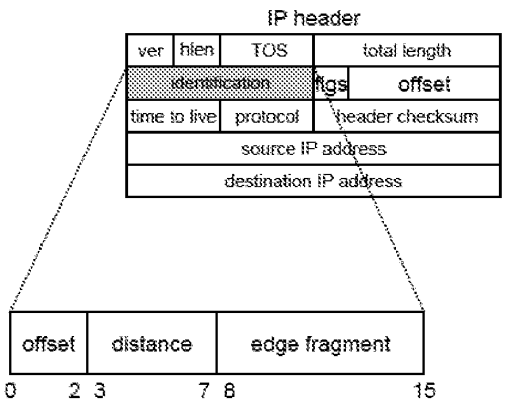


Figure 9: Encoding edge fragments into the IP identification field.

since there are m^k possible combinations of fragments in the worst case. For $h = 32$ and $k = 4$ this means that 100 distinct routers at the same distance (i.e. disjoint attack paths) will be resolved with no errors with a probability of better than 97%. For $h = 32$ and $k = 8$, (the values we use for our implementation) the same certainty can only be provided for 10 distinct routers at the same distance. Our use of the XOR function further complicates reconstruction since all combinations of XOR values must be tried as attack paths diverge. This is somewhat mitigated as the probability of propagating an error from a single edge all the way to the attacker is is extremely small because the resulting edge-id, when XORed with the previous edge-id, must again produce a correct hash.

The most significant drawback to this scheme is the large number of combinations that must be considered as the multiple attack paths diverge. While these combinations can be computed off-line, for large values of k and m even this can become intractable. For example, even with $k = 8$ and $m = 10$, if the separate attack paths diverge such that there are 10 completely independent edges per attacker, this will require roughly a billion combinations to be considered. Consequently, there is a design tension in the size of k – per-packet space overhead is reduced by a larger k, while computational overhead and robustness benefits from a smaller k.

5.2 IP header encoding

To allow for practical deployment requires that we "overload" existing header fields in a manner that will have minimal impact on existing users. This is a difficult task, especially given that even after prodigious effort we require 16 bits of space. Nonetheless, we believe it possible to obtain this space by overloading the 16-bit IP identification field. This field is currently used to differentiate IP fragments that belong to different packets. We describe our proposed encoding below, and then discuss the issues of backwards-compatibility that it raises. However, we note that because the issue of backwards-compatible encoding is largely separate from our traceback algorithms, we could adopt any reasonable encoding that comes to light.

Figure 9 depicts our choice for partitioning the identification field: 3 offset bits to represent 8 possible fragments, 5 bits to represent

the distance, and 8 bits for the edge fragment. We use a 32-bit hash, which doubles the size of each router address to 64 bits. This implies that 8 separate fragments are needed to represent each edge – each fragment indicated by a unique offset value. Finally, 5 bits is sufficient to represent 31 hops, which is more than almost all Internet paths [7, 42, 17].[5]

The observant reader will note that this layout is chosen to allow the highest performance software implementation of our algorithm, which already had a low per-packet router overhead. In the common case, the only modification to the packet is to increment its distance field. Because of its alignment within the packet, this increment precisely offsets the required decrement of the time-to-live field implemented by each router [1]. Consequently, the header checksum *does not need to be altered at all* and the header manipulation overhead could be even lower than in current software-based routers – simply an addition to the distance field, a decrement to the TTL field, and a comparison to check if either has overflowed. In the worst case, our algorithm must read the IP identification field, lookup an edge fragment and XOR it, and fold the write-back into the existing checksum update procedure (a few ALU operations). Of course, for modern ASIC-based routers these optimizations are unnecessary.

As we reuse of the IP identification field, we must address issues of backwards-compatibility for IP fragment traffic. Ultimately, there is no perfect solution to this problem and we are forced to make compromises that disadvantage fragmented traffic. Fortunately, recent measurements suggest that less than 0.25% of packets are fragmented [40, 10]. Moreover, it has long been understood that network-layer fragmentation is detrimental to end-to-end performance [27] so modern network stacks implement automatic MTU discovery to prevent fragmentation regardless of the underlying media [30]. Consequently, we believe that our encoding will interoperate seamlessly with existing protocol implementations in the vast majority of cases.

However, there is a small but real fraction of legitimate traffic that is fragmented, and we wish to ensure that it is not affected by our modifications to the extent that this is possible. Normally if a packet is fragmented, its identification field is copied to each fragment so the receiver can faithfully reassemble the fragments into the original packet. Our marking procedure can violate this property in one of two ways: by writing different values into the identification fields of fragments from the same datagram or by writing the same values into the identification fields of fragments from different datagrams. These two problems present different challenges and have different solutions.

First, a datagram may be fragmented *upstream* from a marking router. If the fragment is subsequently marked and future fragments from the same datagram are not marked consistently then reassembly may fail or data may be corrupted. While the simplest solution to this problem is to simply not mark fragments, an adversary would quickly learn to evade traceback by exploiting this limitation. In fact, some current denial-of-service attacks already use IP fragments to exploit errors in host IP reassembly functions [12]. Instead, we propose an alternative marking mechanism for fragments. We use a separate marking probability, q, for fragments. When we

[5] It is also reasonable to turn off marking on any routers that cannot be directly connected to an attacking host (e.g. core routers). This both reduces the convergence time, and increases the "reach" of the distance field.

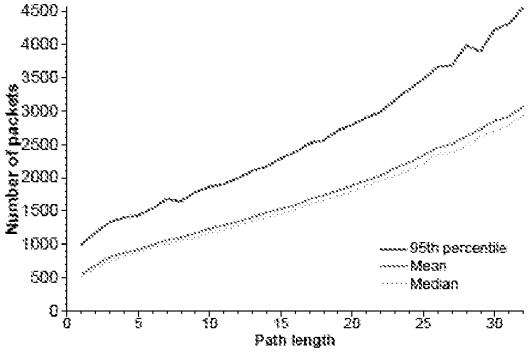


Figure 10: Experimental results for number of packets needed to reconstruct paths of varying lengths. The marking probability, p, is set to $\frac{1}{25}$. Each path length result represents the results of 1,000 independent simulation runs.

decide to mark a fragment, we prepend a new ICMP "echo reply" header, along with the *full* edge data – truncating the tail of the packet. This ICMP packet is considered "marked" and its distance field is set to zero, thereby guaranteeing that the distance field reflects the number of edges traversed on the way to the victim. The packet is consequently "lost" from the standpoint of the receiver, but the edge information is delivered in a way that does not impact legacy hosts. Because we can use the full edge sampling algorithm, q can be more than an order of magnitude smaller than p and yet achieve the same convergence time. This solution increases the loss rate of fragmented flows somewhat (more substantially for longer paths) but preserves the integrity of the data in these flows.

A more insidious problem is presented by fragmentation that occurs *downstream* from a marking router. If a marked packet is fragmented, but one of the fragments is lost, then the remaining fragments may linger in the victim's reassembly buffer for an extended period [5]. Future packets marked by the same router can have the same IP identification value and consequently may be incorrectly reassembled with the previous fragments. One possibility is to leave this problem to be dealt with by higher layer checksums. However, not all higher layer protocols employ checksums, and in any case it is dangerous to rely on such checksums because they are typically designed only for low residual error rates. Another solution is to set the *Don't Fragment* flag on every marked packet. Along rare paths that require fragmentation, this solution will degrade communication between hosts not using MTU path discovery, and may filter marked packets if a reduced MTU edge is close to the victim, but it will never lead to data corruption.

5.3 Experience

We have implemented the marking and reconstruction portions of our algorithm and have tested it using a simulator that creates random paths and originates attacks. In Figure 10 we graph the mean, median and 95th percentile for the number of packets required to reconstruct paths of varying lengths over 1,000 random test runs for each length value. We assume a marking probability of $\frac{1}{25}$. Note that while the convergence time is theoretically exponential in the path length, all three lines appear linear due to the finite path length and appropriate choice of marking probability.

We see that most paths can be resolved with between one and two

thousand packets, and even the longest paths can be resolved with a very high likelihood within four thousand packets. To put these numbers in context, most flooding-style denial of service attacks send many hundreds or thousands of packets each second. The analytic bounds we described earlier are conservative, but in our experience they are no more than 30% higher than our experimental results.

6. LIMITATIONS AND FUTURE WORK

There are still a number of limitations and loose ends in our approach. We discuss the most important of these here:

- backwards compatibility,
- distributed attacks,
- path validation, and
- approaches for determining the attack origin.

6.1 Backwards compatibility

The IP header encoding as we have described it has several practical limitations. It negatively impacts users that require fragmented IP datagrams and is currently incompatible with parts of IPsec [28] (the Authentication Header provides cryptographic protection for the identification field and therefore the field cannot be safely modified by routers). These problems are hardly unique to traceback and are inherent limitations that come about from attempting to coexist with or co-opt protocol features that did not anticipate a new use. One way to address this issue, originally proposed by John Hawkinson, is to selectively enable traceback support in response to operational needs. A "request for traceback" from a particular network could be encoded as a BGP attribute in the network's route advertisement. Routers receiving such an advertisement would enable traceback support on packets destined for that network. Since a network requesting such support is presumably already suffering under an attack, any minor service degradation for fragmented or authenticated flows would be acceptable.

Finally, our scheme does not address implementation in IPv6, the proposed successor to IPv4, which does not have an identification field [18]. While we do not attempt to propose a complete encoding here, we believe that the same techniques we have proposed could also be employed within IPv6, perhaps by overloading the 24-bit *flow label* field (without any further modifications this would result in roughly a factor of three increase in the number of packets required to reconstruct a path).

6.2 Distributed attacks

For large distributed attacks, the practical implementation we have described has serious limitations due to the difficulty in correctly grouping fragments together. Consequently, the probability of misattributing an edge, as well as the amount of state needed to evaluate this decision, increases very quickly with the fan-out of an attack. There is significant future work in designing alternative encoding methods that scale their robustness as they receive more data.

6.3 Path validation

Some number of the packets sent by the attacker are unmarked by intervening routers. The victim cannot differentiate between these packets and genuine marked packets. Therefore an attacker could insert "fake" edges by carefully manipulating the identification fields in the packets it sends. While the distance field prevents an attacker from spoofing edges between it and the victim – what we call the *valid suffix* – nothing prevents the attacker from spoofing extra edges past the end of the true attack path.

There are several ways to identify the valid suffix within a path generated by the reconstruction procedure. With minimal knowledge of Internet topology one can differentiate between routers that belong to transit networks (e.g. ISPs) and those which belong to stub networks (e.g. enterprise networks). Generally speaking, a valid path will never enter a stub network and then continue into a transit network. Moreover, simple testing tools such as `traceroute` should enable a victim to determine if two networks do, in fact, connect. More advanced network maps [8, 23] can resolve this issue in an increasing number of cases.

A more general mechanism is to provide each router with a "secret" that is sent along with each marked packet (perhaps in the single unallocated bit in the IP flags field). When the victim wants to validate a router in the path, it contacts the associated network (possibly out of band, via telephone or e-mail) and obtains the secret used by the router at the time of the attack. To guard against replay, the secret can be time-varying and hashed with the packet contents. Since the attacker will not know the router's secret, they will not be able to include the proper bit in their forged edge-id fragments. By eliminating edge-ids for which the secret in their constituent fragments can not be validated, we can prune a candidate attack path to only include the valid suffix.

6.4 Attack origin detection

While our IP-level traceback algorithm could be an important part of the solution for stopping denial-of-service attacks, it is by no means a complete solution. Our algorithm attempts to determine the approximate origin of attack traffic – in particular, the earliest traceback-capable router involved in forwarding attack traffic from the source that directly generated it. As mentioned earlier, there are a number of reasons why this may differ from the true source of the attack: attackers can hide their true identities by "laundering" attacks through third parties, either indirectly (e.g. smurf attacks [13] or DNS reflectors [15]) or directly via compromised "stepping stone" machines or IP-in-IP tunnels. While there is ongoing work on following attackers through intermediate hosts [45, 39], there are still significant challenges in developing a generally applicable and universally deployable solution to this problem. One interesting possibility enabled by the packet marking approach is to extend traceback across "laundering points". For example, identifying marks could be copied from a DNS request packet into the associated DNS reply, thereby allowing the victim to trace the full causal path.

Even in absence of such "laundering", our approach does not reveal the actual host originating the attack. Moreover, since hosts can forge both their IP source address and MAC address the origin of a packet may never be explicitly visible. On shared media such as FDDI rings, this problem can only be solved by explicit testing. However, on point-to-point media, the input port a packet arrives on is frequently enough to determine its true origin. On other media, there may be a MAC address, cell number, channel, or other hint that would help to locate the attack origin. In principle, our algorithm could be modified to report this information by occasionally marking packets with a special edge-id representing a link between the router and the input port on which the packet arrived (or other

"hint" information). We have not explored the design of such a feature in any depth.

Finally, traceback is only effective at finding the source of *attack traffic*, not necessarily the *attacker* themselves. Stopping an attack may be sufficient to eliminate an immediate problem, but long term disincentives may require a legal remedy and therefore the forensic means to determine an attacker's identity. Even with perfect traceback support, unambiguously identifying a sufficiently skilled and paranoid attacker is likely to require cooperation from law enforcement and telecommunications organizations.

7. CONCLUSION

In this paper we have argued that denial-of-service attacks motivate the development of improved traceback capabilities and we have explored traceback algorithms based on packet marking in the network. We have shown that this class of algorithm, best embodied in *edge sampling*, can enable efficient and robust multi-party traceback that can be incrementally deployed and efficiently implemented. As well, we have developed variant algorithms that sacrifice convergence time and robustness for reduced per-packet space requirements. Finally, we have suggested one potential deployment strategy using such an algorithm based on overloading existing IP header fields and we have demonstrated that this implementation is capable of fully tracing an attack after having received only a few thousand packets. We believe our solution represents a valuable first step towards an automated network-wide traceback facility. Several areas remain to be addressed in future work, such as the combination of widely distributed attacks and points of indirection such as reflectors.

Acknowledgments

This paper has benefited from conversations with many different people – far more than can be acknowledged completely here. Still we would like to particularly thank Bill Cheswick and Hal Burch for early access to their work in this area, Sean McCreary and KC Claffy for access to their packet trace data, Alec Wolman for help with jgraph, Przemek Pardyak, Mike Swift, Neil Spring, Gretta Bartels, Robert Grimm and Geoff Voelker for commenting on early drafts of the paper, and constructive feedback from John Hawkinson, Robert Stone, Jeff Mogul, Robert Moskowitz, Greg Minshall, Tony Li, Curtis Villamizar, Steve Corbato, and countless others. Finally, we thank both the anonymous reviewers and our shepherds, Vern Paxson and Craig Partridge, for their efforts.

8. REFERENCES

[1] F. Baker. Requirements for IP Version 4 Routers. RFC 1812, June 1995.

[2] G. Banga, P. Druschel, and J. Mogul. Resource Containers: A New Facility for Resource Management in Server Systems. In *Proceedings of the 1999 USENIX/ACM Symposium on Operating System Design and Implementation*, pages 45–58, Feb. 1999.

[3] S. M. Bellovin. Security Problems in the TCP/IP Protocol Suite. *ACM Computer Communications Review*, 19(2):32–48, Apr. 1989.

[4] S. M. Bellovin. ICMP Traceback Messages. Internet Draft: draft-bellovin-itrace-00.txt, Mar. 2000.

[5] R. Braden. Requirements for Internet Hosts – Communication Layers. RFC 1122, Oct. 1989.

[6] H. Burch and B. Cheswick. Tracing Anonymous Packets to Their Approximate Source. Unpublished paper, Dec. 1999.

[7] R. L. Carter and M. E. Crovella. Dynamic Server Selection Using Dynamic Path Characterization in Wide-Area Networks. In *Proceedings of the 1997 IEEE INFOCOM Conference*, Kobe, Japan, Apr. 1997.

[8] B. Cheswick and H. Burch. Internet Mapping Project. `http://cm.bell-labs.com/who/ches/map/index.html`, 2000.

[9] Cisco Systems. Configuring TCP Intercept (Prevent Denial-of-Service Attacks). Cisco IOS Documentation, Dec. 1997.

[10] K. Claffy and S. McCreary. Sampled Measurements from June 1999 to December 1999 at the AMES Inter-exchange Point. Personal Communication, Jan. 2000.

[11] Computer Emergency Response Team. CERT Advisory CA-96.26 Denial-of-Service Attack via pings. `http://www.cert.org/advisories/CA-96.26.ping.html`, Dec. 1996.

[12] Computer Emergency Response Team. CERT Advisory CA-97.28 IP Denial-of-Service Attacks. `http://www.cert.org/advisories/CA-97.28.smurf.html`, Dec. 1997.

[13] Computer Emergency Response Team. CERT Advisory CA-98.01 "smurf" IP Denial-of-Service Attacks. `http://www.cert.org/advisories/CA-98.01.smurf.html`, Jan. 1998.

[14] Computer Emergency Response Team. CERT Advisory CA-2000-01 Denial-of-Service Developments. `http://www.cert.org/advisories/CA-2000-01.html`, Jan. 2000.

[15] Computer Emergency Response Team. CERT Incident Note IN-2000-04 Denial-of-Service Attacks using Nameservers. `http://www.cert.org/incident_notes/IN-200-04.html`, Apr. 2000.

[16] Computer Security Institute and Federal Bureau of Investigation. 1999 CSI/FBI Computer Crime and Security Survey. Computer Security Institute publication, Mar. 1999.

[17] Cooperative Associationfor Internet Data Analysis. Skitter Analysis. `http://www.caida.org/Tools/Skitter/Summary/`, 2000.

[18] S. Deering. Internet protocol, version 6 (ipv6). RFC 2460, Dec. 1998.

[19] W. Feller. *An Introduction to Probability Theory and Its Applications (2nd edition)*, volume 1. Wiley and Sons, 1966.

[20] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing. RFC 2267, Jan. 1998.

[21] J. Glave. Smurfing Cripples ISPs. Wired Technolgy News: (`http://www.wired.com/news/news/technology/story/9506.html`), Jan. 1998.

[22] I. Goldberg and A. Shostack. Freedom Network 1.0 Architecture and Protocols. Zero-Knowledge Systems White Paper, Nov. 1999.

[23] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery. In *Proceedings of the 2000 IEEE INFOCOM Conference*, Tel Aviv, Israel, Mar. 2000.

[24] L. T. Heberlein and M. Bishop. Attack Class: Address Spoofing. In *1996 National Information Systems Security Conference*, pages 371–378, Baltimore, MD, Oct. 1996.

[25] J. D. Howard. *An Analysis of Security Incidents on the Internet*. PhD thesis, Carnegie Mellon University, Aug. 1998.

[26] P. Karn and W. Simpson. Photuris: Session-Key Management Protocol. RFC 2522, Mar. 1999.

[27] C. Kent and J. Mogul. Fragmentation Considered Harmful. In *Proceedings of the 1987 ACM SIGCOMM Conference*, pages 390–401, Stowe, VT, Aug. 1987.

[28] S. Kent and R. Atkinson. Security architecture for the internet protocol. RFC 2401, Nov. 1998.

[29] C. Meadows. A Formal Framework and Evaluation Method for Network Denial of Service. In *Proceedings of the 1999 IEEE Computer Security Foundations Workshop*, Mordano, Italy, June 1999.

[30] J. Mogul and S. Deering. Path MTU Discovery. RFC 1191, Nov. 1990.

[31] R. T. Morris. A Weakness in the 4.2BSD Unix TCP/IP Software. Technical Report Computer Science #117, AT&T Bell Labs, Feb. 1985.

[32] V. Paxson. End-to-End Routing Behavior in the Internet. *IEEE/ACM Transactions on Networking*, 5(5):601–615, Oct. 1997.

[33] C. Perkins. IP Mobility Support. RFC 2002, Oct. 1996.

[34] J. Postel. Internet Protocol. RFC 791, Sept. 1981.

[35] M. G. Reed, P. F. Syverson, and D. M. Goldschlag. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998.

[36] E. C. Rosen, Y. Rekhter, D. Tappan, D. Farinacci, G. Fedorkow, T. Li, and A. Conta. MPLS Label Stack Encoding. Internet Draft: draft-ietf-mpls-label-encaps-07.txt (expires March 2000), Sept. 1998.

[37] G. Sager. Security Fun with OCxmon and cflowd. Presentation at the Internet 2 Working Group, Nov. 1998.

[38] O. Spatscheck and L. Peterson. Defending Against Denial of Service Attacks in Scout. In *Proceedings of the 1999 USENIX/ACM Symposium on Operating System Design and Implementation*, pages 59–72, Feb. 1999.

[39] S. Staniford-Chen and L. T. Heberlein. Holding Intruders Accountable on the Internet. In *Proceedings of the 1995 IEEE Symposium on Security and Privacy*, pages 39–49, Oakland, CA, May 1995.

[40] I. Stoica and H. Zhang. Providing Guaranteed Services Without Per Flow Management. In *Proceedings of the 1999 ACM SIGCOMM Conference*, pages 81–94, Boston, MA, Aug. 1999.

[41] R. Stone. CenterTrack: An IP Overlay Network for Tracking DoS Floods. In *to appear in Proceedings of thje 2000 USENIX Security Symposium*, Denver, CO, July 2000.

[42] W. Theilmann and K. Rothermel. Dynamic Distance Maps of the Internet. In *Proceedings of the 2000 IEEE INFOCOM Conference*, Tel Aviv, Israel, Mar. 2000.

[43] C. Villamizar. Personal Communication, Feb. 2000.

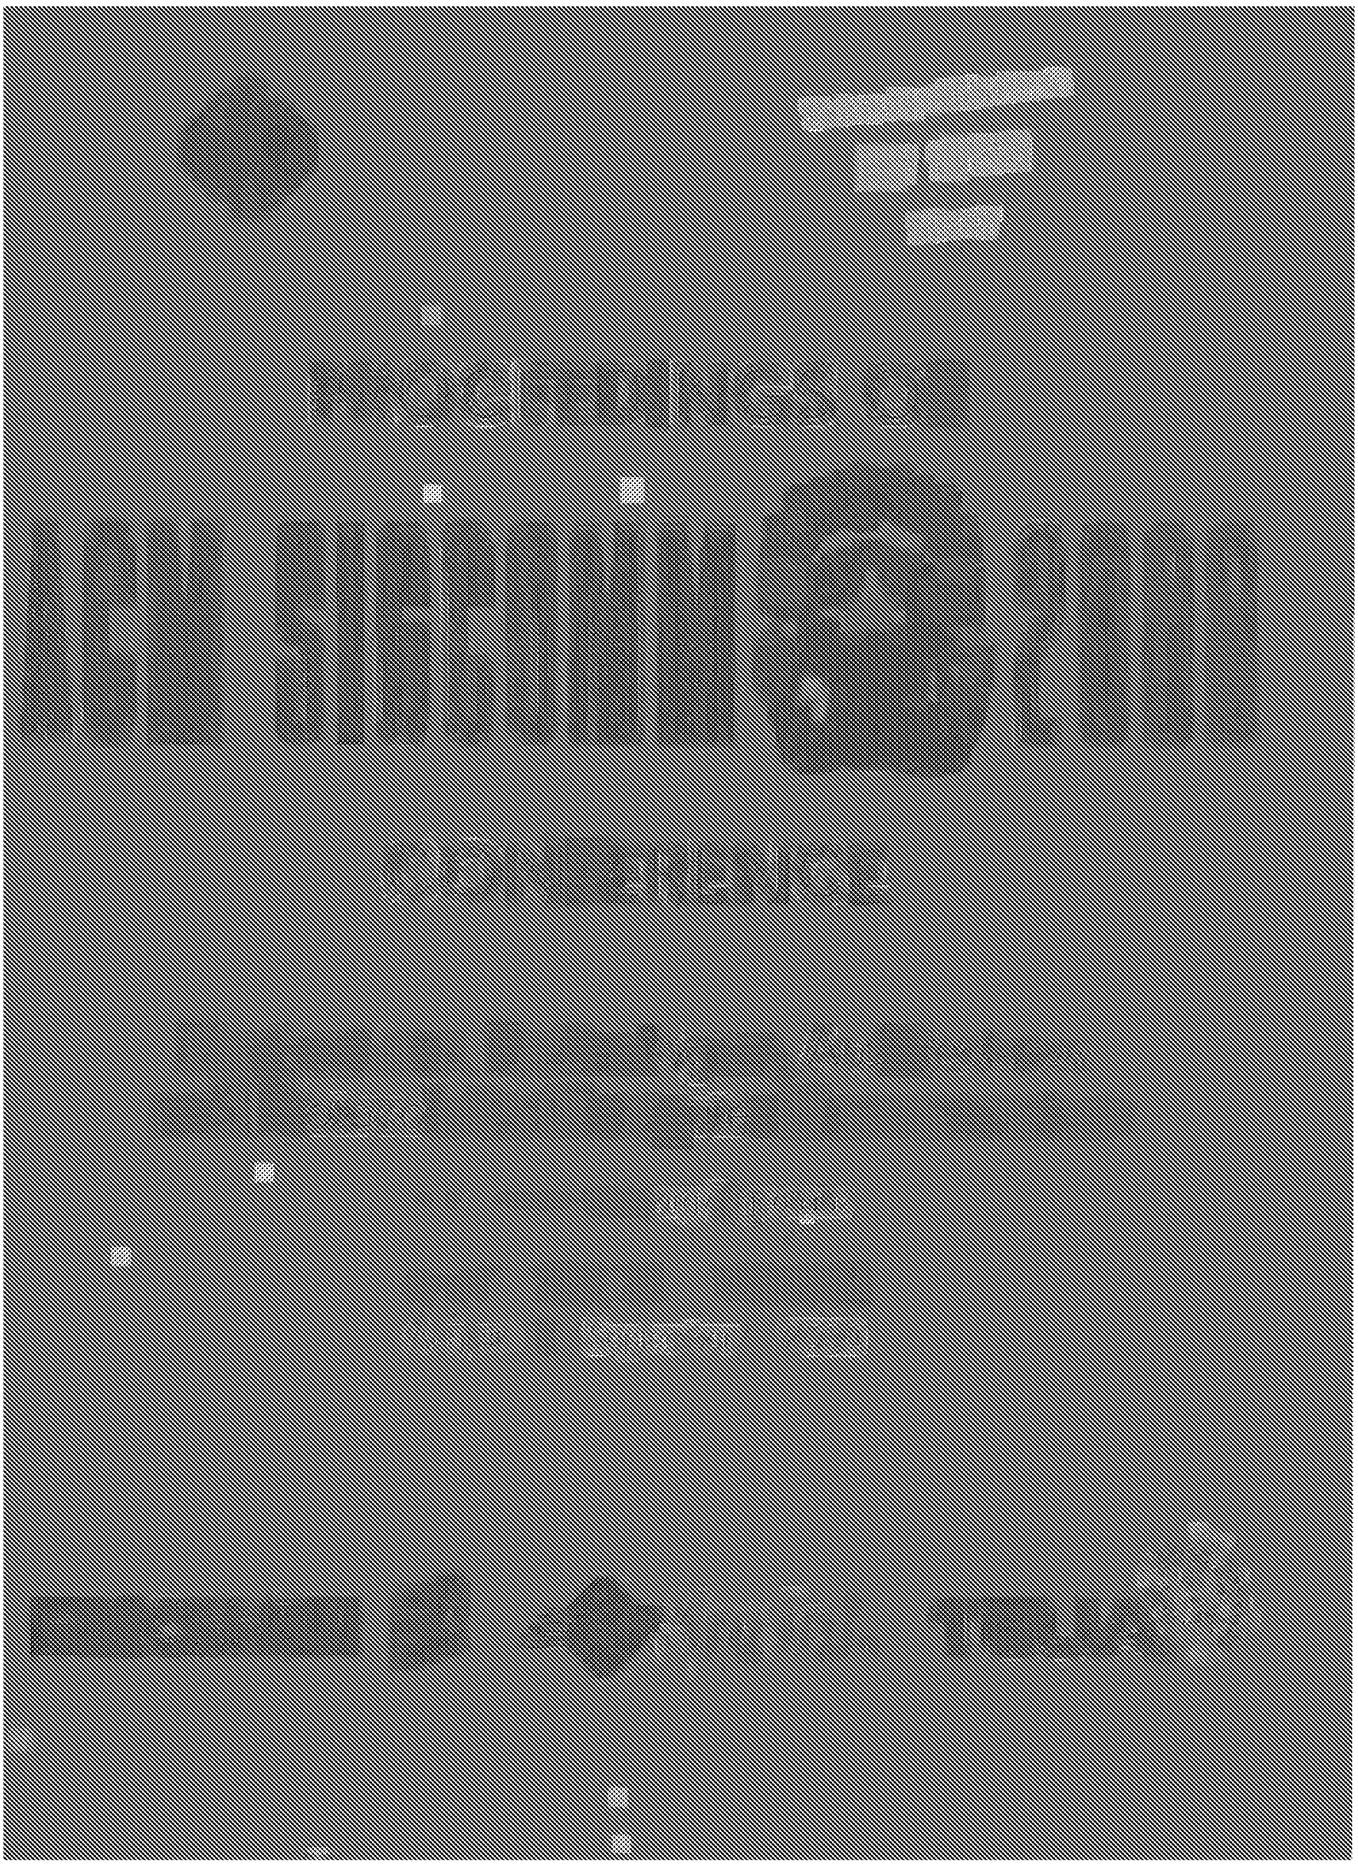
[44] M. Vivo, E. Carrasco, G. Isern, and G. O. Vivo. A review of port scanning techniques. *ACM Computer Communications Review*, 29(2):41–48, Apr. 1999.

[45] Y. Zhang and V. Paxson. Stepping Stone Detection. In *to appear in Proceedings of thje 2000 USENIX Security Symposium*, Denver, CO, July 2000.

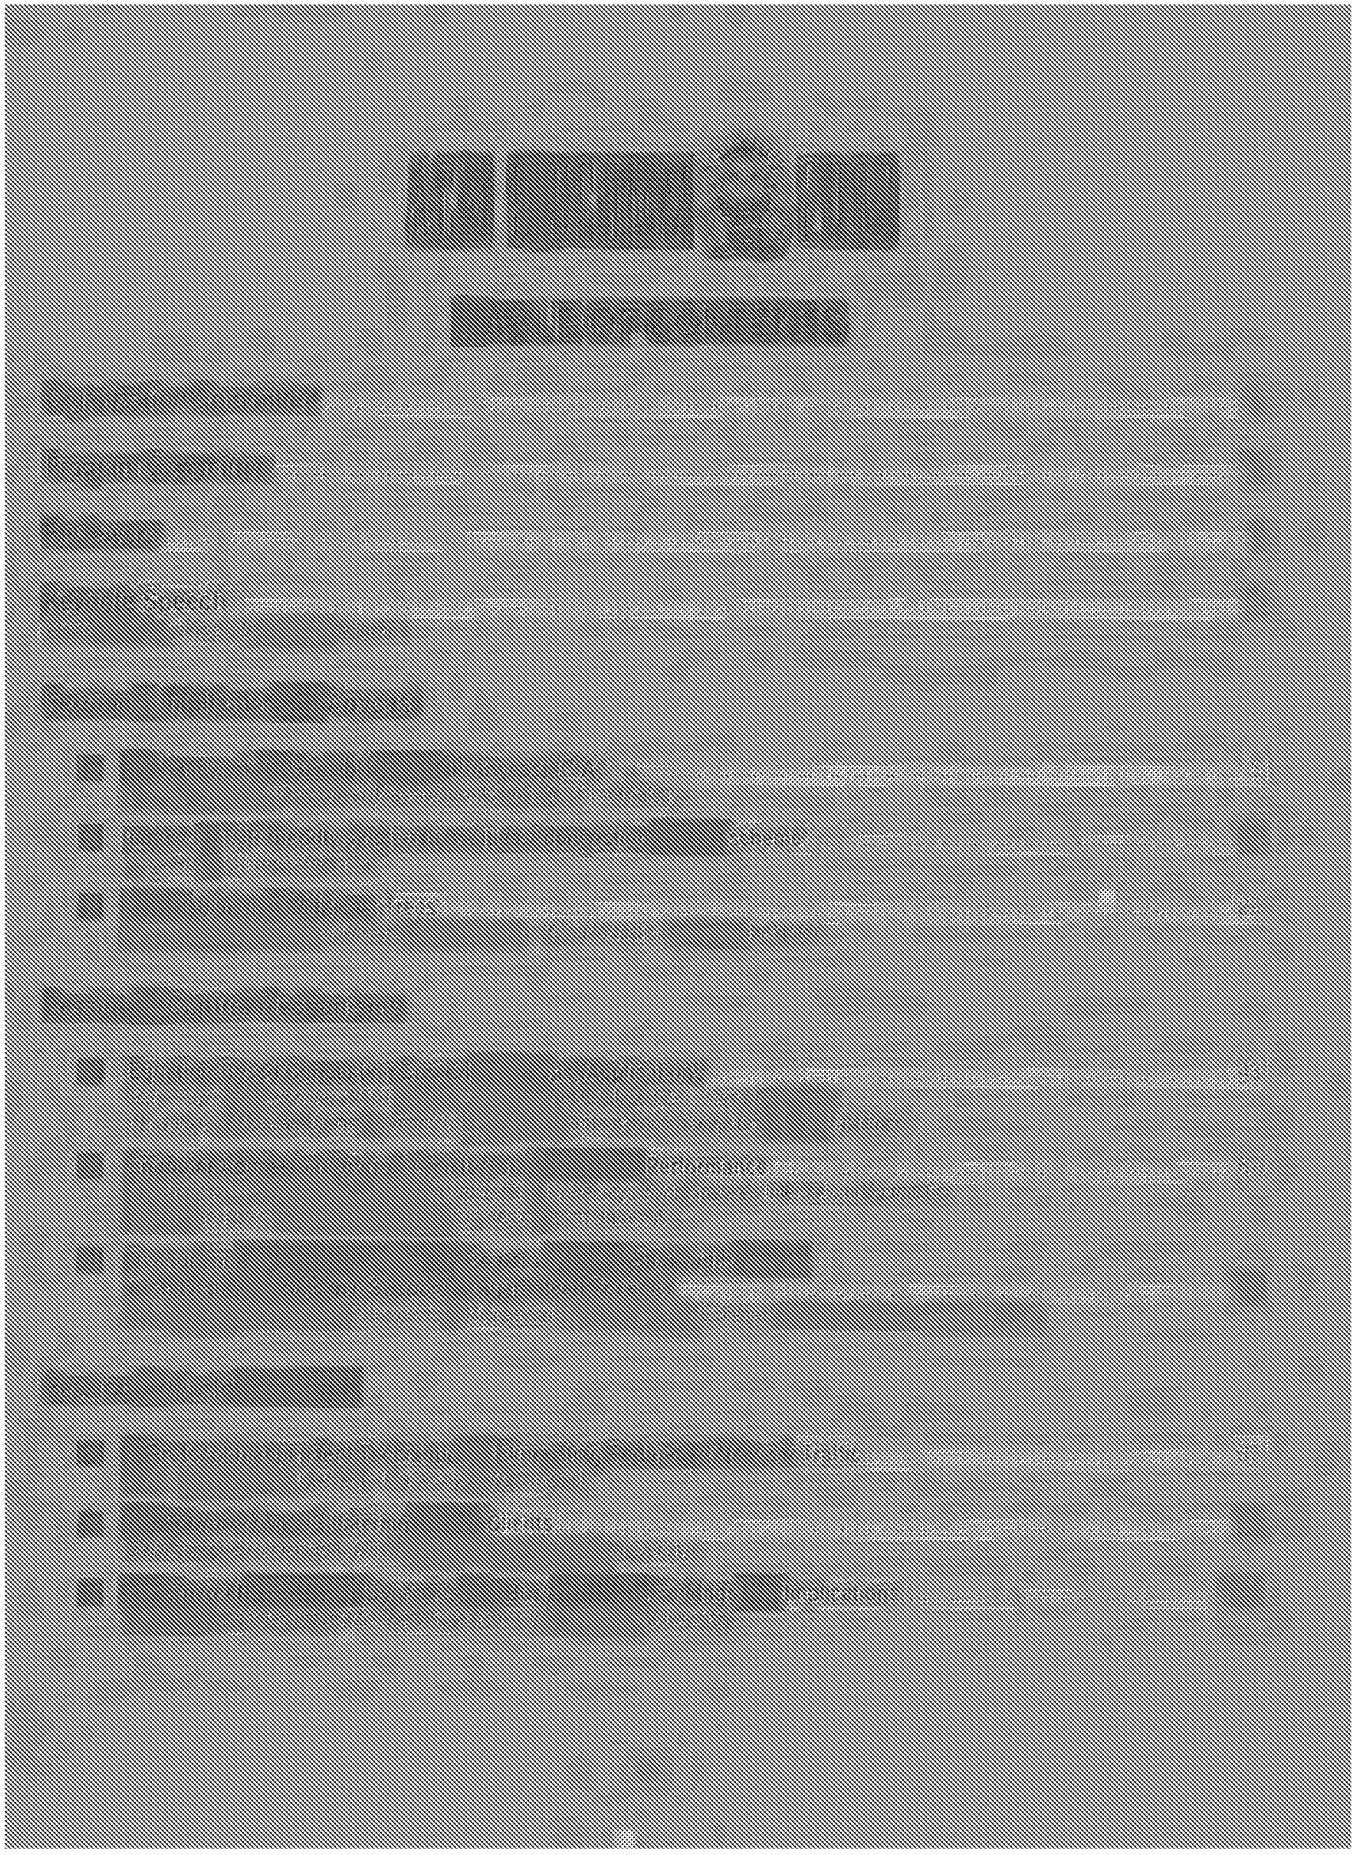
Declaration of Dr. Hall-Ellis

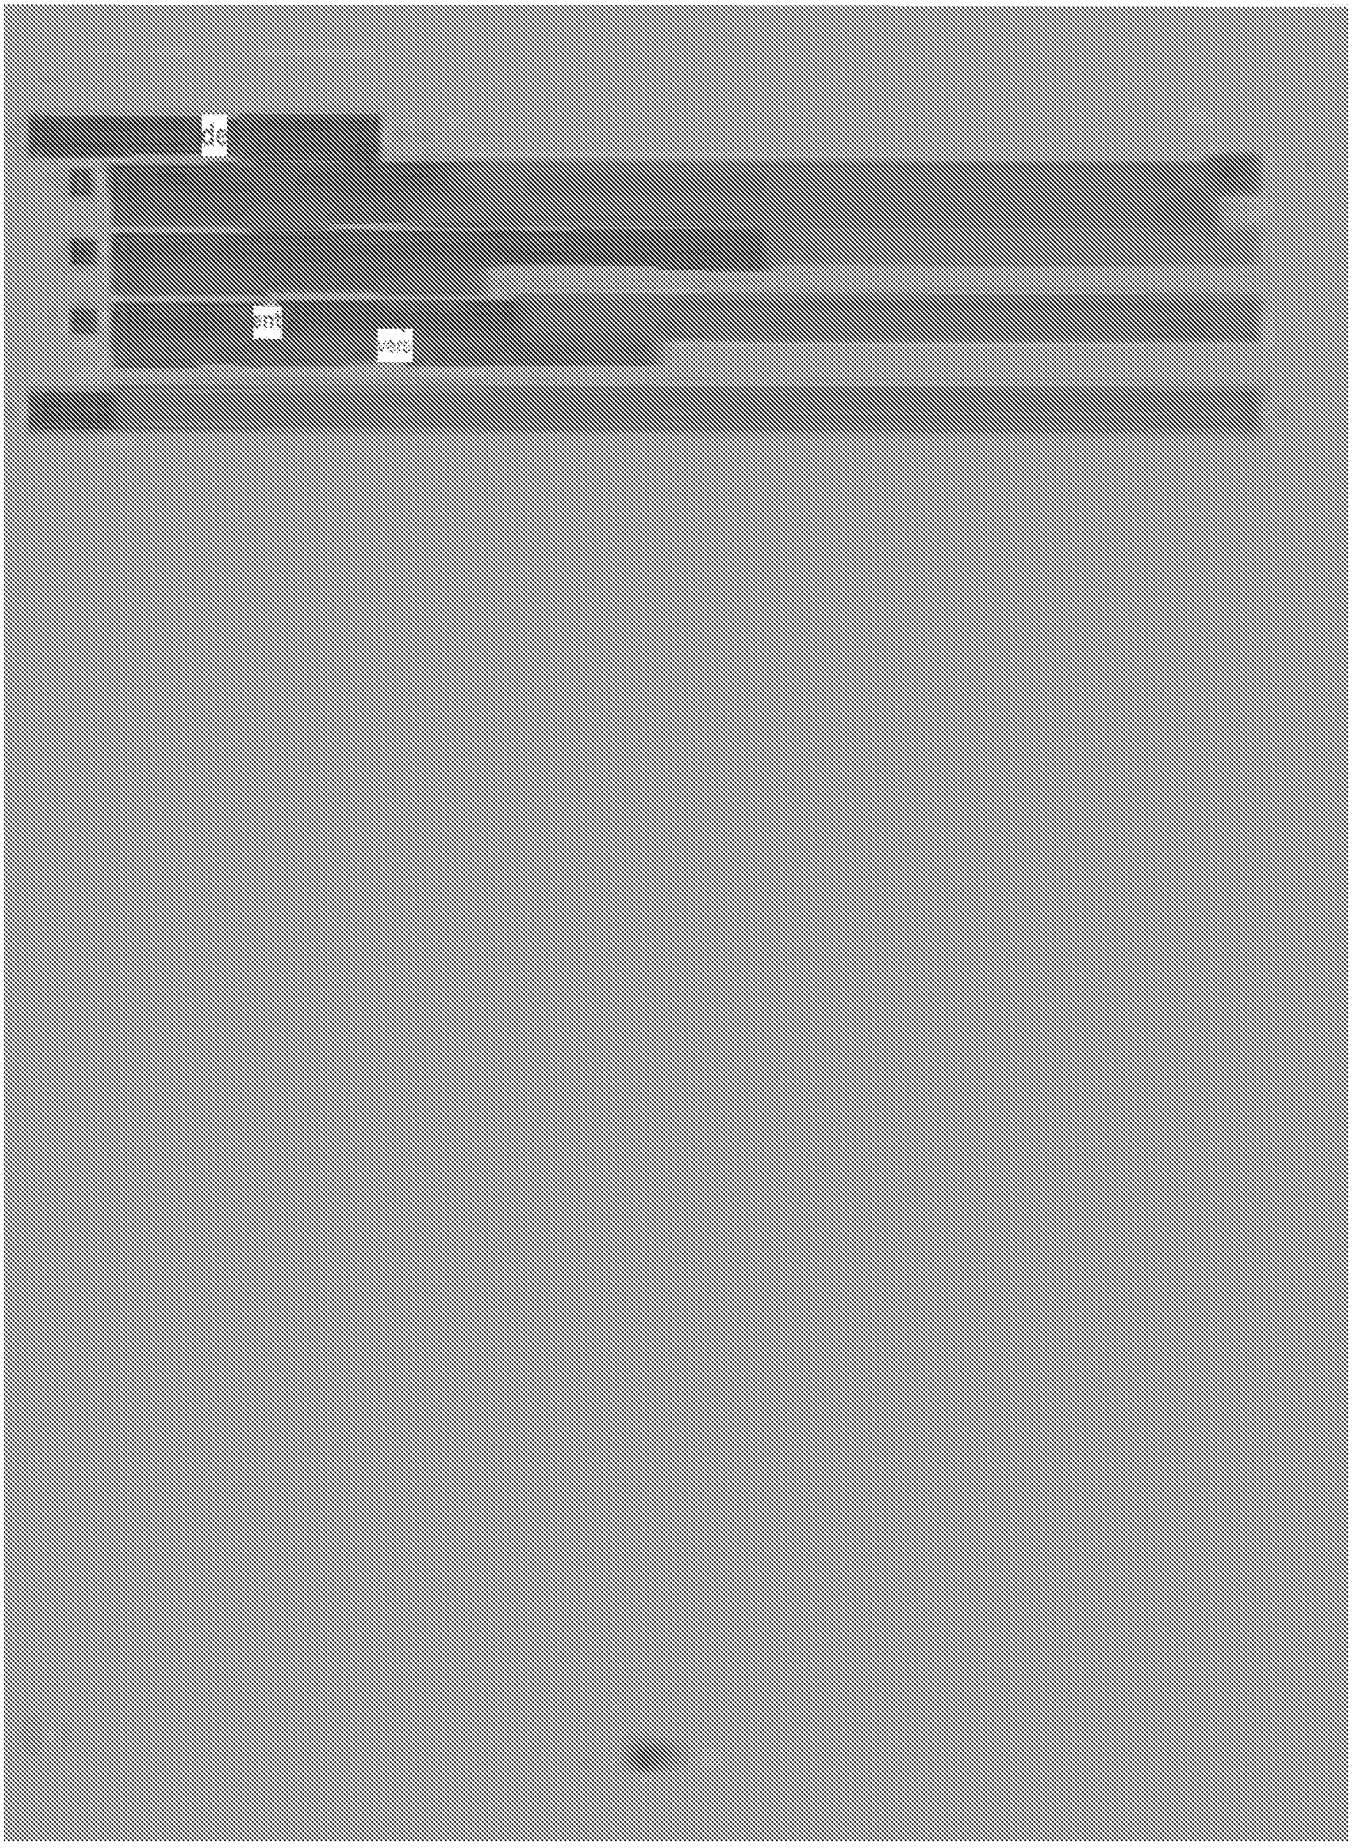
EXHIBIT 1040

Excerpt of *ACM SIGCOMM 2000 Conference: Applications, Technologies, Architectures, and Protocols for Computer Communication*, Stockholm, Sweden, August 28, 2000 to September 1, 2000, obtained from University of Cincinnati Library








3d
3c
3b
3a
3f
3e

This material may be protected by Copyright law (Title 17 U.S. Code)

Practical Network Support for IP Traceback

Stefan Savage, David Wetherall, Anna Karlin and Tom Anderson

Department of Computer Science and Engineering
University of Washington
Seattle, WA, USA

Abstract

This paper describes a technique for tracing anonymous packet flooding attacks in the Internet back towards their source. This work is motivated by the increased frequency and sophistication of denial-of-service attacks and by the difficulty in tracing packets with incorrect, or "spoofed", source addresses. In this paper we describe a general purpose traceback mechanism based on probabilistic packet marking in the network. Our approach allows a victim to identify the network path(s) traversed by attack traffic without requiring interactive operational support from Internet Service Providers (ISPs). Moreover, this traceback can be performed "post-mortem" - after an attack has completed. We present an implementation of this technology that is incrementally deployable, (mostly) backwards compatible and can be efficiently implemented using conventional technology.

1. INTRODUCTION

Denial-of-service attacks consume the resources of a remote host or network, thereby denying or degrading service to legitimate users. Such attacks are among the hardest security problems to address because they are simple to implement, difficult to prevent, and very difficult to trace. In the last several years, Internet denial-of-service attacks have increased in frequency, severity and sophistication. Howard reports that between the years of 1989 and 1995, the number of such attacks reported to the Computer Emergency Response Team (CERT) increased by 50 percent per year [25]. More recently, a 1999 CSI/FBI survey reports that 32 percent of respondents detected denial-of-service attacks directed against their sites [16]. Even more worrying, recent reports indicate that attackers have developed tools to coordinate distributed attacks from many separate sites [14].

Unfortunately, mechanisms for dealing with denial-of-service have not advanced at the same pace. Most work in this area has focused on *tolerating* attacks by mitigating their effects on the victim [38, 2, 26, 29, 9]. This approach can provide an effective stop-gap measure, but does not eliminate the problem nor does it discourage attackers. The other option, and the focus of this paper, is to trace attacks back towards their origin - ideally stopping an attacker at the source.

A perfect solution to this problem is complicated by the potential use of indirection to "launder" the true causal origin of an attack. For example, an attack may consist of packets sent from many different slave machines, themselves under the control of a remote master machine. Such indirection may be achieved either explicitly (by compromising the individual slave hosts directly) or implicitly (by sending false requests to the slaves on behalf of the victim - a so-called *reflector*). More challenging still, the true origin and identity of the attacker can be similarly concealed through chains of false computer accounts, call forwarding, and so forth. Consequently, we regard a complete solution - particularly one able to address the forensic needs of law enforcement - as an open problem.

Instead, we address the more limited operational goal of simply identifying the machines that *directly* generate attack traffic and the network path this traffic subsequently follows. We call this the *traceback problem* and it is motivated by the operational need to control and contain attacks. In this setting, even incomplete or approximate information is valuable because the efficacy of measures such as packet filtering improve as they are applied further from the victim and closer to the source.

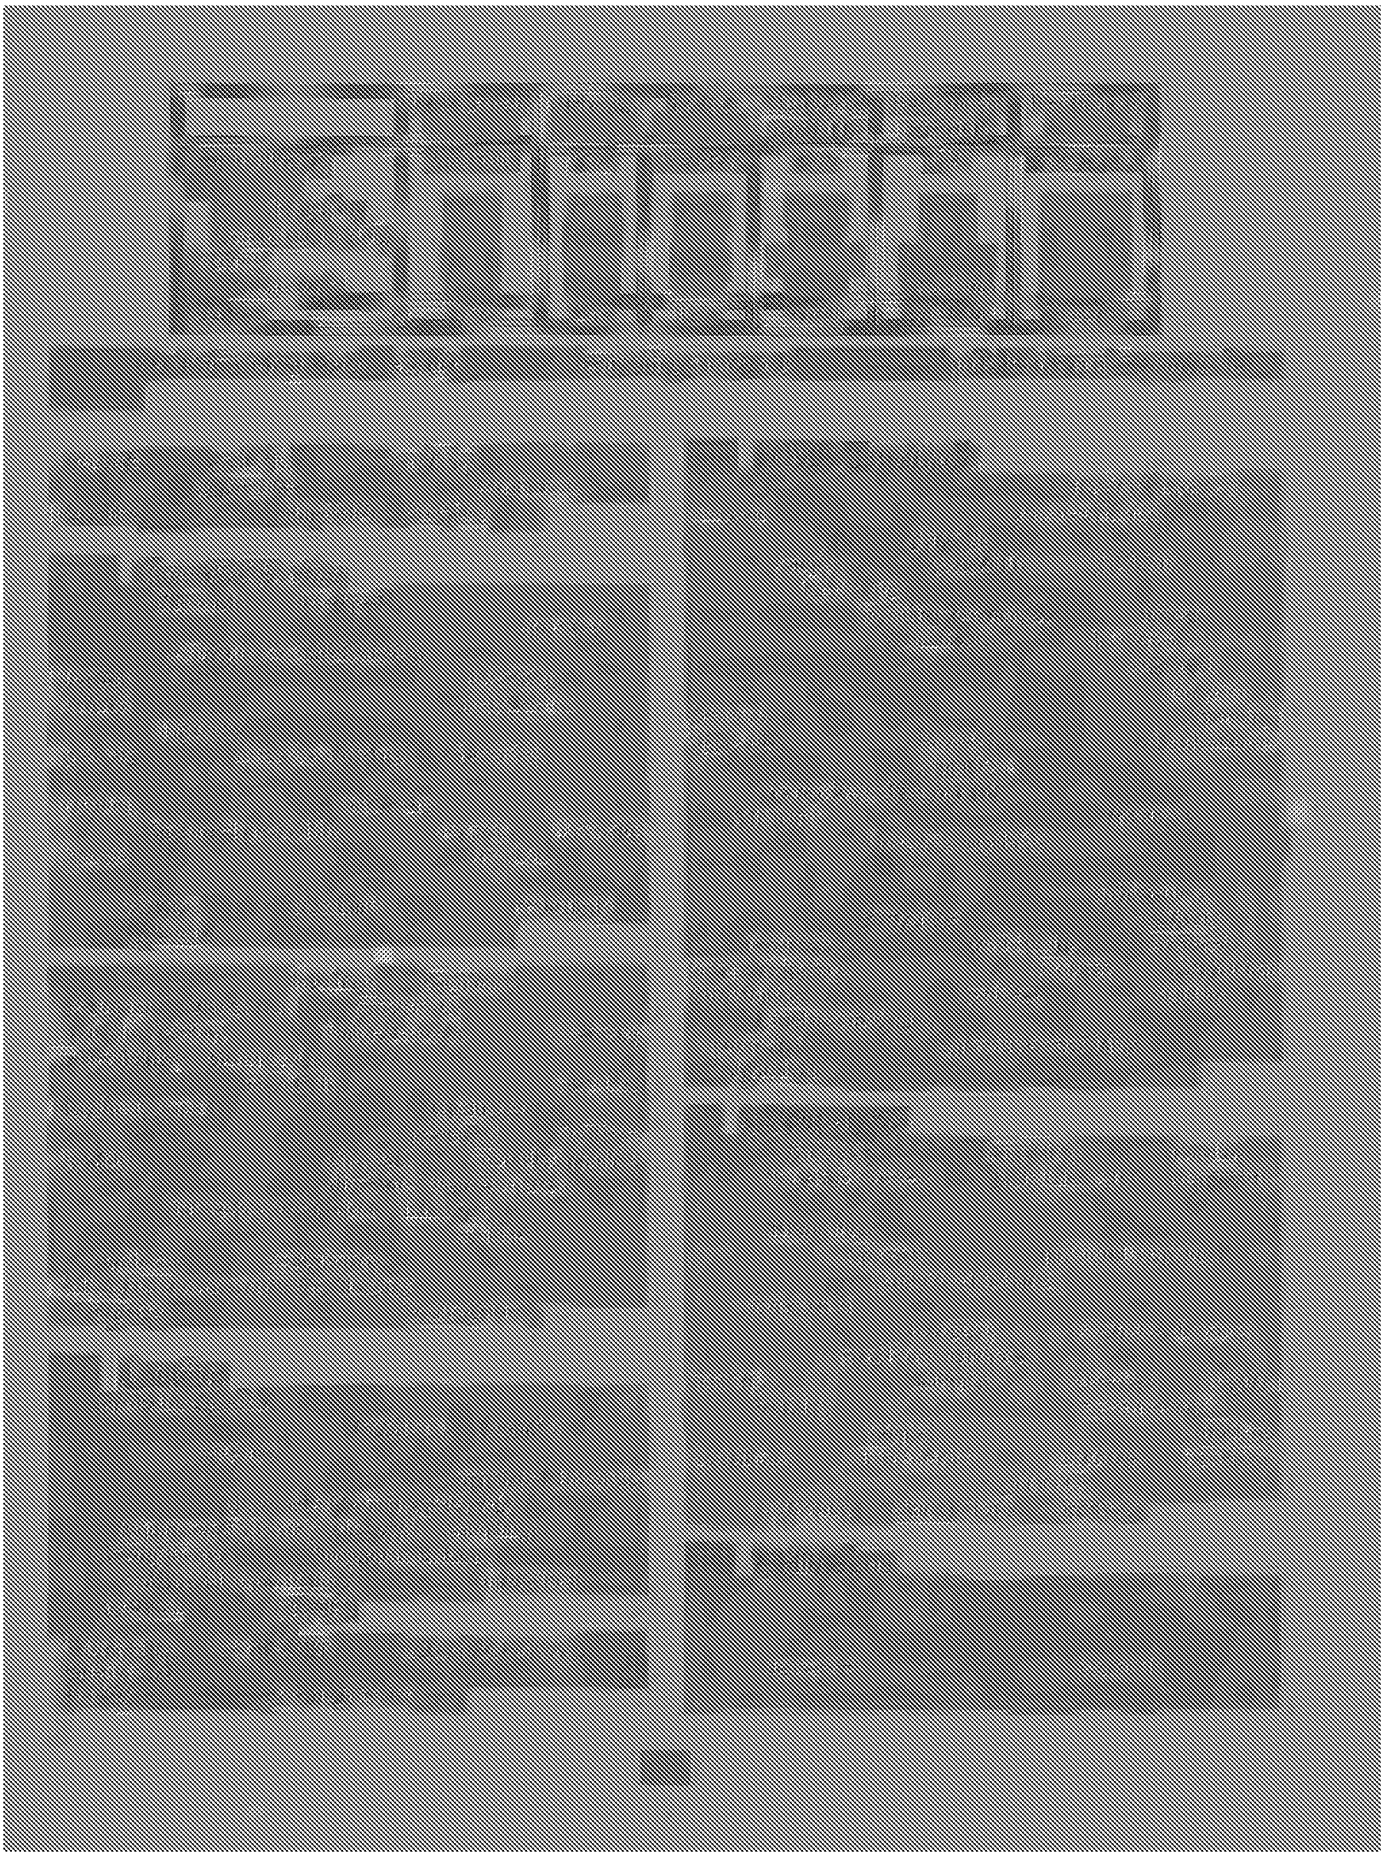
However, even for our restricted problem, determining the source generating attack traffic is surprisingly difficult due to the stateless nature of Internet routing. Attackers routinely disguise their location using incorrect, or "spoofed", IP source addresses. As these packets traverse the Internet their true origin is lost and a victim is left with little useful information. While there are several ad hoc traceback techniques in use, they all have significant drawbacks that limit their practical utility in the current Internet.

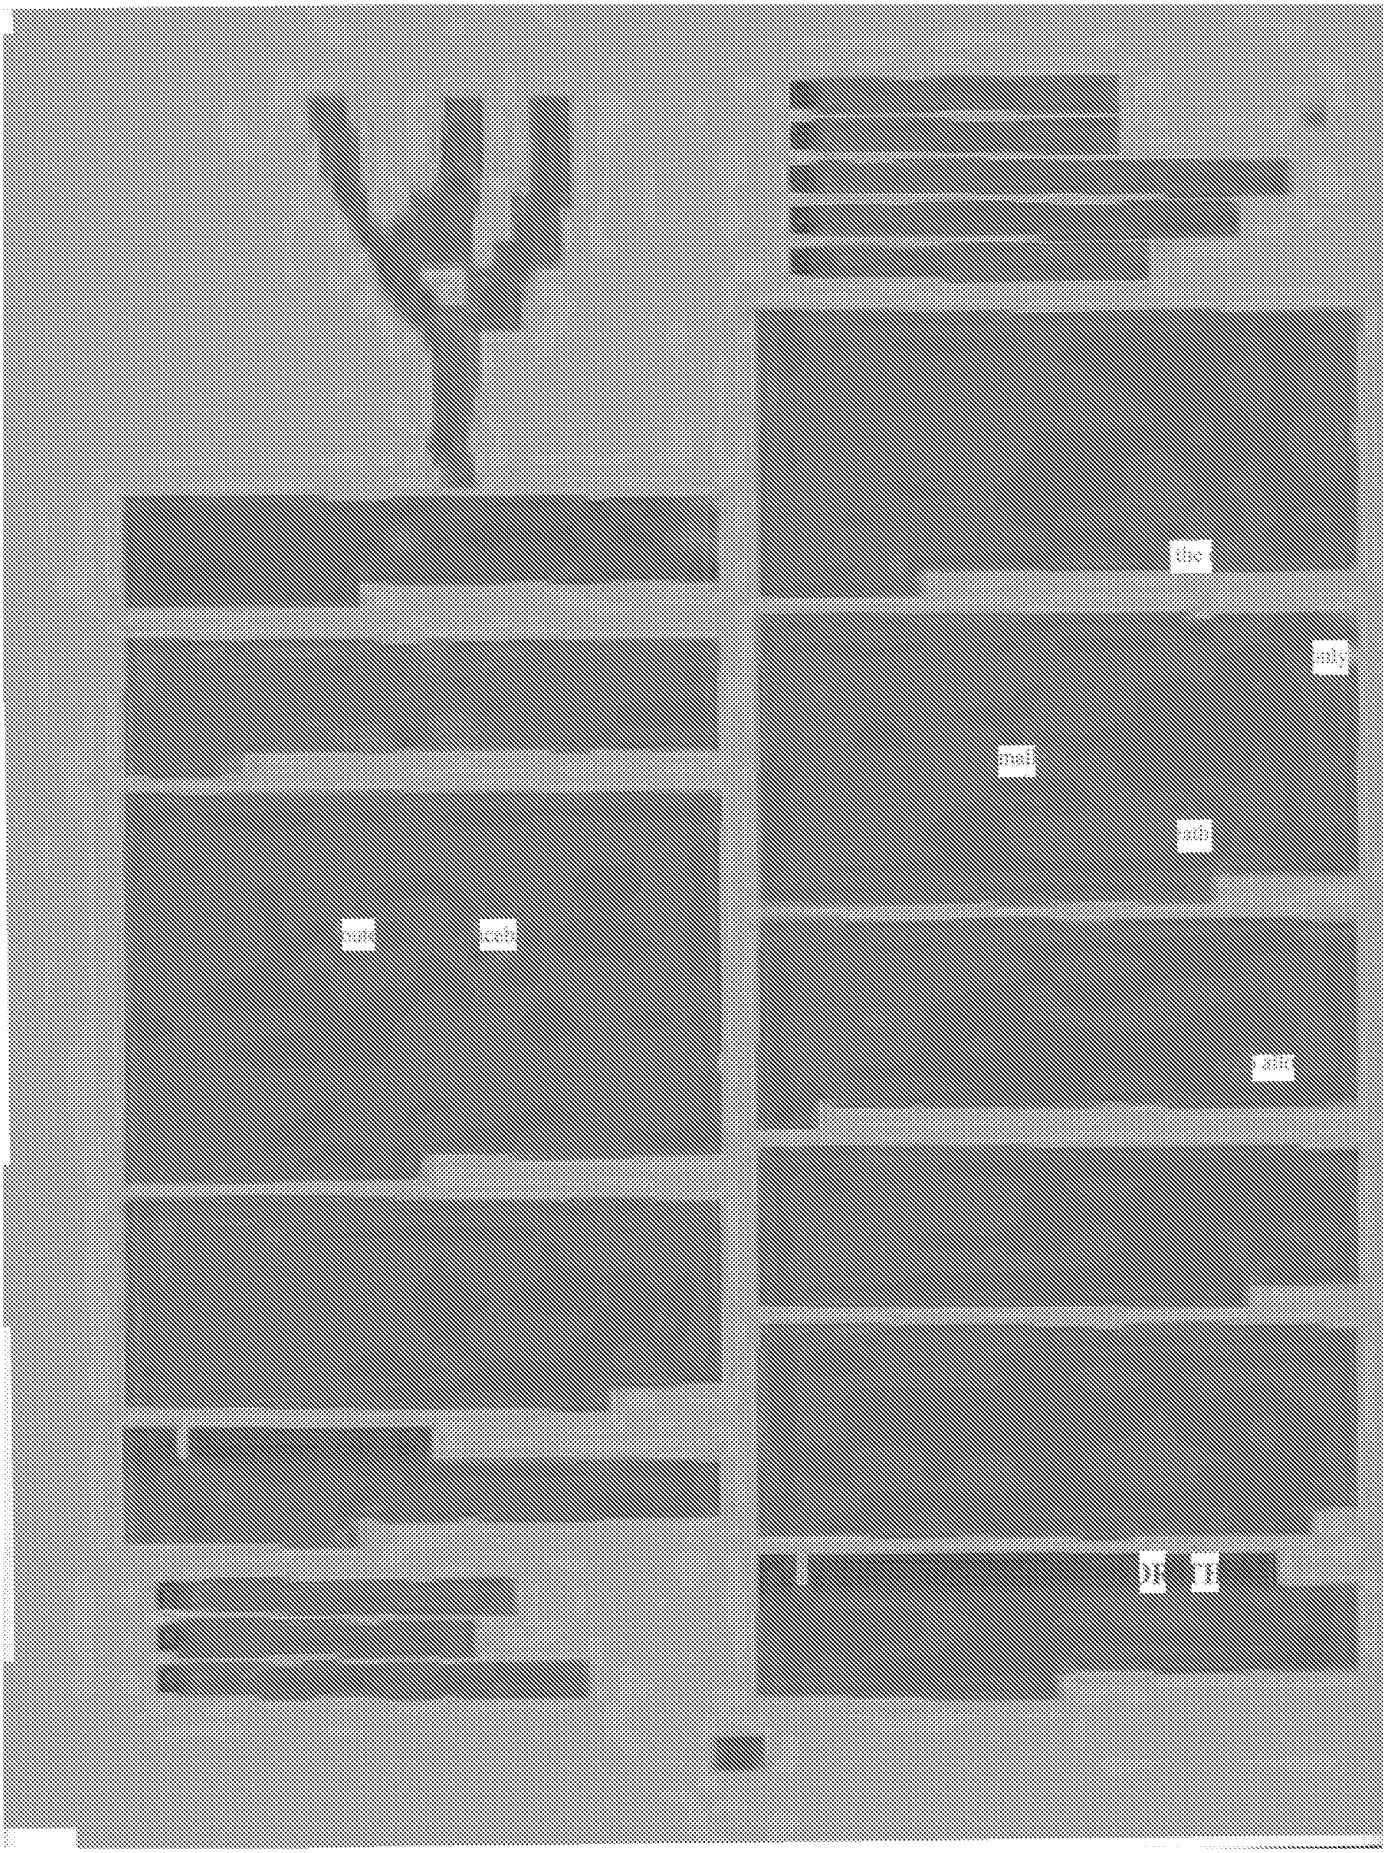
In this paper we present a new approach to the traceback problem that addresses the needs of both victims and network operators. Our solution is to probabilistically mark packets with partial path information as they arrive at routers. This approach exploits the observation that attacks generally comprise large numbers of packets. While each marked packet represents only a "sample" of the path it has traversed, by combining a modest number of such packets a victim can reconstruct the entire path. This allows victims to locate the approximate source of attack traffic without requiring the assistance of outside network operators. Moreover, this determination can be made even after an attack has completed. Both facets of our solution represent substantial improvements over existing capabilities for dealing with flooding-style denial-of-service attacks.

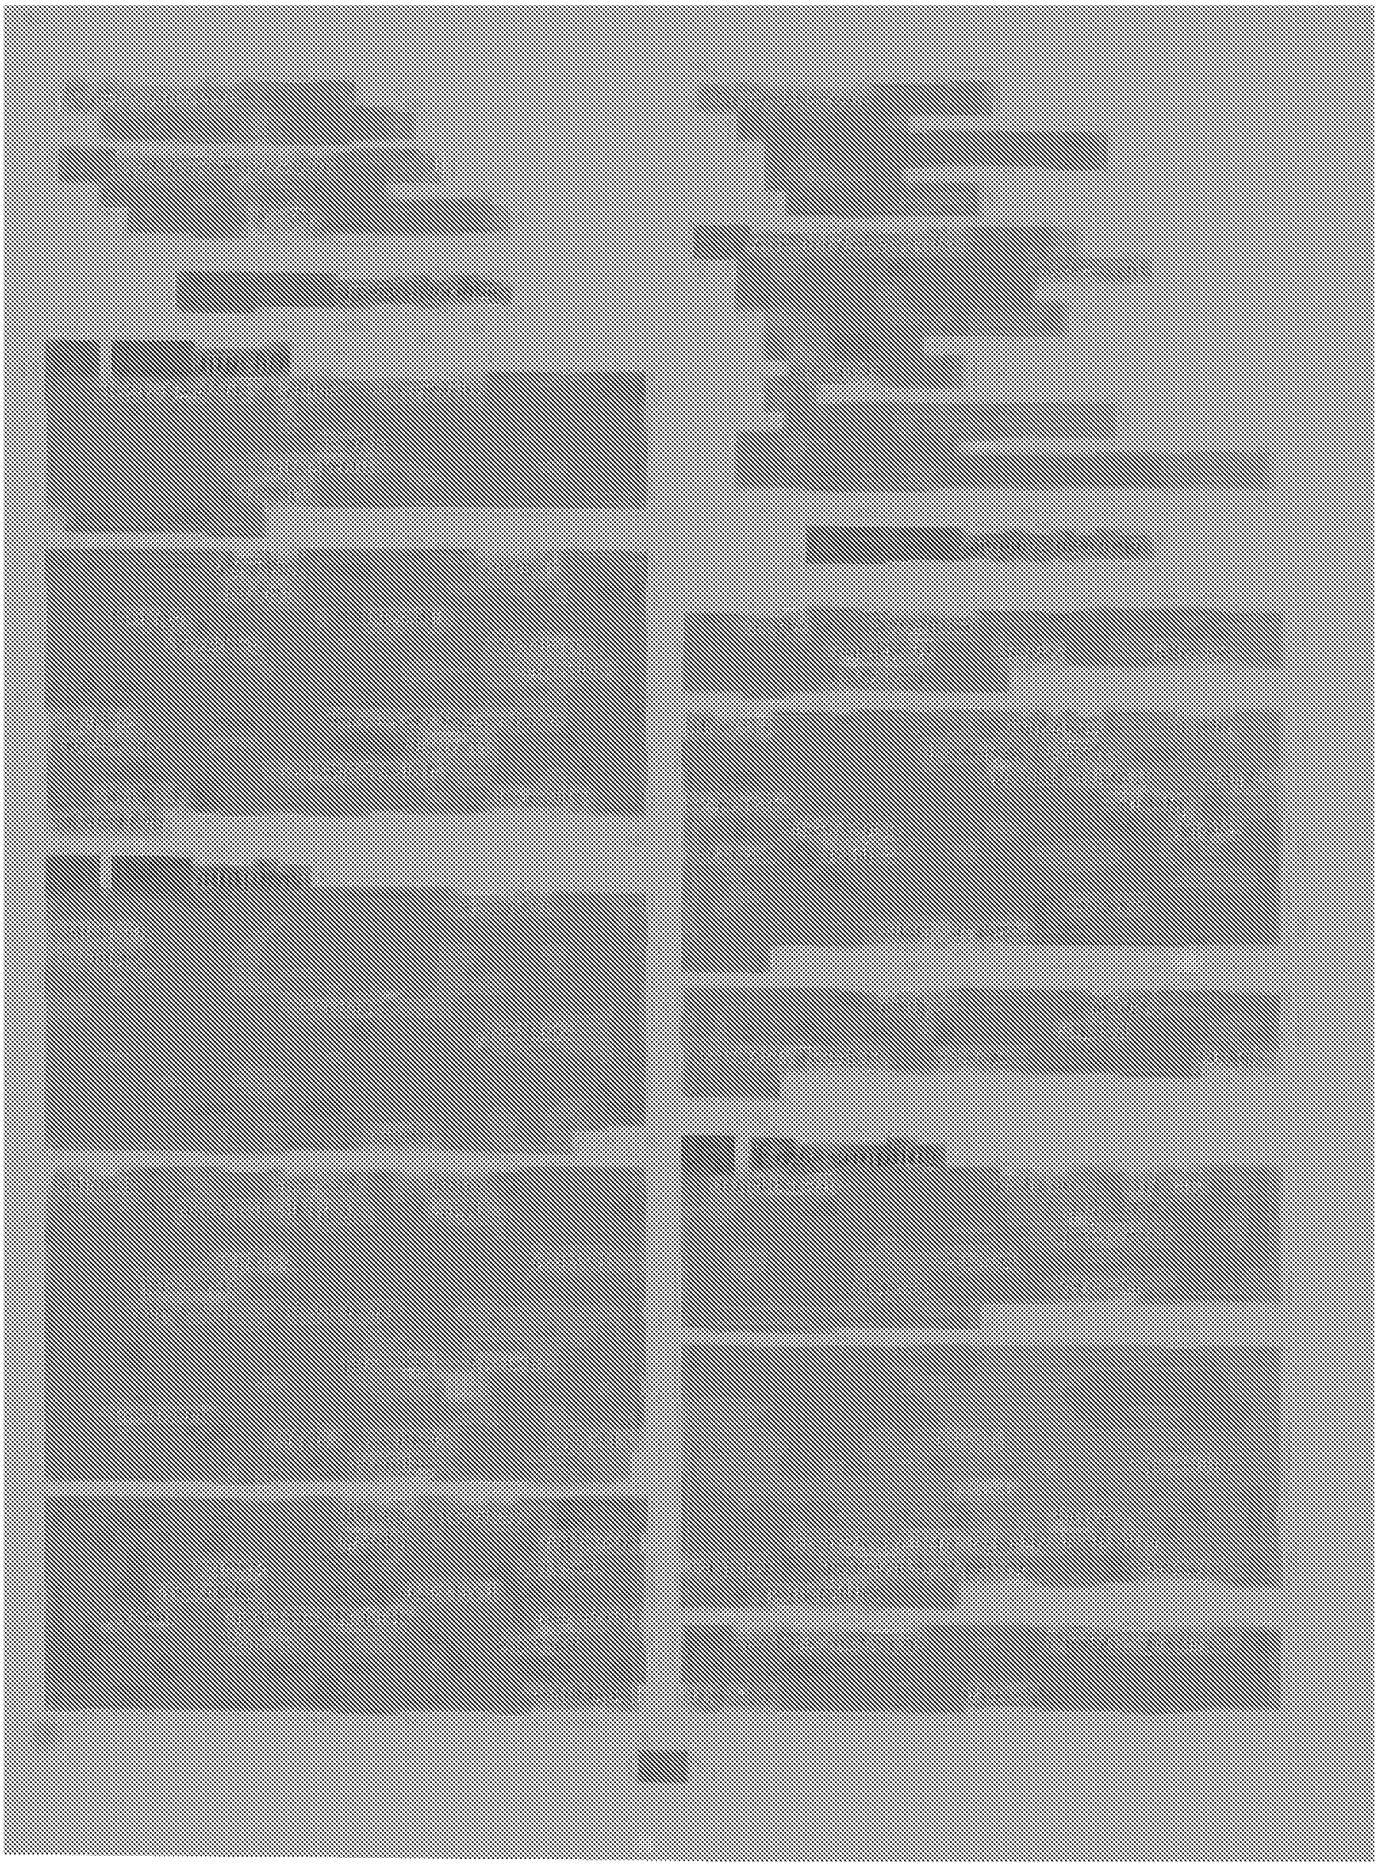
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM'00, Stockholm, Sweden.
Copyright 2000 ACM 1-58113-224-7/00/0008...$5.00.

Cloudflare, Inc. Exhibit 1003 Page 1420

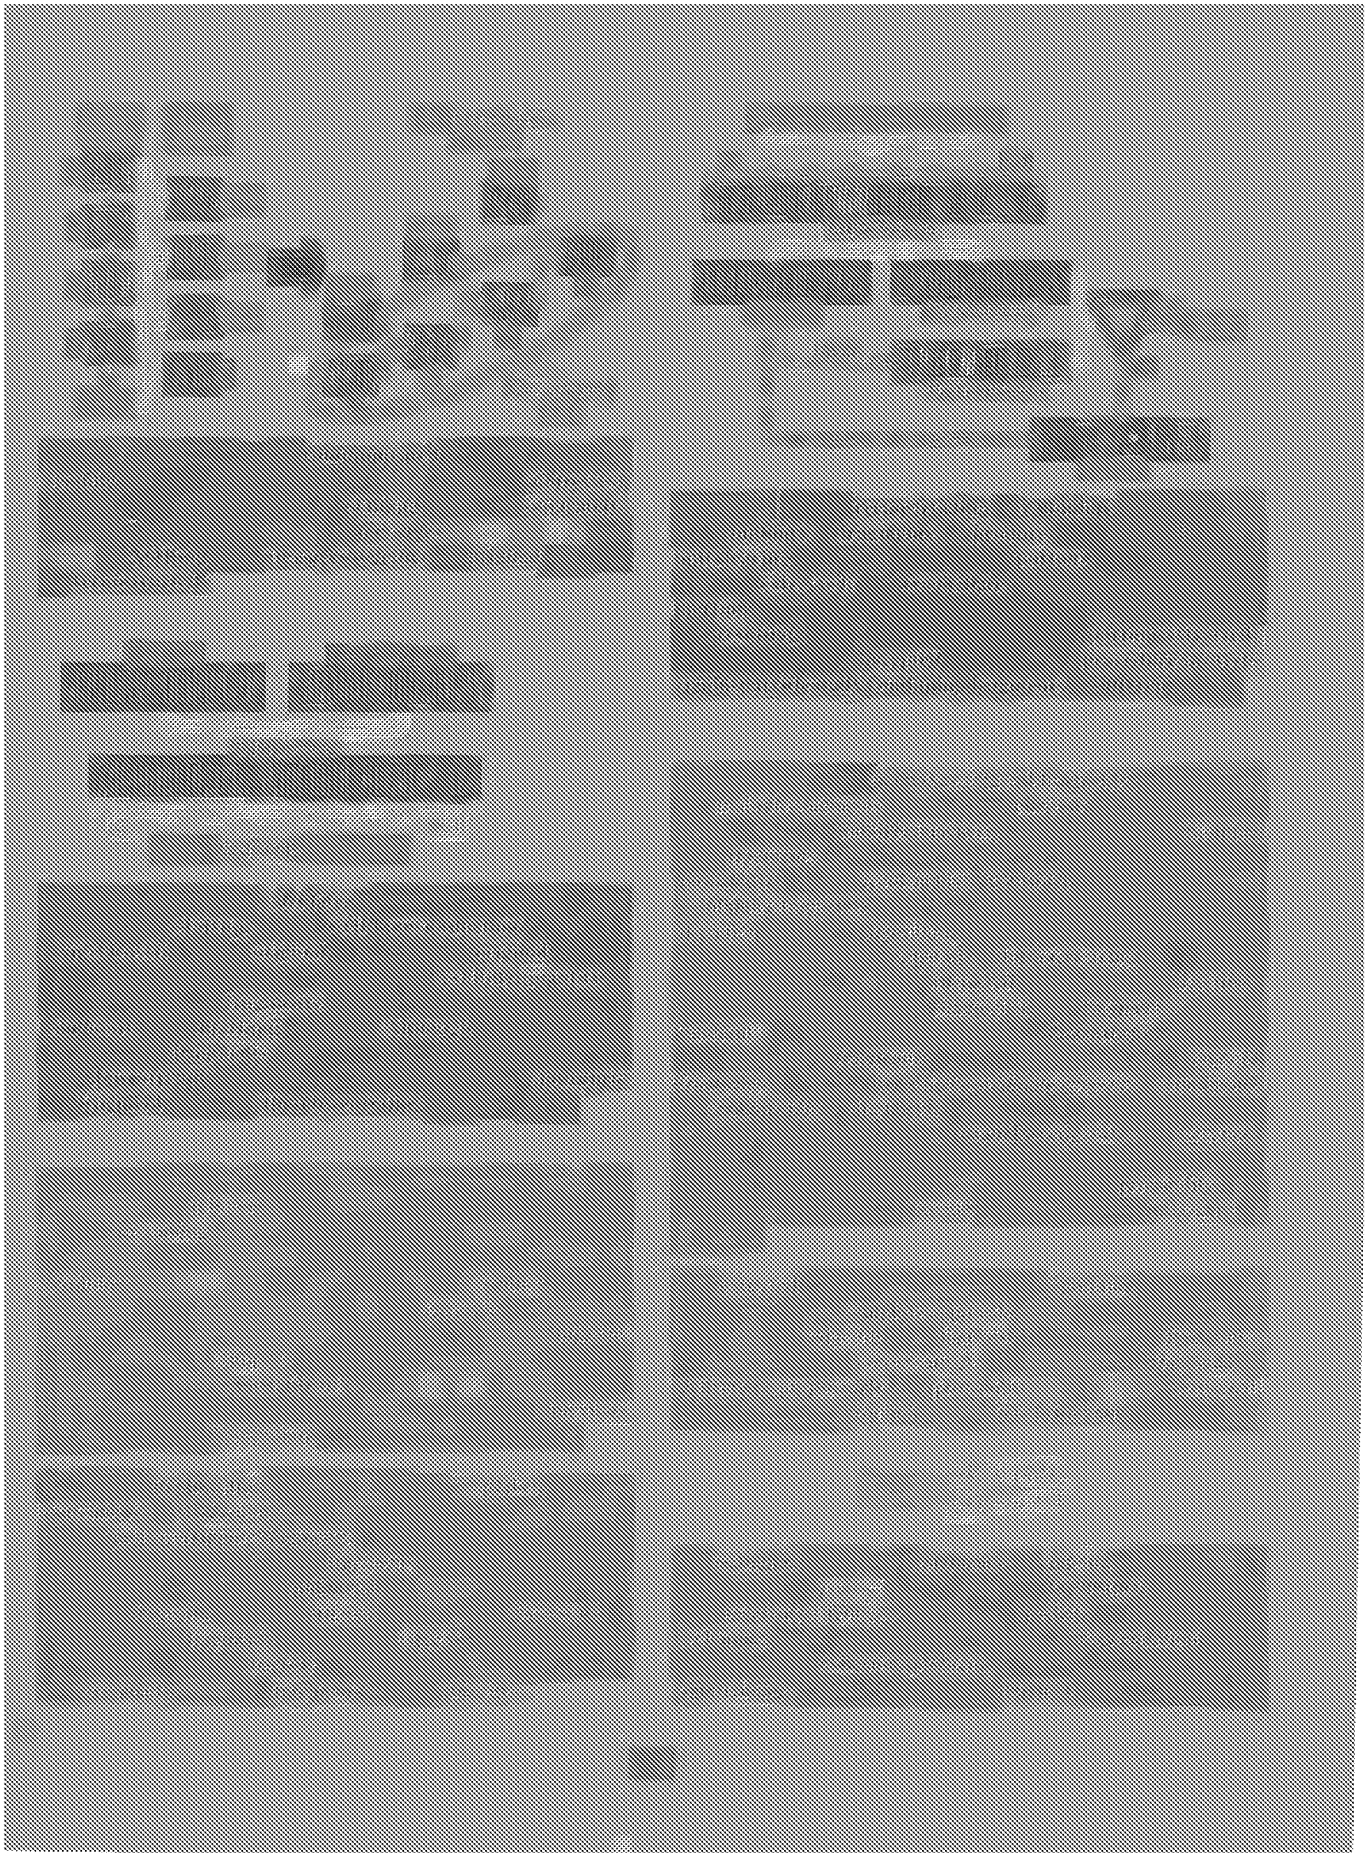


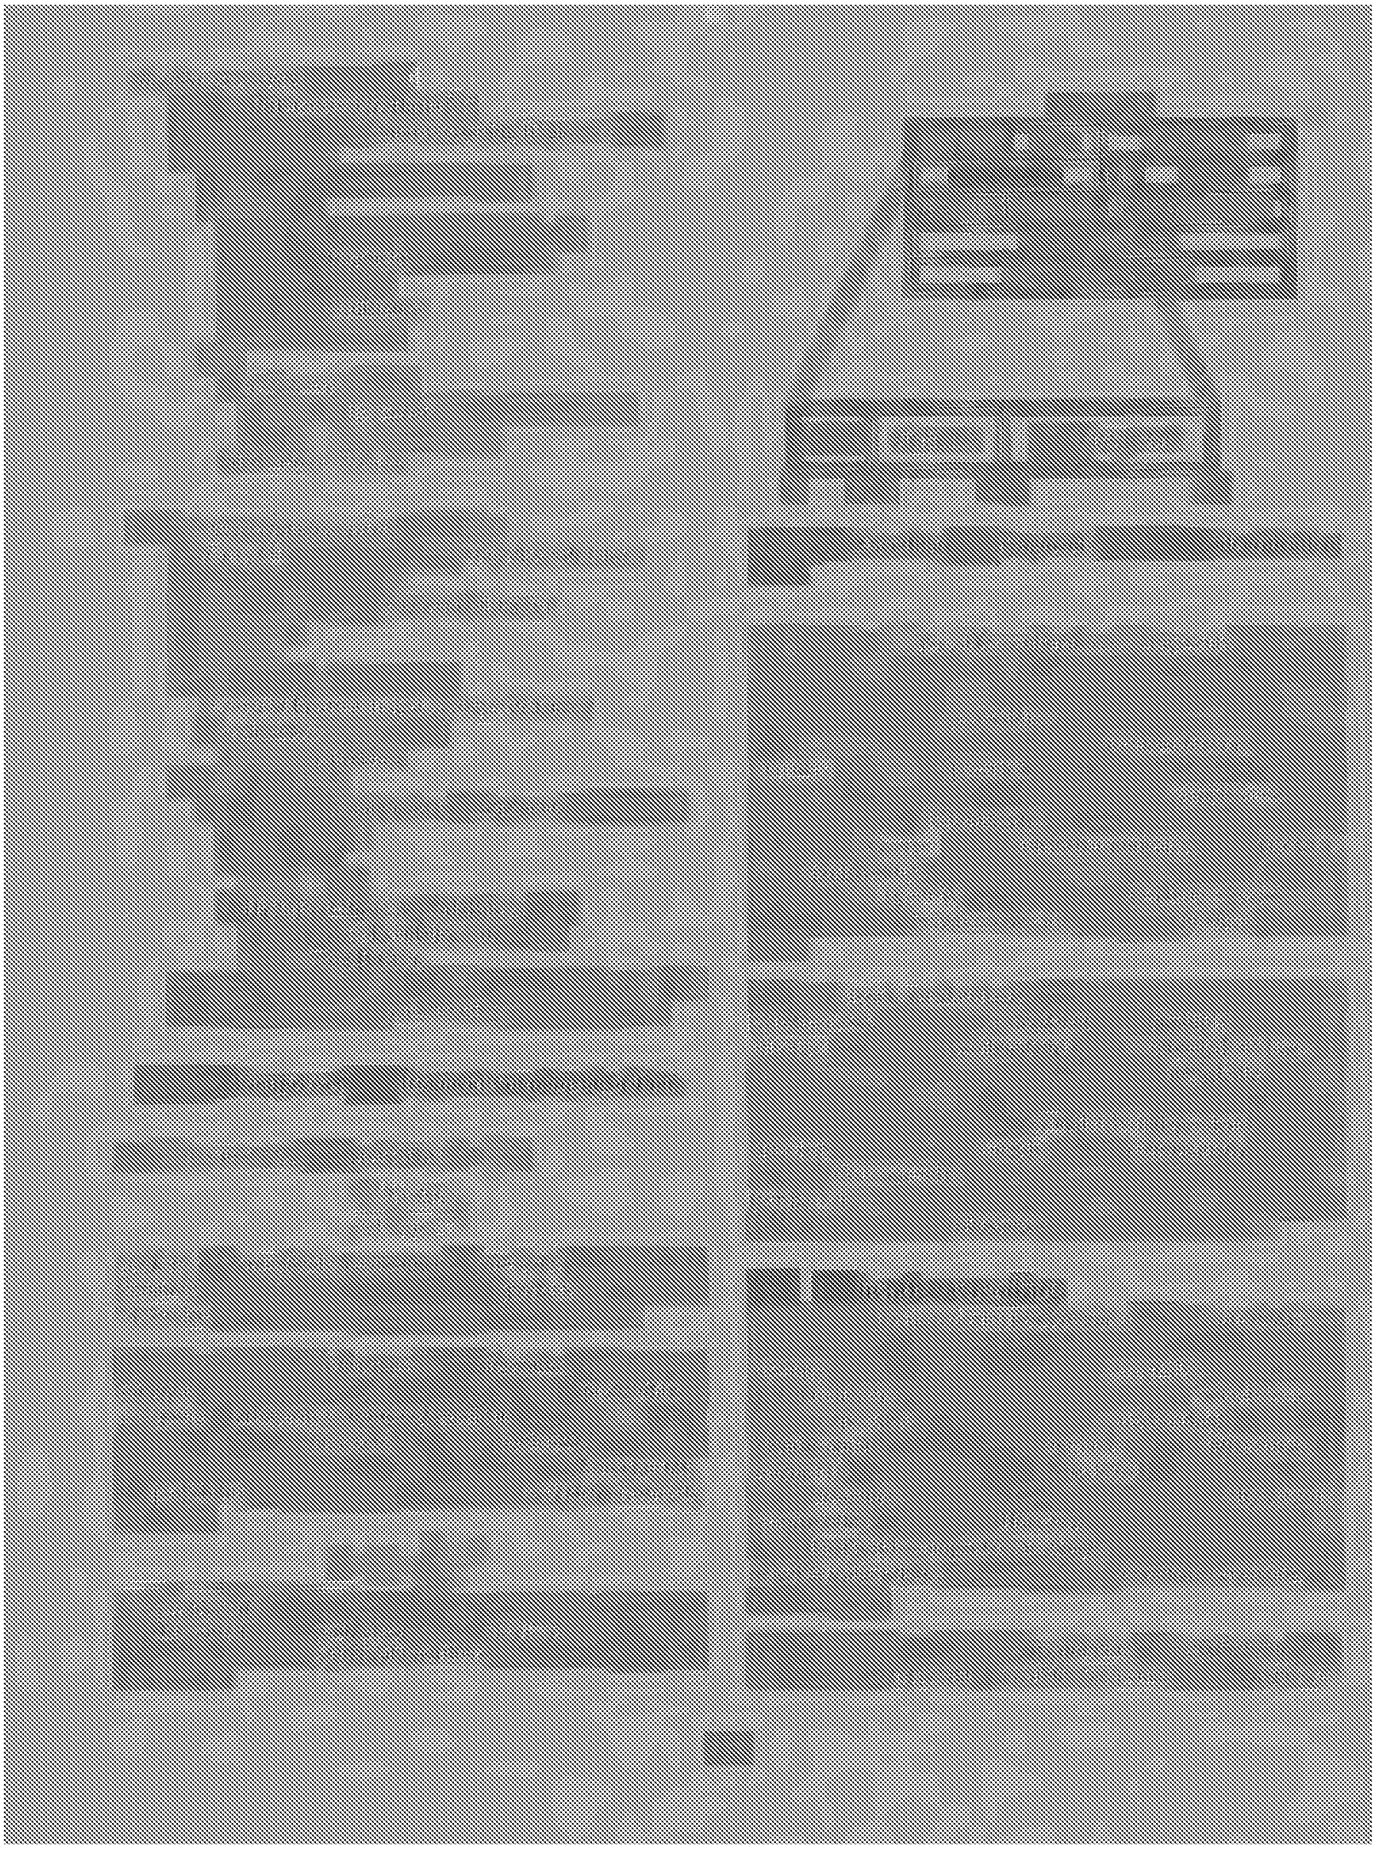


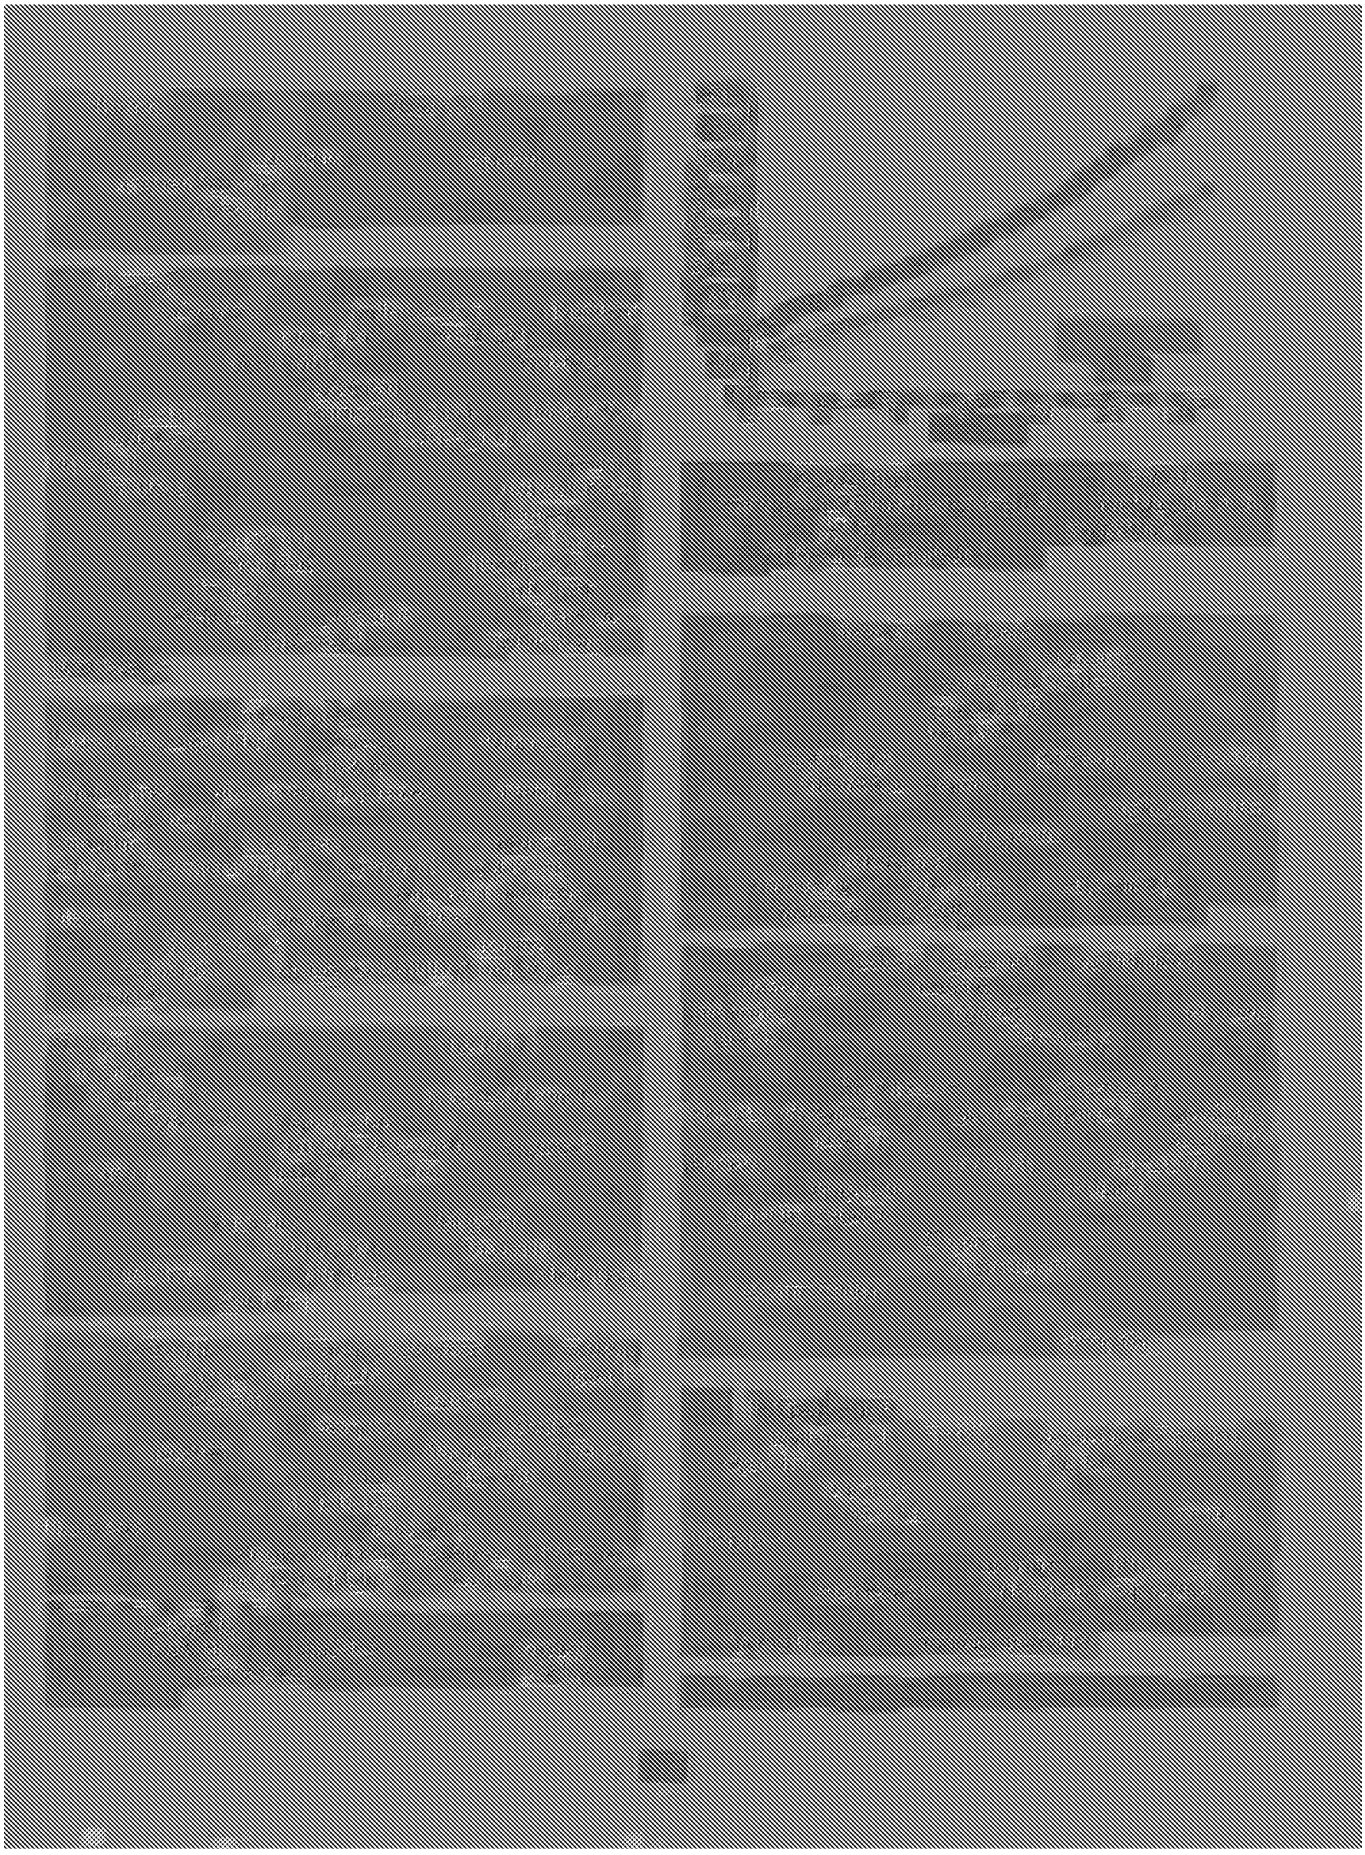






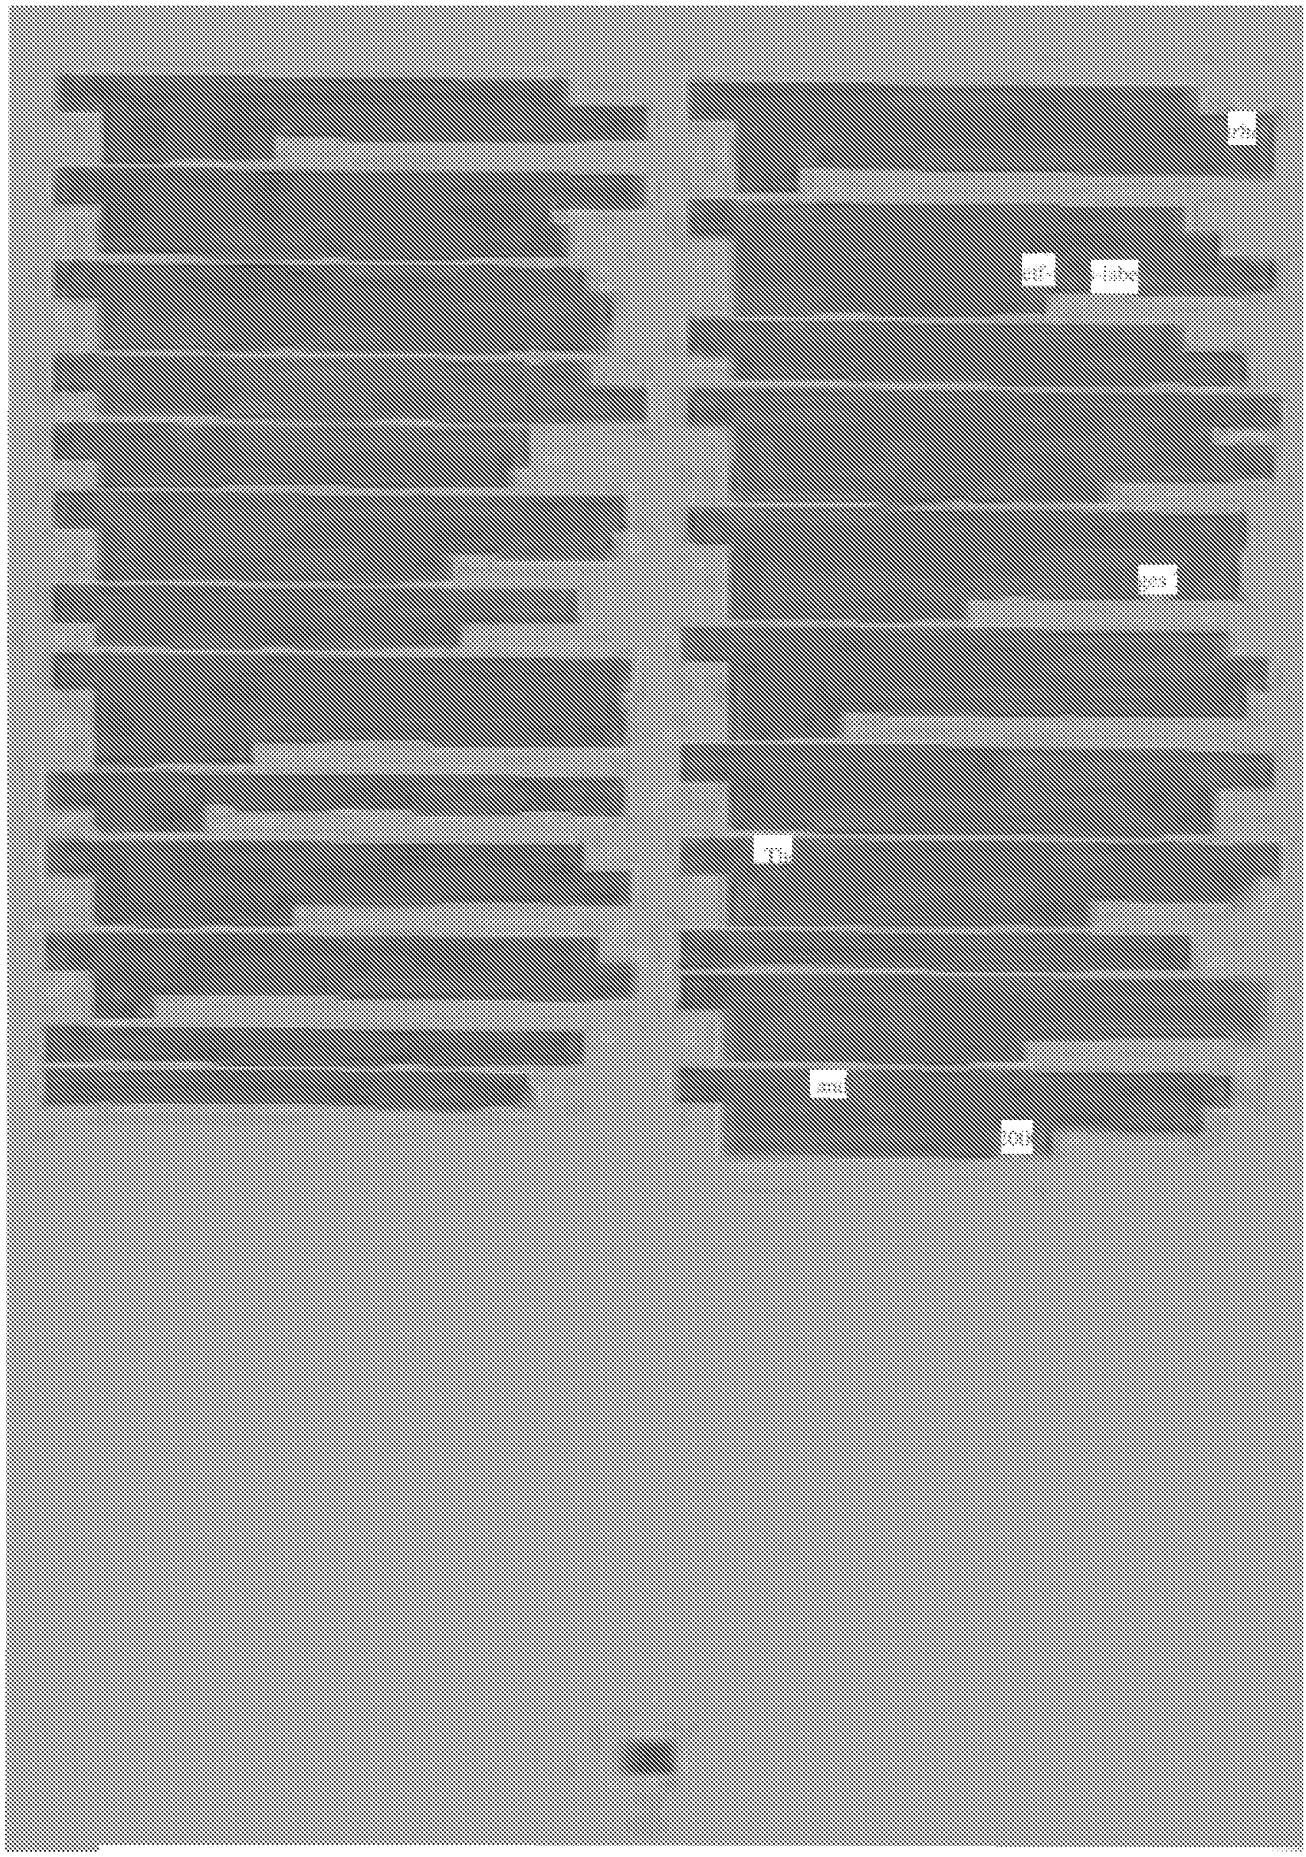












Declaration of Dr. Hall-Ellis

ATTACHMENT 1A

CINCINNATI

Library Website | Ask the Libraries

- Catalog
- Libraries
- Reserves
- OhioLINK
- Off Campus Access
- My Library Record
- Help

Click & Collect - Items with the status of "Click and Collect" are available for request. Items from one library location cannot be requested for pickup at another library location. For details visit https://libraries.uc.edu/about/covid-19.html

Start Over | Previous Display | Return to List | Limit/Sort

(Search History)

```
LEADER 00000cam  2200000Ia 4500
001    45172567
003    OCoLC
005    20020808172642.0
008    001014s2000    nyu           000 0 eng d
019    47191367
020    1581132247
040    OGE|cOGE|dOCL|dCUS|dCAI|dCIN
111 2  ACM SIGCOMM Conference|d(2000 :|cStockholm, Sweden)
245 10 ACM SIGCOMM 2000 Conference :|bapplications, technologies,
       architectures, and protocols for computer communications,
       Stockholm, Sweden, August 28 to September 1, 2000
246 1  |iAlso known as:|aSIGCOMM '00
246 14 Proceedings of ACM SIGCOMM 2000 Conference
246 30 Applications, technologies, architectures, and protocols
       for computer communications
246 30 SIGCOMM 2000 Conference
260    New York :|bACM,|cc2000
300    xii, 348 p. :|bill. ;|c28 cm
440  0 Computer communication review,|x0146-4833 ;|vv. 30, no. 4
500    "October 2000."
504    Includes bibliographical references and index
530    Also available on the World Wide Web via ACM Digital
       Library with title: Proceedings of the Conference on
       SIGCOMM 2000
650  0 Computer networks|vCongresses
650  0 Computer network protocols|vCongresses
650  0 Data transmission systems|vCongresses
710 2  Association for Computing Machinery.|bSpecial Interest
       Group on Data Communications
740 0  Proceedings of the Conference on SIGCOMM 2000
910    MARS
```

LOCATION	CALL NO.	STATUS
CEAS Journals	TK5105.5 .C63694 v.30(2000)	LIB USE ONLY

Cloudflare, Inc. Exhibit 1003 Page 1434

Declaration of Dr. Hall-Ellis

ATTACHMENT 1B

OCLC 45172567 **No holdings in XXX - 25 other holdings**

Books

						Rec Stat	c	**Entered**	20001014	**Replaced**	20190311232322.9
Type	a	**ELvl**	I	**Srce**	d	**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	1	**Biog**		**MRec**		**Ctry**	nyu
		Cont	b	**GPub**		**LitF**	0	**Indx**	1		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	2000		

040 OGE $b eng $c OGE $d OCL $d CUS $d CAI $d OCLCQ $d BAKER $d YDXCP $d ZWZ $d OCLCQ $d OCLCF $d OCLCQ $d OCLCO $d OCL $d OCLCO $d OCLCQ
019 47191367
020 1581132247
020 9781581132243
050 4 TK5105.5 $b .A26 2000
050 4 QA1 $b .A85458 v.30, no. 4
082 04 004.6
090 $b
111 2 ACM SIGCOMM Conference $d (2000 : $c Stockholm, Sweden)
245 10 ACM SIGCOMM 2000 Conference : $b applications, technologies, architectures, and protocols for computer communications, Stockholm, Sweden, August 28 to September 1, 2000.
246 30 Applications, technologies, architectures, and protocols for computer communications
246 14 Proceedings of ACM SIGCOMM 2000 Conference
246 30 SIGCOMM 2000 Conference
246 1 $i Also known as: $a SIGCOMM '00
260 New York : $b ACM, $c ©2000.
300 xii, 348 pages : $b illustrations ; $c 28 cm.
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
490 1 Computer communication review, $x 0146-4833 ; $v v. 30, no. 4
500 "October 2000."
504 Includes bibliographical references and index.
530 Also available on the World Wide Web via ACM Digital Library with title: Proceedings of the Conference on SIGCOMM 2000.
650 0 Computer networks $v Congresses.
650 0 Computer network protocols $v Congresses.
650 0 Data transmission systems $v Congresses.
650 7 Computer network protocols. $2 fast $0 (OCoLC)fst00872279
650 7 Computer networks. $2 fast $0 (OCoLC)fst00872297
650 7 Data transmission systems. $2 fast $0 (OCoLC)fst00887993
650 7 Réseaux d'ordinateurs $x Congrès. $2 ram
650 7 Protocoles de réseaux d'ordinateurs $x Congrès. $2 ram
655 7 Conference papers and proceedings. $2 fast $0 (OCoLC)fst01423772
710 2 Association for Computing Machinery. $b Special Interest Group on Data Communications.
740 0 Proceedings of the Conference on SIGCOMM 2000.

830 0 Computer communication review ; $v v. 30, no. 4. $x 0146-4833
856 41 $z Full text of the Conference proceedings (2000) $u
http://www.acm.org/pubs/contents/proceedings/comm/347059/
856 4 $u http://www.acm.org/pubs/contents/proceedings/series/comm $z Accessible en INTRANET

Declaration of Dr. Hall-Ellis

ATTACHMENT 1C

LIBRARY OF CONGRESS AUTHORITIES

Library buildings are closed to the public until further notice, but LC Authorities is available. More.

LC control no.: sh2008101222
LCCN Permalink: https://lccn.loc.gov/sh2008101222
HEADING: Computer networks Congresses

000 00490cz a2200169n 450
001 7437752
005 20110729160105.0
008 080206|| anannbabn |n ana
010 __ |**a** sh2008101222
035 __ |**a** (DLC)376758
035 __ |**a** (DLC)sh2008101222
040 __ |**a** DLC |**b** eng |**c** DLC
150 __ |**a** Computer networks |**v** Congresses
667 __ |**a** Record generated for validation purposes.
670 __ |**a** Work cat.: Combinatorial and algorithmic aspects of networking, c2005
906 __ |**t** 8888 |**u** te00 |**v** 0
953 __ |**a** te00

Previous Next

Save, Print or Email Records (View Help)

Select Download Format: Text (Labelled Display) Press to SAVE or PRINT

Email Text (Labelled Display) to: Press to SEND EMAIL

Library of Congress
URL: https://www.loc.gov/

Mailing Address:
101 Independence Ave, S.E.
Washington, DC 20540

Library of Congress Authorities
URL: http://authorities.loc.gov/
Library of Congress Online Catalog
URL: https://catalog.loc.gov/

Questions, comments, error reports: Contact Us

Declaration of Dr. Hall-Ellis

ATTACHMENT 1D

LIBRARY OF CONGRESS AUTHORITIES

Library buildings are closed to the public until further notice, but LC Authorities is available. More.

LC control no.: sh2008101491
LCCN Permalink: https://lccn.loc.gov/sh2008101491
HEADING: Computer network protocols Congresses

000 00511cz a2200169n 450
001 7438020
005 20110729160116.0
008 080206|| anannbabn |n ana
010 __ |**a** sh2008101491
035 __ |**a** (DLC)377027
035 __ |**a** (DLC)sh2008101491
040 __ |**a** DLC |**b** eng |**c** DLC
150 __ |**a** Computer network protocols |**v** Congresses
667 __ |**a** Record generated for validation purposes.
670 __ |**a** Work cat.: Communications and multimedia security issues of the new century, 2001
906 __ |**t** 8888 |**u** te00 |**v** 0
953 __ |**a** te00

Previous Next

Save, Print or Email Records (View Help)

Select Download Format: Text (Labelled Display) Press to SAVE or PRINT

Email Text (Labelled Display) to: Press to SEND EMAIL

Library of Congress
URL: https://www.loc.gov/

Mailing Address:
101 Independence Ave, S.E.
Washington, DC 20540

Library of Congress Authorities
URL: http://authorities.loc.gov/
Library of Congress Online Catalog
URL: https://catalog.loc.gov/

Questions, comments, error reports: Contact Us

Declaration of Dr. Hall-Ellis

ATTACHMENT 1E

LIBRARY OF CONGRESS AUTHORITIES

Library buildings are closed to the public until further notice, but LC Authorities is available. More.

LC control no.: sh2008102044
LCCN Permalink: https://lccn.loc.gov/sh2008102044
HEADING: Data transmission systems Congresses

000 00485cz a2200169n 450
001 7439990
005 20110729160139.0
008 080207|| anannbabn |n ana
010 __ |**a** sh2008102044
035 __ |**a** (DLC)377610
035 __ |**a** (DLC)sh2008102044
040 __ |**a** DLC |**b** eng |**c** DLC
150 __ |**a** Data transmission systems |**v** Congresses
667 __ |**a** Record generated for validation purposes.
670 __ |**a** Work cat.: IEEE Intelligent Network 2001 Workshop, c2001
906 __ |**t** 8888 |**u** te00 |**v** 0
953 __ |**a** te00

Previous Next

Save, Print or Email Records (View Help)

Select Download Format: Text (Labelled Display) Press to SAVE or PRINT

Email Text (Labelled Display) to: Press to SEND EMAIL

Library of Congress
URL: https://www.loc.gov/

Mailing Address:
101 Independence Ave, S.E.
Washington, DC 20540

Library of Congress Authorities
URL: http://authorities.loc.gov/
Library of Congress Online Catalog
URL: https://catalog.loc.gov/

Questions, comments, error reports: Contact Us

Declaration of Dr. Hall-Ellis

EXHIBIT 1041

K. Park et al., "A Proactive Approach to Distributed DoS Attack Prevention Using Route-Based Packet Filtering," Department of Computer Sciences, Purdue University, Dec. 3, 2000

A Proactive Approach to Distributed DoS Attack Prevention using Route-Based Packet Filtering*

Kihong Park† Heejo Lee‡
Network Systems Lab
Department of Computer Sciences
Purdue University
West Lafayette, IN 47907
{park,hlee}@cs.purdue.edu

CSD-TR 00-017 and
CERIAS-TR 2001-40
December 3, 2000

Abstract

Denial-of-service (DoS) attack on the Internet has become a pressing problem. In this paper, we describe, analyze and evaluate route-based distributed packet filtering (DPF), a novel approach to distributed DoS (DDoS) attack prevention. We show that DPF achieves proactiveness and scalability, and we show that there is an intimate relationship between the effectiveness of DPF at mitigating DDoS attack and power-law network topology.

The salient features of this work are two-fold. First, we show that DPF is able to proactively filter out a significant fraction of spoofed packet flows and prevent attack packets from reaching their targets in the first place. The IP flows that cannot be proactively curtailed are extremely sparse such that their origin can be localized—i.e., IP traceback—to within a small, constant number of candidate sites. We show that the two proactive and reactive performance effects can be achieved by implementing route-based filtering on less than 20% of Internet autonomous system (AS) sites. Second, we show that the two complementary performance measures are dependent on the properties of the underlying AS graph topology. In particular, we show that the power-law structure of Internet AS topology leads to connectivity properties which are crucial in facilitating the observed performance effects.

As a DDoS prevention architecture, DPF is able to emulate the IP traceback prowess of probabilistic packet marking, while alleviating the latter's three principal weaknesses: (i) need to inscribe link information in the IP packet header, (ii) reactiveness—traceback occurs after the impact of DoS attack has been felt—and (iii) scalability where the effort needed to achieve IP traceback grows proportionally with the number of attack hosts engaged in a DDoS attack.

*This work was supported in part by NSF grant EIA-9972883.
†Additionally supported by NSF grants ANI-9714707, ANI-9875789 (CAREER), ESS-9806741, ANI-0082861 (ITR), and grants from the Purdue Research Foundation, Santa Fe Institute, Sprint, and Xerox.
‡Additionally supported by the Center for Education and Research in Information Assurance and Security (CERIAS).

1 Introduction

1.1 Background

Denial of service (DoS) is a pressing problem on the Internet as evidenced by recent attacks on commercial servers and ISPs and their consequent disruption of services [8]. DoS attacks [4, 11, 17, 28] consume resources associated with various network elements—e.g., Web servers, routers, firewalls, and end hosts—which impedes the efficient functioning and provisioning of services in accordance with their intended purpose. Their impact is more pronounced than network congestion due to the concentrated and targeted nature of resource depletion and clogging, which not only impacts quality of service (QoS) but can affect the very availability of services. When the attack is distributed—e.g., affected by multiple compromised hosts on the Internet—then its impact can be proportionally severe.

Susceptibility to DoS is an intrinsic problem of any service provisioning system where, at a minimum, the occurrence of a potentially valid event (e.g., service request, TCP SYN packet) must be processed to ascertain its validity. Even though the resource expenditure associated with processing a single event may be negligible, when this is multiplied by the large factors enabled by the high bandwidth of modern broadband networks, its impact can be significant no matter how small the individual processing overhead. As with prank telephone calls or ringing of door bells in days gone by, an effective means of preventing DoS attacks from occurring in the first place—also the only fundamental solution given the intrinsic susceptibility of service provisioning systems to DoS—lies in identification of the attacker which admits assigning commensurate costs (e.g., legal or economical) to the perpetrating entity. Even if the attack was instituted from compromised hosts intruded by an attacker, if the physical source of DoS traffic can be identified, then at the very least the invaded network element can be isolated or shut down, and in some instances, the attacker's identity can be further traced back by state information remnant on the compromised system.

In this paper we address two complementary problems and goals: (1) proactive prevention of spoofed IP packets from reaching their destination, and (2) reactive source identification (i.e., IP traceback) of spoofed IP flows. We describe a novel approach to proactive and reactive distributed DoS (DDoS) attack prevention—route-based distributed packet filtering—and evaluate its efficacy in Internet autonomous system (AS) topologies.

1.2 New Contributions

Route-based distributed packet filtering (DPF) uses routing information to determine if a packet arriving at a router—e.g., border router at an AS—is valid with respect to its inscribed source/destination addresses, given the reachability constraints imposed by routing and network topology. A single AS—there were 4872 autonomous systems on the Internet in 1999 [16]—can only exert a limited impact with respect to identifying and discarding forged IP flows. This is similar to the limitation of firewalls whose filtering rules reflect access constraints local to the network system being guarded. At the other extreme, if all autonomous systems and their routers implement route-based packet filtering then no

Cloudflare, Inc. Exhibit 1003 Page 1446

spoofed IP flows can escape, but its ultimate effect is not much different from that achievable by perfect ingress filtering.

Route-based DPF's main strength lies in the fact that with partial coverage or deployment—about 18% in Internet AS topologies—a synergistic filtering effect can be achieved whose collective filtering action proactively prevents spoofed IP flows from reaching other autonomous systems in the first place. This is akin to setting up road blocks at certain intersection points in a city to apprehend bank robbers where the bank robbers are constrained to take the public transportation system whose routes, in turn, are constrained by the physical street network (topology) and routing policy/algorithm imposed by the municipal transportation department. As with routing, route-based filtering occurs at two time scales—packet forwarding/discard based on table look-up (fast) and filter table update (slow)—and thus its forwarding/discard function can be performed close to line speed subject to generic processing overhead. That is, the core filtering function itself is not subject to DoS attack[1].

Perfect proactive protection, due to intrinsic connectivity properties of Internet topology, cannot be achieved with small coverage. However, its effect is strong enough such that those spoofed IP flows that cannot be prevented from reaching their targets are sufficiently sparse and, as a consequence, their origin can be localized to within a small, constant number of sites (less than 5 for Internet AS topologies). Thus, as with probabilistic packet marking (PPM), effective IP traceback—strictly speaking, "AS traceback" since the granularity of our study is targeted at AS graphs—is achieved which serves as a deterrent in addition to facilitating reactive recovery. Our approach can also be applied to intra-domain router graphs, although in this paper we focus on AS graphs for comparative performance evaluation purposes using Internet AS topologies.

An interesting aspect of the performance evaluation side of our work is that effectiveness of both proactive and reactive filtering depend intimately on the connectivity structure of the underlying AS graph topology. In particular, we show that power-law Internet AS topology [6] is crucial in facilitating small coverage with strong proactive and reactive filtering effect. After defining relevant performance metrics, we show how route-based DPF depends on topology, filter placement, and multi-path routing using both Internet AS [16] and artificially generated topologies [12, 14].

An important feature of route-based DPF is its scalability with respect to distributed DoS attack. In PPM, attack site localization efficiency deteriorates proportionally with the number of attack hosts [20]. In route-based DPF, the fraction of AS from which spoofed IP flows can reach other AS is a small subset (less than 12%) which makes harnessing attack sites when engaging in DDoS attack more difficult for the attacker. Furthermore, in route-based DPF a single spoofed IP packet arriving at a target suffices to trace back the origin of the packet. In PPM, sampling constraints allow an attacker to distribute an attack flow targeted at a common victim among many attack hosts such that the emanating individual spoofed IP flows are difficult to traceback due to their thinness.

[1] Route table and filtering table updates are potential targets of DoS attack. The large time scale associated with route table updates—filter table updates are triggered by the same events—and the preventative effect of route-based DPF with respect to protecting routers from DoS attack jointly help alleviate the control plane protection problem.

Cloudflare, Inc. Exhibit 1003 Page 1447

From an implementation perspective, DPF does not require expending IP header bits to encode link information as PPM does nor the generation of ICMP messages (as in a messaging-based version of PPM [2]). On the other hand, computing appropriate filtering tables alongside existing inter-domain routing protocols (e.g., BGP) is a nontrivial problem due to the destination-based structure of Internet routing protocols. We submit that few fundamental solutions exist for preventing distributed DoS attacks, and this paper's main contribution lies in advancing a scalable architecture for DDoS attack prevention that is effective for Internet AS topology. The architecture is implementable in IP internetworks if source-based routing information is made available to the routers (e.g., the intra-domain link-state routing protocol RIP uses global topology information). The specific implementation approaches for inter-domain protocols—in particular, BGP—and their trade-offs are challenging problems unto their own and outside the scope of this paper.

The rest of the paper is organized as follows. In the next section, we give a summary of related works. In Section 3, we define route-based packet filtering, the key notions, and performance metrics. We also discuss the issues surrounding performance evaluation including benchmarking with Internet AS topology, filter placement, and multi-path routing. In Section 4, we present performance results based on benchmark experiments with both real and artificial network topologies. We conclude with a discussion of our results.

2 Related Work

Several types of DoS attacks have been identified [8, 17, 28] with the most basic DoS attack demanding more resources than the target system or network can supply. Resources may be network bandwidth, file system space, processes, or network connections [17]. While host-based DoS attacks are more easily traced and managed, network-based DoS attacks which exploit weaknesses of the TCP/IP protocol suite [15], represent a more subtle and challenging threat [17, 23]. Network-based DoS attacks, by default, employ spoofing to forge the source address of DoS packets, and thereby hide the identity of the physical source [4]. Previous works have focused on detecting DoS attacks and mitigating their detrimental impact upon the victim [1, 13, 24, 26]. This approach does not eliminate the problem, nor does it deter potential attackers.

A number of recent works have studied source identification—also called IP traceback [23]—which span a range of techniques with their individual pros and cons. In link testing, the physical source of an attack is identified by tracing it back hop-by-hop through the network [27]. Traceback is typically performed manually and recursively repeated at the upstream router until the originating host is reached. The drawbacks of link testing include multiple branch points, slow traceback during an attack, communication overhead due to message exchange, and administrative constraints between network operators [27]. The audit trail approach facilitates tracing via traffic logs at routers and gateways [22]. This method is conducive to off-line traceback of DoS attacks. A principal weakness, however, is the high storage and processing overhead incurred at routers which can exert a significant

Cloudflare, Inc. Exhibit 1003 Page 1448

burden. In behavioral monitoring [17], the likely behavior of an attacker during a DoS attack is monitored to identify the source. For example, an attacker may perform DNS requests to resolve the name of the target host which may not be resident in its local name server's cache. During a DoS attack, an attacker may try to gauge the impact of the attack using various service requests including Web and ICMP echo requests. Thus, logging of such events and activities can reveal information about the attacker's source.

In packet-based traceback, packets are marked with the addresses of intermediate routers, in some sense, an inverse operation of source routing and similar to the IP Record Route option [21]. The victim uses information inscribed in packets to trace the attack back to its source. A related method is generating information packets—separate from data packets—that convey analogous path information as ICMP traceback messages to the victim [2]. In these methods, overhead in the form of variable-length marking fields that depend on path length or traffic overhead due to extra messaging packets are incurred. Probabilistic packet marking [3, 23] has been proposed for achieving the best of both worlds—space efficiency in the form of constant marking field and processing efficiency in the form of minimal router support—at the expense of introducing uncertainty due to probabilistic sampling of a flow's path. The effectiveness of probabilistic packet marking was analyzed when considering the intrinsic vulnerability of marking field spoofing [20] and shown that the attacker's location can be localized to within 5 equally likely sites on the Internet under single-source attack. Improved marking schemes including authentication were studied in [25]. In spite of its efficiency properties, PPM has several drawbacks: (i) spoofed packets are allowed to exert their debilitating influence on server resources before being reactively curtailed; (ii) bits in the IP header must be expended to inscribe link information; and (iii) uncertainty of IP traceback amplifies proportionally with the number of hosts partaking in the distributed DoS attack. We show that route-based distributed packet filtering, in addition to matching the IP traceback prowess of PPM, solves its three weaknesses.

Packet filtering is a network security mechanism for controlling what data can flow to and from a network affected routers or firewalls [31]. Filtering decisions, typically, are made based on packet content including source/destination addresses and port numbers. As a means of preventing network-based DoS attacks, ingress filtering in border gateways has been proposed for limiting IP source address spoofing [5, 7, 29]. Ingress filtering requires a prolonged period to be broadly deployed on the Internet, and even then, it is subject to attacks from AS that are not compliant (see Section 4.5 for a discussion of its performance effects).

3 Route-based Distributed Packet Filtering

3.1 Route-based Detection of Spoofed IP Packets

Consider the undirected graph—interpreted as an AS graph—shown in Figure 3.1, which depicts the routes from node 2 to all other nodes (solid arrows). Assume a host belonging to AS 7 is attempting a DoS attack targeted at a server residing in AS 4 by using a forged source IP address belonging to AS 2.

Cloudflare, Inc. Exhibit 1003 Page 1449

A border router belonging to AS 6 at the peering point with AS 7—if cognizant of the route topology—would recognize that a packet originating from AS 2 destined to AS 4 would not enter through link $(7,6)$ implying that its source address must be spoofed. Such packets could be discarded at AS 6 thus proactively protecting AS 4 from the DoS attack. Note that in this specific instance AS 6 only need inspect the source IP address to determine that no packet from AS 2—irrespective of destination IP address—can arrive on link $(7,6)$. This example serves to illustrate the potential opportunities available by exploiting routing information to identify and filter spoofed packets at forwarding points in the system. We remark that the above description—from an inter-domain IP routing perspective—

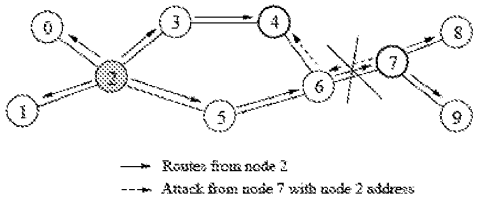


Figure 3.1: Illustration of route-based packet filtering executed at node 6. Note 7 uses IP address belonging to node 2 when attacking node 4.

is imprecise. First, an edge in the AS graph between a pair of nodes is in general a set of peering point connections and all corresponding border routers must carry out the specified filtering tasks. Second, two or more IP prefixes belonging to the same destination AS may lead to different paths on an AS topology. This is incorporated in our AS model by allowing multi-path routing. Third, we ignore potentially relevant classification of AS nodes into stub, multi-homed, and transit AS where only the latter may engage in routing (i.e., in the sense of packet forwarding). When we speak of an AS node performing route-based filtering, it must be understood that the finer resolution picture is more complex, although logical consistency between the two descriptions is achieved.

Consider the case where the attack host residing in AS 7 uses an IP address belonging to AS 8 when attacking servers in AS 4. The gateway at AS 6 cannot unambiguously decide that the IP packet with source address in AS 8 is spoofed since it may indeed have come from AS 8 (and forwarded by AS 7). This demonstrates that performing route-based filtering at a *single* site can achieve only so much. Route-based distributed packet filtering aims to achieve a synergistic, proactive filtering effect through the collective action of a small number of AS nodes. The key objectives of DPF can be summarized as follows: (i) maximize proactive filtering of spoofed IP packets; (ii) for bogus packets that do get through, minimize the number of sites that could have sent the packets (IP traceback); achieve objectives (i) and (ii) while minimizing the number of sites at which route-based filtering is carried out; (iv) in tandem with (iii), find the optimum sites where filtering is to be performed.

Cloudflare, Inc. Exhibit 1003 Page 1450

3.2 Maximal and Semi-maximal Filters

Let $G = (V, E)$ be an undirected graph representing Internet AS topology. We remark that our framework and conclusions may be carried over to router topologies *within* an AS, however, presently little is known about the internal structure of large, commercial autonomous systems[2], and testing needs to await further measurement studies. Let $\mathcal{L}(u, v)$ denote the set of all loop-free paths from u to v where $u, v \in V$. A routing algorithm and its computed routes lead to a subset $R(u, v) \subseteq \mathcal{L}(u, v)$. An IP packet $M(s, t)$ with source IP address s and destination IP address t is routed through the network according to $R(s, t)$. If $|R(s, t)| > 1$, we assume a separate network mechanism resolves selection among multiple paths. Performance results for multi-path routing are discussed in Section 5.7.

A *filter* $F_e : V^2 \to \{0, 1\}$ is a function defined for link $e = (u, v) \in E$ where this is interpreted to mean that a gateway router in v acting as a peering point inspects an IP packet $M(s, t)$ arriving on e, then decides whether to forward the packet ($F_e(s, t) = 0$), or filter—i.e., discard—the packet ($F_e(s, t) = 1$). We call F_e a *route-based packet filter with respect to* R if

$$F_e(s, t) = 0 \quad \text{for } e \in R(s, t).$$

To avoid cluttering, with a slight abuse of notation, we use "$e \in R(s, t)$" to mean that link e is on some path belonging to $R(s, t)$. Thus a route-based filter is "safe" in the sense that it does not discard packets that are potentially consistent with respect to R as judged locally at link e. A route-based filter is *maximal* if it satisfies $F_e(s, t) = 0$ if, and only if, there exists a path in $R(s, t)$ with e as one of its links. Thus a maximal route-based filter carries out all the filtering of spoofed IP traffic that is possible without adversely affecting routing of non-spoofed IP packets as determined by R. If a set of route-based filters collectively were "perfect" in the sense that no spoofed datagram is allowed to reach its destination, then this may be viewed as providing a form of authentication service. Computing a maximal route-based filter—e.g., represented as a table—is straightforward, but it requires in general $O(n^2)$ space ($n = |V|$) which is an overwhelming burden to place on routers that are expected to perform fast table look-up.

A *semi-maximal* filter is a maximal filter which uses only the source IP address of a packet to carry out its filtering (i.e., a projection of F_e). In other words, $\hat{F}_e(s, t)$ is a *semi-maximal filter with respect to* R if

$$\hat{F}_e(s, t) = \begin{cases} 0, & \text{if } e \in R(s, v) \text{ for some } v \in V; \\ 1, & \text{otherwise.} \end{cases}$$

Hence, its filtering capability is, in general, less than that of its maximal counterpart, i.e., $\hat{F}_e(s, t) \leq F_e(s, t)$. Although we lose in potential filtering power—by not much as shown in Section 5—a semi-maximal filter can be represented by a filtering table in linear space which brings it in the domain

[2]Router topologies may obey power-law connectivity structure similar to AS topologies [6, 18]. There are, however, semantic differences between AS and router topologies—e.g., geographical distance between two nodes in an AS graph may not be meaningful—which have to be taken into consideration when advancing interpretations.

Cloudflare, Inc. Exhibit 1003 Page 1451

of feasibility if not practicality. As with techniques for speed-up of routing table look-up, further optimizations will be needed to affect practical implementations. Protocol implementation issues are discussed in Section 6.

3.3 Performance Measures for Distributed Packet Filtering

3.3.1 Filtering Effect: Attacker and Victim Perspectives

A (semi) maximal filter is *distributed* if it is executed at more than one node in V. We will use T to denote a subset $T \subseteq V$ of nodes where filtering is performed. We call $\gamma = |T|/|V|$ the *coverage ratio*. To quantify and measure the collective filtering effect of route-based DPF—including IP traceback—we define a set of performance metrics that is used in the rest of the paper. First, we define two families of variables $S_{a,t}$ and $C_{s,t}$ $(a, s, t \in V)$ which are then used to define other high-level measures of more direct interest and relevance to quantifying DDoS mitigation.

$S_{a,t}$ denotes the set of nodes—more precisely, the set of IP addresses belonging to an AS node in $S_{a,t}$—that an attacker at AS a can use as spoofed source IP addresses to reach t without being cut-off by filters executed at autonomous systems in T. By definition, $a \in S_{a,t}$ for all $a, t \in V$. The larger the set $S_{a,t}$, the more options an attacker at a has in terms of forging the IP source address field with a bogus address which will go undetected and unhindered with respect to R at filters in T. Whereas $S_{a,t}$ is defined from the attacker's perspective, $C_{s,t}$ captures the victim's perspective and denotes the set of nodes that could have sent an IP packet $M(s,t)$ with spoofed source IP address s and destination address t which did not get filtered on its way. We allow $s \in C_{s,t}$ for all $s, t \in V$ in the definition. The larger $C_{s,t}$, the more uncertain the victim at t is upon receiving spoofed packet $M(s,t)$ with respect to its true origin. If $|C_{s,t}| = 1$, then this means that IP address s cannot be used by any attacker $a \in V$ (outside of s itself) to mount a spoofed DoS attack aimed at t. Figure 3.2 illustrates the impact of route-based distributed filtering on curtailing the attacker's ability to engage in spoofing. Without

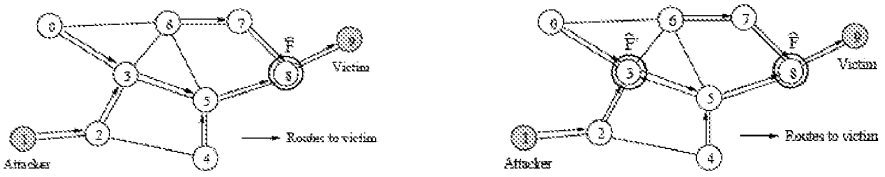


Figure 3.2: Left: With route-based filtering executed at node 8, the spoofable address range at attack site 1 is reduced from $S_{1,9} = \{0, 1, 2, 3, 4, 5, 6, 7, 8\}$ to $\{0, 1, 2, 3, 4, 5\}$. Right: Distributed filtering with filter $\hat{F}$ at AS 3, the spoofable range further reduces to $S_{1,9} = \{1, 2\}$.

route-based filtering, an attacker residing at AS 1 can disguise himself with undetectable spoofed IP addresses belonging to AS 0–8, i.e., $S_{1,9} = \{0, 1, \ldots, 8\}$, when attacking a server in AS 9. With route-based filtering at AS 8, the spoofable address range shrinks to $\{0, 1, \ldots, 5\}$. With distributed filtering at AS 8 and AS 3, $S_{1,9} = \{1, 2\}$.

Cloudflare, Inc. Exhibit 1003 Page 1452

3.3.2 Proactive Filtering Measures

The most immediate—but also practically useless—proactive filtering effect is captured by $\Phi_1(\tau)$ which is defined as $\Phi_1(\tau) = |\{t : \forall\ a \in V,\ |S_{a,t}| \leq \tau\}|\ n^{-1}$. The range of τ is $\tau \geq 1$. Thus, $0 \leq \Phi_1(1) \leq 1$ denotes the fraction of AS's that cannot be reached by spoofed packets from anywhere. The closer $\Phi_1(1)$ is to 1, the fewer the number of AS's exposed to DoS attack. For $\tau \geq 2$, $\Phi_1(\tau)$ has a less relevant and directly interpretable meaning. In spite of its appealing semantic relevance, we will show that $\Phi_1(1)$ is near zero for Internet AS topologies when the coverage ratio γ is not near 1, and thus of little import as a performance measure.

A more subtle, but practically relevant, proactive performance measure is given by

$$\Phi_2(\tau) \quad = \quad \frac{|\{a : \forall\ t \in V,\ |S_{a,t}| \leq \tau\}|}{n}.$$

$\Phi_2(1)$ measures the fraction of attack sites from which sending spoofed IP packets targeted at other AS is futile since they will be filtered by nodes in T. Thus, if $\Phi_2(1) = 0.8$, then an attacker wishing to engage in DDoS attack cannot make productive use of attack hosts residing in 80% of all autonomous systems. This imposes an upper bound on the distributedness of DDoS attack achievable by any attacker, severely limiting the latter the closer $\Phi_2(1)$ is to 1. Policy-wise, it is also possible for other AS to be "on guard" with respect to traffic emanating from AS where mounting an attack is feasible. As with $\Phi_1(\tau)$, $\Phi_2(\tau)$ does not have directly relevant semantics for $\tau \geq 2$. Unlike $\Phi_1(1)$, however, $\Phi_2(1)$ achieves large values for Internet AS topologies. It is our principal proactive performance measure.

$\Phi_3(\tau)$, Θ, and $\Psi_2(\tau)$ are auxiliary metrics capturing proactive filtering with less sharply delineated semantics which are defined as $\Phi_3(\tau) = |\{(a,t) : |S_{a,t}| \leq \tau\}|/n(n-1)$, $\Psi_2(\tau) = |\{s : \forall\ t \in V,\ |C_{s,t}| \leq \tau\}|/\ n$, and $\Theta = |\{(a,s,t) : s \in S_{a,t}\}|/n(n-1)^2 = |\{(a,s,t) : a \in C_{s,t}\}|/n(n-1)^2$. $\Phi_3(1)$ denotes the fraction of all attacker-victim AS pairs (out of a total of $n^2 - n$) where the attacker cannot reach the victim with spoofed IP packets. Thus an attacker whose aim is to wreck general havoc on the Internet via DoS attack without specific interest in a particular victim may choose random attack site-victim pairs to do so. The larger $\Phi_3(1)$, the less impact such random DDoS attacks will have. Θ captures the reduced attack volume—ratio of unfilterable packets—when, in addition, attacks are mounted using randomly inscribed source IP addresses. $\Psi_2(\tau)$, viewed from the attacker's perspective, represents the fraction of all (spoofable) IP addresses whose use would allow the victim to localize the attack site to within τ locations.

3.3.3 Reactive Filtering Measure: IP Traceback

The performance measures defined in the previous section are proactive in nature in that they capture how effectively spoofed IP packets are prevented from reaching their destination in the first place by filters in T. Perfect proactivity, as captured by $\Phi_1(1) \gg 0$, however, is intrinsically difficult to achieve in Internet topologies due to their connectivity structure unless the coverage ratio $\gamma \approx 1$. Significant, albeit imperfect, proactive filtering is captured by the quantitative values of $\Phi_2(1)$, $\Phi_3(1)$, Θ, and $\Psi_2(\tau)$, where $\Phi_2(1)$ plays the most relevant role.

Cloudflare, Inc. Exhibit 1003 Page 1453

Since not all spoofed IP packets can be effectively filtered, complementing the proactive performance measures is the reactive metric $\Psi_1(\tau)$ which captures the IP traceback (or source identification) effect:

$$\Psi_1(\tau) \quad = \quad \frac{|\{t : \forall\ s \in V,\ |C_{s,t}| \leq \tau\}|}{n}.$$

For example, $\Psi_1(5)$ represents the fraction of (target) autonomous systems which, when attacked with an arbitrary spoofed IP packet, can resolve the attack location to within 5 possible attack sites. The parameter $\tau \geq 1$—meaningful for values greater than 1—represents the uncertainty associated with IP traceback localization[3]. If $\Psi_1(\tau) = 1$ for τ a small constant, then those spoofed IP flows that cannot be prevented from penetrating the "filter net" spanned by nodes in T can be effectively localized with respect to their true attack origin to within τ candidate sites, i.e., we achieve IP traceback. This can act as a deterrent as well as allow responsive on-line counter measures to be undertaken against the attacking party including isolation of offending flows at appropriate peering points. $\Psi_3(\tau)$ is analogously defined as $\Phi_3(\tau)$ with $C_{s,t}$ in place of $S_{a,t}$ but does not have relevant semantics and is omitted from further consideration.

4 Performance Evaluation and Benchmarking Issues

4.1 Overall Objectives

Formally a route-based (semi) maximal distributed filter $\mathcal{F}$ is given by a triple $\mathcal{F} = \langle G, T, R \rangle$ where $G = (V, E)$ is the AS graph, $T \subseteq V$ the subset of AS where route-based filtering is performed, and R is the routing algorithm. For two route-based DPF's $\mathcal{F} = \langle G, T, R \rangle$ and $\mathcal{F} = \langle G, T', R \rangle$ with $T \subseteq T'$, it can be checked that $|S_{a,t}| \leq |S'_{a,t}|$ and $|C_{s,t}| \leq |C'_{s,t}|$ for all $a, s, t \in V$. This, in turn, implies $\Phi_2(1) \leq \Phi'_2(1)$ and $\Psi_1(\tau) \leq \Psi'_1(\tau)$ for all $\tau \geq 1$. Similar monotonicity properties hold for the other performance metrics. Moreover, $\Phi'_2(1) = \Psi'_1(1) = 1$ if $T' = V$ (i.e., there is a trivial lower bound). Evaluating the effectiveness of $\mathcal{F}$ with respect to the proactive and reactive performance measures entails studying its dependence on topology G, the size of the filter net T, its structure, and routing R. Our goal is to uncover the above relationships and find feasible means to economize T when trying to achieve a target performance.

4.2 Influence of Topology

Empirical evidence shows that Internet AS topology exhibits power-law connectivity [6, 10] which may also hold for router topologies [18]. Power-law graph structure induces "centers" where connectivity is concentrated on a few nodes, with most vertices possessing sparse connectivity (e.g., comprised of AS stubs and non-transit multi-homed AS). A key aspect of our DDoS benchmark evaluation is to

[3] See [20] for a discussion of IP traceback localization issues—also called uncertainty factor—under probabilistic packet marking.

Cloudflare, Inc. Exhibit 1003 Page 1454

ascertain if, and how, topology affects proactive and reactive performance of DPF. We employ 1997–1999 Internet AS topologies taken from NLANR [16] which have been used in other studies aimed at understanding the connectivity structure of Internet topology, especially with respect to its recently discovered power-law property [6]. In addition to actual Internet AS topologies, we use artificial Internet topology generators [12, 14] and random graphs to perform comparative benchmarking. An (unintended) side effect of our study is the reverse evaluation of artificial topology generators with respect to capturing relevant graph properties—above-and-beyond power-law relations—in the context of DPF.

4.3 Filter Placement

In addition to the influence of the *size* of the filter net T on DPF performance, for a given coverage ratio $\gamma = |T|/n$, the selection of the nodes in T is a key performance variable. We consider the effect of choosing T randomly—we sample from V uniformly randomly until a target coverage size $|T|$ is reached—and by more customized design rules, in particular, the case when T forms a vertex cover[4]. It can be checked that T being a vertex cover (VC) is neither a sufficient nor necessary condition for $\Phi_2'(1) = \Psi_1'(\tau) = 1$. However, since a VC forms a cover of all edges in the graph—being VC implies that on any path, at least every other vertex on the path belongs to T—it may be expected that the VC property is conducive to enhancing the performance of DPF. In tandem, the presence of "centers" in power-law graphs leads one to expect that a small coverage ratio γ may be achievable.

Finding a minimal VC in a graph is an NP-complete problem [9]. We use two approximation algorithms—one with a constant factor guarantee and the other a heuristic—for finding small VCs. The first algorithm is a well-known constant-factor approximation scheme whose output is guaranteed to be at most twice as large as an optimal VC [19]. There is a randomization component, and the approximation scheme is run multiple (in our evaluations 10) times with the smallest VC constituting the final output. The second algorithm is a heuristic: little is known rigorously about its behavior although, in practice, it oftentimes outperforms the constant-factor approximation scheme. The heuristic—greedy algorithm—iteratively grows a VC by picking a node which covers the most remaining uncovered edges. The presence of centers in power-law graphs makes it more conducive for the heuristic to find small VCs which is verified in our performance results. We use the minimum VC found by the two algorithms as our T.

We also consider a rank-based placement strategy to isolate VC's tendency to pick "large centers" vis-à-vis its complete edge covering property. We call this algorithm *Rank*, and it simply orders vertices by their degree (in decreasing order) and selects the top γ percentage as elements of T. We show that *Rank*—even with larger T—affects a diminished performance effect than VC.

[4] T is a *vertex cover* of G if every edge in E is incident on some node in T.

Cloudflare, Inc. Exhibit 1003 Page 1455

4.4 Maximal vs. Semi-maximal Filters

Our performance results are for semi-maximal filters which are, in general, less powerful than maximal filters. In comparative evaluations we show that replacing semi-maximal with maximal filters results only in an incremental improvement in proactive/reactive filtering performance. The marginal performance difference—e.g., maximal filtering can localize to within 4 sites for IP traceback instead of 5—justifies the use of semi-maximal filters when performing route-based DPF in addition to its consideration of efficiency.

4.5 Ingress Filtering

Consider the case when the nodes in T were to perform ingress filtering only. Then for coverage ratio $\gamma = |T|/n$, the DDoS prevention performance effect as captured by $\Phi_2(1)$ and $\Psi_1(\tau)$ would be: $\Phi_2(1) = \gamma$, $\Psi_1(\tau) = 0$ for $\tau < n - |T|$ and $\Psi_1(\tau) = 1$ for $\tau \geq n - |T|$. Thus, ingress filtering, unless carried out almost everywhere, is an ineffective DDoS prevention strategy. Even when $\gamma = 0.95$, for the 1999 Internet AS topology where $n = 4872$, IP traceback capability as captured by Ψ_1 incurs a non-constant uncertainty factor of 243 (the trivial number of possible attack sites to investigate when trying to pin down the true attack location). There is little compelling reason for a group of AS in the global Internet to form trusted security partnerships based on mandatory ingress filtering since the collective performance effect is low. In contrast, we show that when AS in T implement route-based DPF, then with $\gamma < 0.2$ coverage $\Phi_2(1) > 0.88$ and $\Psi_1(5) = 1$ for 1997–1999 Internet AS topologies.

4.6 Routing

The set of feasible routes is influenced by topology but, in addition, we consider the impact of having multiple paths from source to destination. Note that $R' \subseteq R$ implies $\Phi_2(1) \leq \Phi'_2(1)$ and $\Psi_1(\tau) \leq \Psi'_1(\tau)$ for all $\tau \geq 1$. Consider routing policies that allow R to have up to m separate paths—not necessarily disjoint—between two nodes. This allows us to evaluate the influence, other things being equal, that the more paths are permitted when routing a packet from source to destination, the more easily the packet can elude route-based filtering when using spoofed source IP addresses. The latter is due to the attack site's spoofable IP address space $S_{a,t}$ having increased. When multi-path routing is performed between two nodes a and t with $|R(a,t)| = m$, we select m shortest paths from $\mathcal{L}(a,t)$. In the case where two or more candidates have the same path length, we choose the path coming first in the canonical (i.e., lexicographic) order. We give special names to two extreme forms of R: *loose* and *tight*. "R=*loose*" means that all possible paths among two nodes can be used for routing, i.e., $R(a,t) = \mathcal{L}(a,t)$. When R allows only a single routing path ($m = 1$), we choose a shortest path between a and t, and denote this case as "R=*tight*."

Cloudflare, Inc. Exhibit 1003 Page 1456

5 Performance Results

5.1 Set-up

We have built a performance evaluation tool called *dpf* which implements the benchmarking set-up described in Section 4. *dpf* consists of three core modules: `cover`, `dpf`, and `stats`. `cover` handles the generation of T with various input specifications including random selection, VC, and rank ordering. `dpf` is the main module which computes $S_{a,t}$ and $C_{s,t}$; its input specification include the filter type, T, and routing algorithm. `stats` takes the output of `dpf` and computes the various performance measures including $\Phi(1)$ and $\Psi(\tau)$. We use topology generators *Inet* [12] and *Brite* [14] to generate benchmark graphs which are included in the test suite.

5.2 Proactive Filtering Effect

5.2.1 Limitations to Perfect Proactivity

$\Phi_1(1)$ measures the fraction of AS which are immune from DoS attack—i.e., no spoofed IP packet can reach them—distributed or single-source. Figures 5.1 (left) and (middle) show $\Phi_1(\tau)$ as a function of τ for different coverage and routing combinations for 1997 Internet AS topology ($|V| = 3015$ and $|E| = 5230$). In Figure 5.1 (middle), $\Phi_1(\tau) = 0$ up to $\tau = 4$. That is, perfect proactivity where there exists at least one AS that is immune from DoS attack from anywhere is unachievable at 18.9% coverage ratio under the best of circumstances. The two graphs show that, overall, $R = tight$ gives better performance than $R = loose$ and, other thing being equal, T being VC—the size of the 1997 Internet AS vertex cover is 18.9%—is more effective than T being random even with higher coverages Rnd30 ($\gamma = 0.3$) and Rnd50 ($\gamma = 0.5$). These plots depict a general trend but are not otherwise very useful since for performance evaluation purposes only $\Phi_1(1)$ has direct relevance.

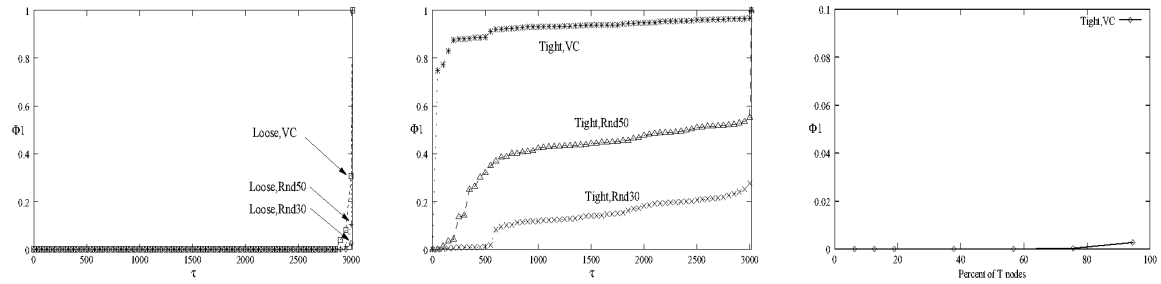


Figure 5.1: 1997 Internet AS topology. Left: $\Phi_1(\tau)$ for $R = loose$. Middle: $\Phi_1(\tau)$ for $R = tight$. Right: $\Phi_1(1)$ as a function of $|T|$ while maintaining VC property.

Figure 5.1 (right) shows that the limitation to achieving perfect proactivity does not change when the VC is grown to larger sizes up to 100%. Although eventually $\Phi_1(1)$ becomes positive when coverage is above 90%, its value is negligible to warrant the high cost of almost full coverage. Perfect

Cloudflare, Inc. Exhibit 1003 Page 1457

proactivity as captured by $\Phi_1(1)$ is intrinsically difficult to attain, and should not be construed as a viable performance goal.

5.2.2 DDoS and Proactive Filtering

$\Phi_2(1)$ measures the fraction of AS from which DoS attacks cannot be launched since all spoofed packets—whoever their target—will be detected and filtered. $\Phi_2(1)$ thus puts an upper bound on the distributedness of DDoS attacks. Figure 5.2 (left) and (middle) show $\Phi_2(\tau)$ as a function of τ for $R = loose$ and $R = tight$. As with Φ_1, $\Phi_2(\tau)$ for $\tau \geq 2$ does not have a concrete, relevant meaning and are shown to depict the general trend. Figure 5.2 (right) is the more relevant plot which shows $\Phi_2(1)$

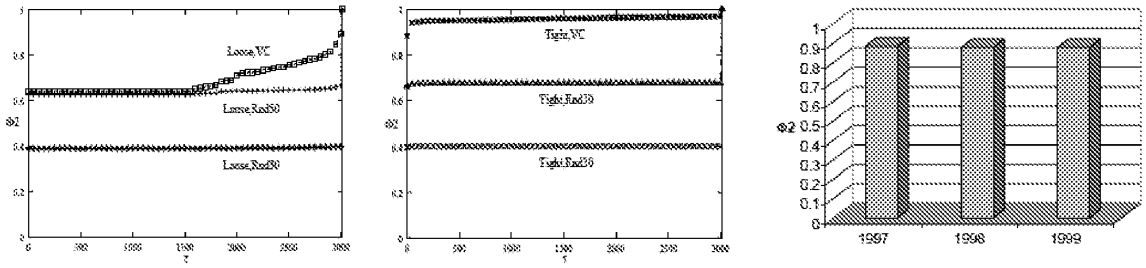


Figure 5.2: 1997 Internet AS topology. Left: $\Phi_2(\tau)$ as a function of τ for $R = loose$. Middle: Corresponding graph for $R = tight$. Right: $\Phi_2(1)$ for 1997–1999 Internet AS topologies.

for Internet AS topologies during 1997–1999. $\Phi_2(1)$ achieves a value of around 88% during the three years which implies that only 12% of all autonomous systems can be used by attackers to launch DoS attacks. Since the number of AS's has grown from 3015 in 1997 to 3878 in 1998 to 4872 in 1999, the absolute number of possible attack sites has grown commensurately. However, as a percentage, viable attack sites have remained well-behaved at 12%.

Figure 5.3 (left) shows $\Phi_3(\tau)$ as a function of τ with $\Phi_3(1) = 0.96$. That is, only 4% of all source-destination AS pairs are feasible attack AS/victim AS combinations from the attacker's perspective, where spoofed packets emitted from the attack AS can reach the victim AS. For example, an attacker who tries to enlist attack hosts in a DDoS attack by intruding these hosts will waste 96% of its effort

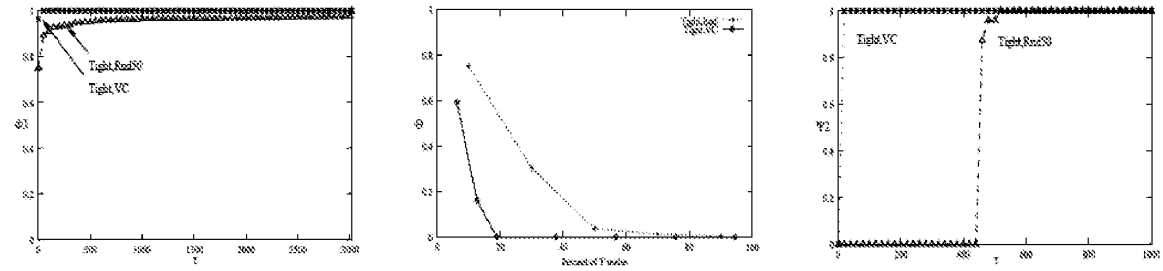


Figure 5.3: 1997 Internet AS topology. Left: $\Phi_3(\tau)$. Middle: Θ as a function of $|T|$. Right: $\Psi_2(\tau)$.

Cloudflare, Inc. Exhibit 1003 Page 1458

if the source-destination AS's are chosen randomly. Thus proactive filtering erects barriers in terms of effort and cost to mounting effective DoS attacks which, in turn, can act as a deterrent in addition to its primary curtailing effect. Figure 5.3 (middle) shows Θ, the coarsest measure, which represents the fraction of source, destination, and spoof address triples (a, t, s) where a host residing at AS a is able to send an IP packet to target AS t with spoofed source IP address s. We observe that for coverage above 20%, the fraction of forgeable triplets shrinks to near 0. This means that if, in addition to a and t, the spoof address s is randomly generated, then the spoofed IP packet has almost zero chance of reaching its target. Figure 5.3 (right) shows $\Psi_2(\tau)$ as a function of τ. For an uncertainty factor of $\tau = 20$, the fraction of traceable—i.e., to within τ sites—IP addresses is nearly 1. Collectively, these results show that the attacker's effort, resources, and sophistication needed to launch a successful DDoS attack is significant and brought about by route-based DPF's proactive filtering effect.

5.3 Reactive Filtering Effect: IP Traceback

As shown in the previous section, eliminating *all* spoofable IP flows is an unrealistic goal given its intrinsic difficulty with respect to Internet AS connectivity. A different consequence of proactive filtering is the more subtle, complementary effect where spoofed IP flows that cannot be prevented from penetrating the network system can be localized to within a few possible sites. This is affected by DPF filtering sufficiently many flows such that the remaining spoofable IP flows form a sparse subset which, in turn, facilitates source identification, i.e., IP traceback.

Figure 5.4 shows $\Psi_1(\tau)$ as a function of τ for $R =$ *loose*, *tight*, and $T =$ VC, Rnd30, Rnd50. The general trend shows that $\Psi_1(\tau)$ undergoes a sharp transition at some τ value, especially for $T =$ VC and $R =$ *tight*. Figure 5.5 (left) shows $\Psi_1(\tau)$ for 1997–1999 Internet AS topologies for $1 \leq \tau \leq 10$.

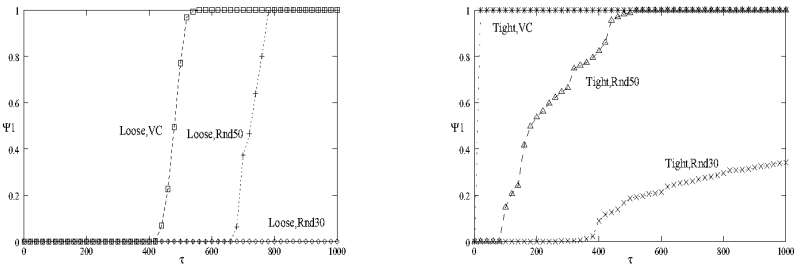


Figure 5.4: 1997 Internet AS topology. Left: $\Psi_1(\tau)$ for $R =$ *loose*. Right: $\Psi_1(\tau)$ for $R =$ *tight*.

We observe that across 1997, 1998, and 1999, $\Psi_1(5)$ is preserved—i.e., every attack can be localized to within 5 candidate sites—and the only performance difference occurs for $\tau < 5$ where $\Psi_1(\tau) < 1$. IP traceback is achieved "instantly" and thus allows speedy on-line response by the attacked site with respect to actions against the perpetrating attack site. Compared to probabilistic packet marking, route-based DPF is proactive even with respect to IP traceback since a *single* spoofed IP packet suffices to reveal the attacker's AS location to within a small constant number of locations. In PPM, a

Cloudflare, Inc. Exhibit 1003 Page 1459

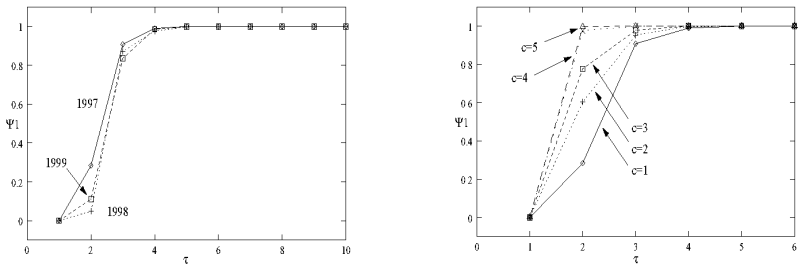


Figure 5.5: Left: $\Psi_1(\tau)$ for 1997–1999 Internet AS topologies. Right: Shape of $\Psi_1(\tau)$ for $|T| = c \cdot |VC|$ with dilation factor $c = 1, 2, \ldots, 5$.

sufficient number of DoS attack packets must be received before the attack path can be reconstructed by the probabilistically inscribed link values in the IP datagram [20, 23]. Thus, not only can route-based DPF emulate the IP traceback prowess of PPM, it can do so more efficiently and speedily than PPM. Figure 5.5 (right) shows the marginal benefit of increasing the number of nodes in T after achieving T =VC. We observe that increasing the size of the vertex cover as represented by the dilation factor $c = |T|/|VC|$ has only an incremental effect. This shows that much of the IP traceback effect is attained at the smaller vertex cover size (18.9%) which facilitates economy of coverage and deployment.

5.4 Maximal Filters vs. Semi-maximal Filters

All the results reported in this paper are, by default, based on semi-maximal filters. To ascertain the potential performance loss due to not using maximal filters, we compare filtering performance with respect to $\Psi_1(\tau)$ and $\Phi_2(1)$. Figure 5.6 (left) shows $\Psi_1(\tau)$ for 1997 Internet AS topology as a function of τ when performing route-based DPF with maximal versus semi-maximal filters under $R = \textit{tight}$ and T being VC. We observe that the performance difference in IP traceback capability as captured by $\Psi_1(\tau)$ is minimal. For example, for $\tau = 5$, there is no performance difference. Figure 5.6 (right) compares $\Phi_2(1)$ for maximal and semi-maximal filters which, in fact, are equal. Thus the marginal performance difference coupled with space efficiency warrants the use of semi-maximal filters when implementing route-based DPF.

5.5 Impact of Network Topology

5.5.1 Internet AS Topology

Figure 5.7 shows the vertex cover sizes (expressed as a percentage) and $\Psi_1(5)$ values for 1997–1999 Internet AS topologies. We observe that $|\text{VC}|/n$—as well as $\Psi_1(5)$ and $\Phi_2(1)$—remain invariant over 1997–1999. In the rest of this section we focus on $\Psi_1(\tau)$ and discuss the results for $\Phi_2(1)$ when their performance is qualitatively different. The size of the vertex cover plays an important role as an intermediate indicator and facilitator of filtering performance. In fact, the smaller the VC, the better the filtering performance in spite of the small coverage ratio—note that $\gamma = |\text{VC}|/n$—which indicates

Cloudflare, Inc. Exhibit 1003 Page 1460

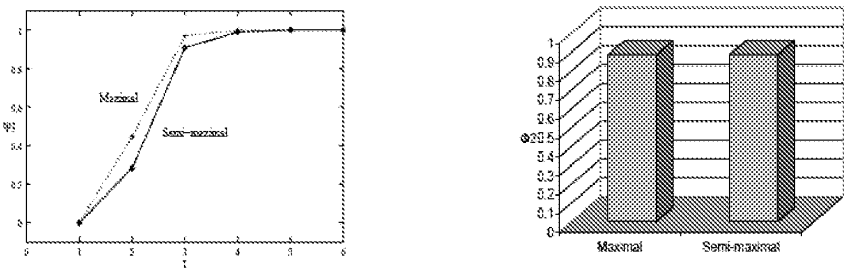


Figure 5.6: 1997 Internet AS topology. Left: Comparison of $\Psi_1(\tau)$ for maximal and semi-maximal filters. Right: Corresponding comparison of $\Phi_2(1)$ for maximal and semi-maximal filters.

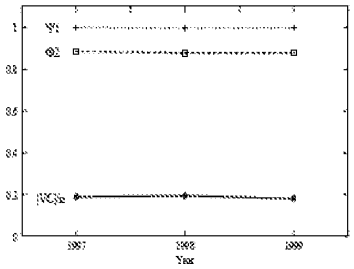


Figure 5.7: Vertex cover size $|VC|/n$ and $\Psi_1(5)$, $\Phi_2(1)$ for 1997–1999 Internet AS topologies.

that the VC property and its relative size is a useful indicator of connectivity property relevant to DPF performance.

5.5.2 Random Topology

We generate p-random graphs by connecting two nodes with link probability p. For a given Internet AS graph, we generate its corresponding random graph by setting $p = \frac{2e}{n(n-1)}$ where $e = |E|$. The specification and p values for 1997–1999 Internet AS topologies are shown in Figure 5.8 (left). The two families of graphs differ only in their connectivity pattern. Figure 5.8 (middle) shows vertex cover size

Year	n	e	p
1997	3015	5230	0.001151
1998	3878	7080	0.000942
1999	4872	9254	0.000780

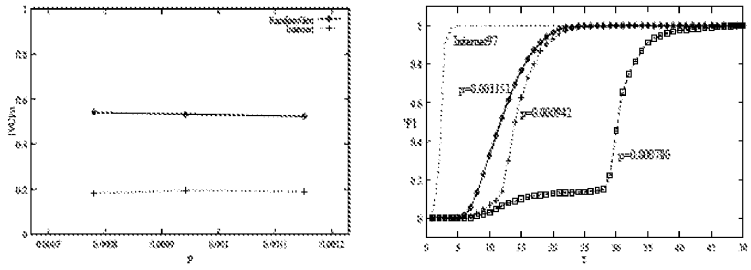


Figure 5.8: Left: Link probabilities for random graphs corresponding to 1997–1999 Internet AS topologies. Middle: $|VC|/n$ as a function of p and comparison with Internet AS. Right: $\Psi_1(\tau)$ plot.

Cloudflare, Inc. Exhibit 1003 Page 1461

of the generated random graphs and corresponding Internet AS topologies. On average, the VC sizes of the random graphs are 2.5 larger than their Internet AS counterparts. Figure 5.8 (right) shows $\Psi_1(\tau)$ as a function of τ for different topologies. In spite of engaging more nodes when performing filtering, the performance as captured by $\Psi_1(\tau)$ is significantly less than that of Internet AS topology. Moreover, the performance difference amplifies as the size of the graph increases. Recall that the performance values for 1997–1999 Internet AS topologies (see Figure 5.5 (left)) stayed invariant.

5.5.3 Inet Topology Generator

We use Inet 2.0 [12][5], a network topology generator, for generating artificial topologies closer to the Internet in their connectivity structure than random graphs. Inet is designed to generate graph topologies with connectivity properties similar to Internet AS in terms of power-law structure. Figure 5.9 (left) shows the VC sizes of Inet generated graphs and their Internet AS counterparts for 1997–1999. We observe that the VC sizes of Inet graphs are about 50% larger than corresponding Internet AS graphs. Figure 5.9 (middle) shows $\Psi_1(\tau)$ as a function of τ for Inet, Internet AS, and random graphs. We observe, as expected, that filtering performance for Inet graphs is closer to that of Internet AS than random graphs.

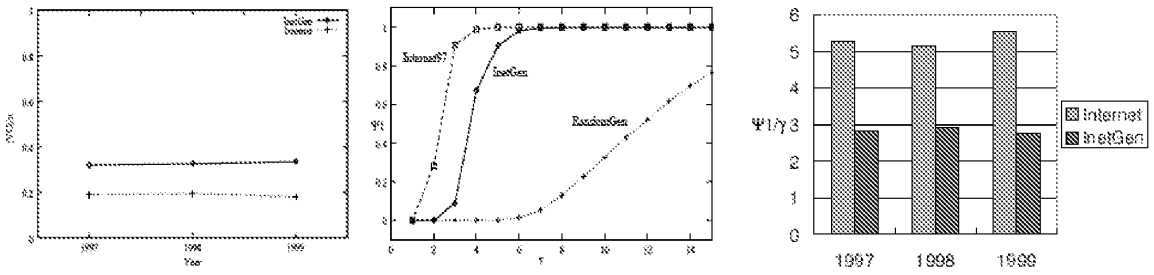


Figure 5.9: Left: VC sizes for Inet graphs and corresponding Internet AS graphs. Middle: Comparison of $\Psi_1(\tau)$ of Inet graph with Internet AS and random graphs. Right: Performance difference between Inet and Internet AS graphs normalized by VC size: $\Psi_1(5)/(|\mathrm{VC}|/n)$.

Figure 5.9 (right) shows *normalized* filtering performance $\overline{\Psi}_1(\tau) = \Psi_1(\tau)/\gamma$ for $\tau = 5$ where the relative size of the filter set is incorporated. Since $\Psi_1(1) = 1$ if $T = V$ no matter what the structure of the underlying topology, $\overline{\Psi}_1$ measures filtering performance per filter node (relative to $|V|$) which is a more accurate metric for comparative evaluation. Figure 5.9 (right) shows that there is significant difference in DPF performance between Inet and Internet AS topologies stemming, in part, from VC size difference. Inet is a topology generator whose primary feature is that of emulating power-law relations for vertex degrees as observed in [6]. The fact that the well-known VC graph property

[5]We also tested with benchmark graphs generated by Inet2.1—it was conveyed to us recently that Inet2.0 had a bug when generating large graphs of size 30K—with similar results. The Inet2.1 graphs resulted in a marginally smaller VC size—less than 2% difference—for graph sizes corresponding to 1997–1999 Internet AS topologies.

Cloudflare, Inc. Exhibit 1003 Page 1462

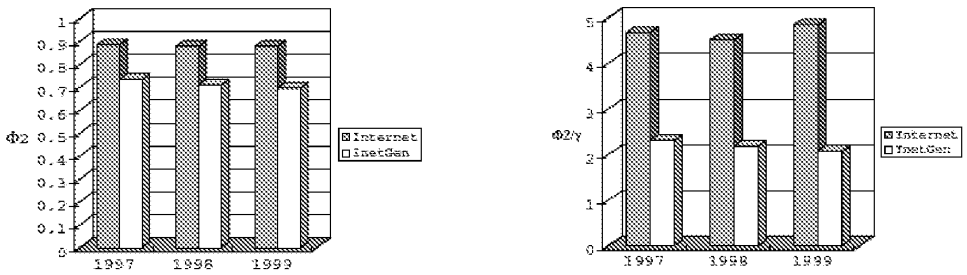


Figure 5.10: Left: Comparison of $\Phi_2(1)$ for Inet and 1997–1999 Internet AS topologies. Right: Corresponding comparison of normalized $\bar{\Phi}_2(1) = \Phi_2(1)/\gamma$.

exhibits nontrivial gaps between Internet AS and Inet topologies indicates that more refined structure may need to be uncovered within the family of power-law graphs to accurately capture the Internet's topological properties.

Figure 5.10 (left) shows a comparison of $\Phi_2(1)$ for Inet and 1997–1999 Internet AS topologies. As with $\Psi_1(\tau)$, we observe a persistent gap in reactive filtering performance which is consistent with the corresponding gap in the VC size. Figure 5.10 (right) shows normalized IP traceback performance $\bar{\Phi}_2(1) = \Phi_2(1)/\gamma$ for the same benchmark set-up which incorporates the size of the filter net in the performance measure. As expected, performance difference is further amplified.

5.5.4 Brite Topology Generator

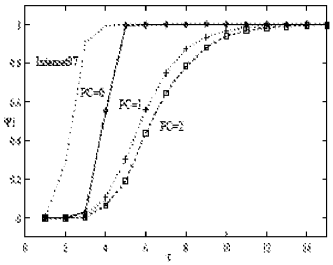
Brite [14] is a network topology generator that, in addition to capturing power-law connectivity structure, seeks to inject spatial proximity in the constructive process. Brite specifies seven parameters: size of higher plane (HS), size of lower plane (LS), number of nodes (n), number of edges added for each new node (m), node placement (NP), preferential connectivity[6] (PC), and incremental growth (IG). When PC=0, a new node is connected to node i with Waxman's probability density [30], $p_i = \alpha e^{-d/(\beta L)}$, where $0 < \alpha, \beta \leq 1$, d is the Euclidean distance between two nodes, and L is the maximum distance between any two nodes. When PC=1, a new node connects to node i with probability $\frac{d_i}{\sum_{j \in C} d_j}$ where d_i is the degree of node i and C is the set of candidate neighbor nodes. With PC=2, the probability of connecting to node i is given by $\frac{p_i d_i}{\sum_{j \in C} p_j d_j}$. Thus PC=0 considers spatial proximity only, PC=1 focuses on power-law structure as captured by node degree distribution, and PC=2 is a hybrid.

Using HS=1000, LS=10, IG=1, and n=3015, test graphs were generated with the three PC options. The specification and results for VC size are shown in Figure 5.11 (left). Figure 5.11 (middle) and (right) show the performance effects with respect to $\Psi_1(\tau)$ and $\Phi_2(1)$, respectively. When PC=0, we observe that the graph generated—in addition to not being power-law—has too small VC (3.6%). Its performance with respect to $\Psi_1(\tau)$ and $\Phi_2(1)$ is closer than that of PC=1 and 2, however, the

[6]The Brite generator [14] had a small bug with respect to option PC=2 which was fixed.

Cloudflare, Inc. Exhibit 1003 Page 1463

Topology	n	e	$\|VC\|/n$
Internet	3015	5230	18.9%
Brite (PC=0)	3029	5978	3.6%
Brite (PC=1)	3006	5935	42.7%
Brite (PC=2)	3002	5908	44.1%



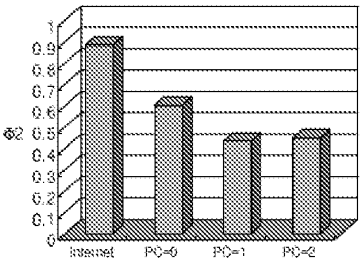
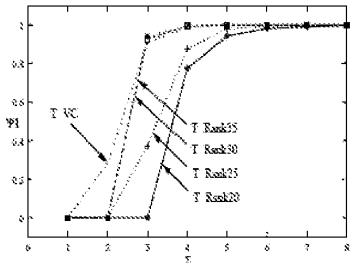


Figure 5.11: Left: VC sizes for Brite graphs with PC=0, 1, and 2. Middle: $\Psi_1(\tau)$ as a function of τ for PC=0, 1, and 2. Right: Corresponding $\Phi_2(1)$ plot.

performance gap from the corresponding Internet AS topology for $\Phi_2(1)$ is significant, being worse than that of the Inet generator. For PC=1 and 2, the VC sizes are too large, and performance for both $\Psi_1(\tau)$ and $\Phi_2(1)$ significantly worse than Internet AS (and Inet). We have tried the Brite generator with other parameter specifications but were unsuccessful in getting topologies that resemble Internet AS both from the VC size and filtering performance perspectives. We have also tried extending option PC=2 by using the weighting $\alpha p_i + (1-\alpha)\frac{d_i}{\sum_{j \in C} d_j}$ to inject both spatial and degree sensitivity in a more controlled fashion. As α increases $|VC|$ monotonically decreases, and for $\alpha = 0.13$ the VC size can be approximated to that of Internet AS with $\Psi_1(\tau)$ close to its Internet AS value. However, the corresponding $\Phi_2(1)$ performance is dismal (about 20%) when compared to Internet AS.

5.5.5 Rank-based Filter Placement

Figure 5.12 shows filtering performance of rank-based filter placement for a range of coverage ratios. Compared to VC-based filtering which has coverage ratio $\gamma = 0.189$, we observe that both $\Psi_1(\tau)$ and $\Phi_2(1)$ show diminished performance even when their coverage ratio is higher. Note that the greedy algorithm for vertex cover selects vertices that are able to cover the maximum number of remaining edges which is not the same as picking maximum degree nodes. The performance gap indicates that selection of high degree vertices as filtering sites is an important—but not the only—effect of VC-based



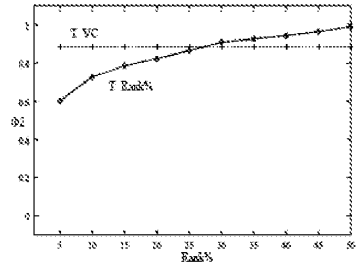


Figure 5.12: 1997 Internet AS topology. Left: Comparison of $\Psi_1(\tau)$ for rank-based filter placement with $\gamma = 0.20, 0.25, 0.30, 0.35$ versus VC-based T. Right: Corresponding comparison for $\Phi_2(1)$.

Cloudflare, Inc. Exhibit 1003 Page 1464

filter selection on DPF performance in Internet AS topology. Indeed, rank-based filter placement can lead to inefficiencies by selecting high-degree nodes that are incident on edges that are already covered. This suggests that optimal filter placement may be an NP-hard problem[7].

5.6 Ingress Filtering

Section 4.5 showed that ingress filtering is not a viable strategy for achieving proactive and reactive filtering performance for DDoS attack prevention. Since AS belonging to T represent "trusted" domains where route-based DPF is guaranteed to be executed at its border routers, ingress filtering was assumed to be carried out by AS belonging to T. It is, however, conceivable that AS in T implement route-based DPF but do not assure ingress filtering. That is, they seek to protect themselves from external DoS attack flows while allowing DoS attacks to occur within their domain including those targeted at other domains. Figure 5.13 shows proactive and reactive filtering performance when AS in T perform route-based DPF but do not perform ingress filtering. Figure 5.13 (left) shows $\Psi_1(\tau)$ for 1997–1999 Internet AS topologies. We observe that there is a performance penalty such that $\Psi_1(5) \neq 1$. On

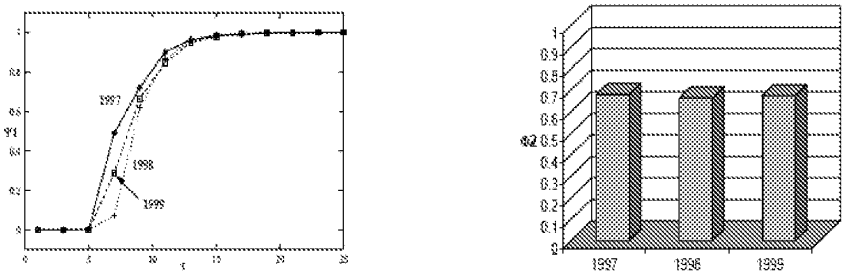


Figure 5.13: Route-based DPF without ingress filtering. Left: $\Psi_1(\tau)$ as a function of τ for 1997–1999 Internet AS topologies. Right: Corresponding $\Phi_2(1)$ values.

the other hand, $\Psi_1(20) = 1$ for all three years. That is, IP traceback can localize the attack site to within 20 locations. This is worse than 5—the number achievable with ingress filtering—however, considering that there were in the range 3000–5000 autonomous systems during 1997–1999, 20 is still a small constant, and thus managable number. Figure 5.13 (right) shows the corresponding $\Phi_2(1)$ values. $\Phi_2(1)$ drops from around 90% to 70% which is still significantly higher than the 20% proactive effect achievable with ingress filtering alone. Interestingly, the performance gap of 20% roughly corresponds to the coverage ratio $\gamma = \lceil T \rceil / n$ for VC in the Internet AS topologies.

5.7 Multi-path Routing

If multiple paths are permitted when routing packets from source to destination, the more easily packets can elude route-based filtering when using spoofed source IP addresses. Figure 5.14 shows the impact

[7]This is an interesting problem to explore in future work. Generally, it may be fruitful to investigate whether network resource allocation problems can be more accurately solved—NP-hardness will remain—for a family of power-law graphs.

Cloudflare, Inc. Exhibit 1003 Page 1465

of multi-path routing on filtering performance. Figure 5.14 (left) and (middle) show that traceback capability as captured by $\Psi_1(\tau)$ decreases gradually as the number of multi-paths allowed is increased. A similar result holds for $\Phi_2(1)$ and is shown in Figure 5.14 (right). Collectively, these performance plots show that presence of multi-paths—a more common phenomenon in Internet AS topologies than in router topologies—has a graded effect and does not significantly impact the effectiveness of route-based DPF.

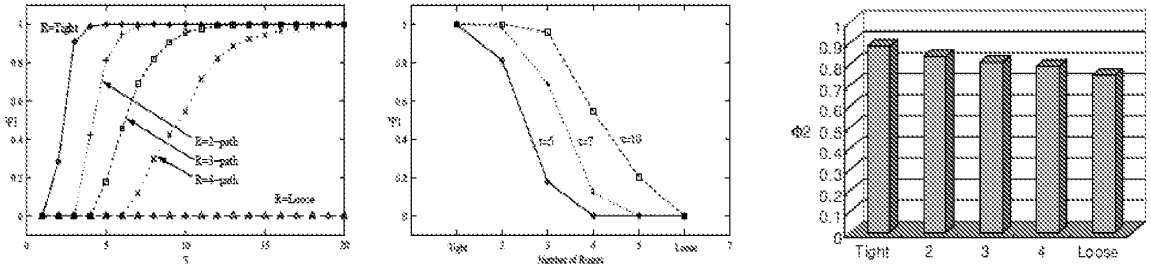


Figure 5.14: Effect of multi-path routing for 1997 Internet AS topology. Left: $\Psi_1(\tau)$. Middle: Distribution of $\Psi_1(5)$ and $\Psi_1(10)$ over the number of multi-paths. Right: $\Phi_2(1)$ for R from *tight* to *loose*.

6 Discussion of Implementation Issues

The most important implementation concern in the context of IP internetworks is not space requirement—many issues are shared with routing table look-up—but the computation of semi-maximal filters at routers belonging to participating AS nodes $T \subseteq V$ given a routing algorithm R. The main difficulty arises from the fact that IP routing—inter- and intra-domain—follows a destination-based approach where routing table update exchanges convey information about destination reachability but not "source reachability." In BGP, for example, an update message contains AS-PATH which is a sequence of AS numbers that identify the autonomous systems—starting with the AS that originated the advertisement of reachability for an IP prefix—from which, if assuming bi-directionality and bounce-back of an "ACK-to-Update" message, source information may be extracted. At a minimum this would require augmentation to BGP, or introduction of a new protocol that interacts with BGP. In addition, some form of compression and encoding may be needed to keep the return messages small—including merger of several ACK-to-Update messages at branch points—given their tendency to grow in number and size the closer they reach the target AS (the essential overhead associated with breaking routing asymmetry). Other challenges arise when trying to construct an accurate global AS routing map.

We do not have an answer to the efficient implementability question for IP internets. This may, perhaps, be route-based DPF's Achilles' heel. We view the contribution of this paper to lie in the definition and evaluation of a scalable DDoS prevention architecture as part of a set of fundamental solutions to the denial-of-service attack problem (of which there are few), and Internet specificity is

Cloudflare, Inc. Exhibit 1003 Page 1466

injected with respect to showing how filtering performance depends on topological properties of Internet AS. We believe that the performance results for route-based DPF are encouraging and suggest that investigation of how to implement route-based DPF so as to minimize overhead and cost for Internet deployment may be worthwhile and should be the focus of future effort. However, a cost-benefit analysis of the gains vis-à-vis the associated costs is a matter of debate and further examination.

7 Conclusion

We have described a proactive/reactive approach to distributed DoS attack prevention based on route-based distributed packet filtering. We have shown route-based DPF's efficacy at proactively curtailing spoofed IP flows from reaching their intended targets, including the drastically reduced Internet AS sites from which such attacks can be launched. We have shown that perfect proactivity—no spoofed IP flow can penetrate—is intrinsically difficult to achieve in Internet AS topologies while maintaining sparse (e.g., 20% or less AS sites deploying route-based DPF) coverage. However, this is mitigated by the fact that those spoofed IP flows that can penetrate the filter net can be localized to within 5 candidate sites which facilitates efficient IP traceback. Compared to probabilistic packet marking, we have shown that route-based DPF is able to solve PPM's three key weaknesses (see Section 2). We have also shown that the filtering effect achieved by route-based DPF is sensitive to the underlying Internet AS connectivity structure. In particular, we have shown that power-law structure of Internet AS topology plays an important role in facilitating proactive/reactive filtering. Finding efficient implementations for computing semi-maximal filters and evaluating the costs associated with deployment and router overhead vis-à-vis the performance benefits of route-based DPF is a major challenge and a task for future work.

References

[1] G. Banga, P. Druschel, and J. Mogul. Resource containers: A new facility for resource management in server systems. In *Proc. of the third USENIX/ACM Symp. on Operating Systems Design and Implementation (OSDI'99)*, pages 45–58, February 1999.

[2] S. Bellovin. ICMP traceback messages, March 2000. Internet Draft: draft-bellovin-itrace-00.txt (expires September 2000).

[3] H. Burch and B. Cheswick. Tracing anonymous packets to their approximate source. In *14th Systems Administration Conference (LISA 2000)*, pages 319–327, 2000.

[4] Computer Emergency Response Team (CERT). CERT Advisory CA-2000-01 Denial-of-service developments, January 2000. http://www.cert.org/advisories/CA-2000-01.html.

[5] CERT/CC, SANS Institute, and CERIAS. Consensus roadmap for defeating distributed denial of service attacks, February 2000. A Project of the Partnership for Critical Infrastructure Security, http://www.sans.org/ddos_roadmap.htm.

Cloudflare, Inc. Exhibit 1003 Page 1467

[6] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proc. of ACM SIGCOMM*, pages 251–262, 1999.

[7] P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing, May 2000. RFC 2827.

[8] Lee Garber. Denial-of-service attacks rip the Internet. *Computer*, pages 12–17, April 2000.

[9] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman and Company, 1979.

[10] R. Govindan and A. Reddy. An analysis of Internet inter-domain topology and route stability. In *Proc. IEEE INFOCOM '97*, 1997.

[11] John Howard. *An Analysis of Security Incidents on the Internet.* PhD thesis, Carnegie Mellon Univerisity, August 1998.

[12] C. Jin, Q. Chen, and S. Jamin. Inet: Internet Topology Generator. Technical Report CSE-TR-443-00, Department of EECS, University of Michigan, 2000.

[13] C. Meadows. A formal framework and evaluation method for network denial of service. In *Proc. of the 1999 IEEE Computer Security Foundations Workshop*, June 1999.

[14] A. Medina and I. Matta. Brite: A flexible generator of Internet topologies. Technical Report BU-CS-TR-2000-005, Boston University, January 2000.

[15] Robert Morris. A weakness in the 4.2BSD Unix TCP/IP software. Technical Report Computer Science #117, AT&T Bell Labs, February 1985.

[16] National Laboratory for Applied Network Research. Routing data, 2000. Supported by NFS, http://moat.nlanr.net/Routing/rawdata/.

[17] NightAxis and Rain Forest Puppy. Purgatory 101: Learning to cope with the SYNs of the Internet, 2000. Some practical approaches to introducing accountability and responsibility on the public internet, http://packetstorm.securify.com/papers/contest/RFP.doc.

[18] J. Pansiot and D. Grad. On routes and multicast trees in the Internet. *Computer Communication Review*, 28(1):41–50, 1995.

[19] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity.* Prentice Hall, Inc., 1982.

[20] K. Park and H. Lee. On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack. Technical Report CSD-TR 00-013, Department of Computer Sciences, Purdue University, June 2000.

[21] Jon Postel. Internet protocol, September 1981. RFC 791.

[22] Glenn Sager. Security fun with OCxmon and cflowd, November 1998. Presentation at the Internet 2 Working Group.

[23] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Practical network support for IP traceback. In *Proc. of ACM SIGCOMM*, pages 295–306, August 2000.

Cloudflare, Inc. Exhibit 1003 Page 1468

[24] C. Schuba, I. Krsul, M. Kuhn, E. Spafford, A. Sundaram, and D. Zamboni. Analysis of a denial of service attack on TCP. In *Proc. of the 1997 IEEE Symp. on Security and Privacy*, pages 208–223, May 1997.

[25] D. Song and A. Perrig. Advanced and authenticated marking schemes for IP traceback. Technical Report UCB/CSD-00-1107, Computer Science Department, University of California, Berkeley, 2000. To appear in IEEE INFOCOM 2001.

[26] O. Spatscheck and L. Peterson. Defending against denial of service attacks in Scout. In *Proc. of the third USENIX/ACM Symp. on Operating Systems Design and Implementation (OSDI'99)*, pages 59–72, February 1999.

[27] Cisco Systems. Characterizing and tracing packet floods using Cisco routers, Aug 1999. http://www.cisco.com/warp/public/707/22.html.

[28] Computer Emergency Response Team. Denial of service, February 1999. Tech Tips, http://www.cert.org/tech_tips/denial_of_service.html, 2nd revision.

[29] Computer Emergency Response Team. Results of the distributed-systems intruder tools workshop, November 1999. http://www.cert.org/reports/dsit_workshop.pdf.

[30] B. Waxman. Routing of multipoint connections. *IEEE Journal of Selected Areas in Communications*, pages 6(9):1617–1622, December 1988.

[31] E. Zwicky, S. Cooper, D. Chapman, and D. Ru. *Building Internet Firewalls.* O'Reilly & Associates, Inc., 2nd edition, 2000.

Cloudflare, Inc. Exhibit 1003 Page 1469

EXHIBIT 1042

T. M. Gil, "Multops: a data structure for denial-of-service attack detection," Dissertation, Submitted to the Division of Mathematics and Computer Sciences of Vrije Universitei (Amsterdam, Netherlands), Dec. 27, 2000

MULTOPS: a data structure for denial-of-service attack detection

by

Thomer M. Gil

tmgil@cs.vu.nl

Submitted to the Division of Mathematics and Computer Science
in partial fulfillment of the requirements for the degree of

Doctorandus Computer Science

at the

VRIJE UNIVERSITEIT

December 2000

© 2000. Thomer M. Gil. All rights reserved.

The author hereby grants to MIT and Vrije Universiteit permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author .....
Division of Mathematics and Computer Science
Vrije Universiteit
Dec 27, 2000

Certified by .....
M. Frans Kaashoek
Associate Professor of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
Thesis Supervisor

Certified by .....
Massimiliano Poletto
Research Staff, Massachusetts Institute of Technology
Thesis Supervisor

Certified by .....
Andrew S. Tanenbaum
Professor of Computer Science, Vrije Universiteit
Thesis Supervisor

MULTOPS: a data structure for denial-of-service attack detection

by

Thomer M. Gil

Submitted to the Division of Mathematics and Computer Science
on Dec 27, 2000, in partial fulfillment of the
requirements for the degree of
Doctorandus Computer Science

Abstract

A denial-of-service (DoS) attack is an attempt by a single person or a group of people to disrupt an online service. In a bandwidth attack, attackers clog links or routers by generating a traffic overload. This can have serious consequences to companies that rely on their online availability to do business. The ubiquity of tools to organize DoS attacks and the determination of some people to wreak havoc make for potential future problems. This thesis proposes a MUlti-Level Tree for Online Packet Statistics (MULTOPS): an attack-resistant data structure enabling routers to detect ongoing bandwidth attacks by searching for significant asymmetries between packet rates to and from different subnets. Statistics are kept in a tree that dynamically adapts its shape to (1) reflect changes in packet rates, and (2) avoid (maliciously intended) memory exhaustion. A MULTOPS is suitable to detect the type of bandwidth attack that occurred on a large scale in February 2000. To remain undetected, the attacker has to launch the attack from a large number of distinct sites which makes mounting the attack more difficult. This will hopefully discourage many attackers.

Thesis Supervisor: M. Frans Kaashoek
Title: Associate Professor of Electrical Engineering and Computer Science, Massachusetts Institute of Technology

Thesis Supervisor: Massimiliano Poletto
Title: Research Staff, Massachusetts Institute of Technology

Thesis Supervisor: Andrew S. Tanenbaum
Title: Professor of Computer Science, Vrije Universiteit

Cloudflare, Inc. Exhibit 1003 Page 1472

To Lisa

"At that time Michael, the great prince who protects your people, will arise. There will be a time of distress such as has not happened from the beginning of nations until then... Those who are wise will shine like the brightness of heavens... Many will go here and there to increase knowledge." (Daniel 12:1-4)

Cloudflare, Inc. Exhibit 1003 Page 1473

Cloudflare, Inc. Exhibit 1003 Page 1474

Contents

1 Introduction **9**

2 Overview of DoS attacks **11**

2.1 IP Spoofing 11
2.2 Distributed DoS attacks 12
2.3 Smurf: ICMP flood 13
2.4 Trinoo: UDP flood 14
2.5 SYN flooding 14
2.6 Stealth bombs 15
2.7 Bandwidth attack classification 16
2.8 Solutions 16

3 Design of MULTOPS **19**

3.1 Overview 19
3.2 Assumptions 20
3.3 Data structure 21
3.4 Algorithm 21
3.5 Rates 24
3.6 Unfolding 24
3.7 Example 24
3.8 Folding 25
3.9 Memory exhaustion attacks 26

4 Implementation of MULTOPS **27**

4.1 Click 27
4.2 IPRateMonitor 28
4.3 Data structure 29
4.4 Algorithm 30
4.5 Unfolding 30

Cloudflare, Inc. Exhibit 1003 Page 1475

4.6 Folding . . . 32
4.7 Memory exhaustion attacks . . . 32
4.8 Measurements . . . 34

5 Detecting bandwidth attacks with MULTOPS **37**
5.1 Flood and D-Flood . . . 37
5.2 Flood+ . . . 38
5.3 Stealth . . . 38
5.4 D-Stealth/Flashcrowd . . . 39
5.5 SYN flooding . . . 40

6 Future work **41**

7 Conclusion **43**

8 Acknowledgements **45**

A Click config file for performance measurement **49**

B Netmasks **51**

Cloudflare, Inc. Exhibit 1003 Page 1476

List of Figures

2-1 ICMP without IP spoofing . . . 13
2-2 ICMP with IP spoofing . . . 13
2-3 "Smurf" attack . . . 14
2-4 SYN flood . . . 15
2-5 Bandwidth attack classification . . . 16

3-1 Location of MULTOPS-equipped routers . . . 20
3-2 MULTOPS . . . 22
3-3 Algorithm for MULTOPS . . . 23

4-1 Simple Click configuration . . . 27
4-2 Click configuration with IPRateMonitor . . . 28
4-3 Counter and Stats . . . 29
4-4 Code executed for each packet . . . 31
4-5 Memory use of IPRateMonitor as a result of folding . . . 33
4-6 Relation between bandwidth and memory use . . . 34
4-7 Packet rate as a function of memory limit . . . 35
4-8 CPU cycles per packet as a function of memory limit . . . 36

5-1 Example bandwidth attack setup . . . 38

Cloudflare, Inc. Exhibit 1003 Page 1477

Cloudflare, Inc. Exhibit 1003 Page 1478

Chapter 1

Introduction

A denial-of-service (DoS) attack is a malicious attempt by a single person or a group of people to cripple an online service. This can have serious consequences for companies such as Amazon and eBay which rely on their online availability to do business. On February 9, 2000, Yahoo, eBay, Amazon.com, E*Trade, ZDNet, Buy.com, the FBI, and several other Web sites fell victim to DoS attacks ([1], [2], [3]) resulting in millions of dollars in damage. In May 2000 the same fate befell Slashdot.org [4].

Sophisticated tools to gain root access to other people's machines are freely available on the Internet ([5], [6]). These tools are easy to use, even for computer illiterates. Once a machine is cracked, it is turned into a "zombie" under control of one "master." The master is operated by the attacker [7]. The attacker can instruct all its zombies to send bogus data to one particular destination. Simultaneously, the resulting traffic can clog links, cause routers near the victim or the victim itself to fail under the load. Similarly, a phone number can be attacked by letting a handful of people continuously call that number. The type of DoS attack that causes problems by generating an overload of traffic is known as a **bandwidth attack**. This thesis focuses on bandwidth attacks.

Several reasons underlie the absence of a solution against bandwidth attacks. Both IP and TCP can be used as dangerous weapons quite easily. Since all Web traffic is TCP/IP based, attackers can release their malicious packets on the Internet without being conspicuous or easily traceable. It is the mass of all packets together that poses a threat rather than single characteristics of individual packets. A bandwidth attack solution is, therefore, more complex than a straightforward filter in a router.

A key problem to tackle when solving bandwidth attacks is **attack detection**. Detecting a bandwidth attack might be easy in the vicinity of the victim, but gets more difficult as the distance (i.e., hop count) to the victim increases. In addition, an attack detection mechanism must be able to establish the source(s) of the attack. Furthermore, any mechanism to detect bandwidth attacks must be robust, for it will be a likely target of attacks itself. In short, a bandwidth attack detection mechanism must be:

- **sensitive**: detect bandwidth attacks
- **accurate**: no false alarms

Cloudflare, Inc. Exhibit 1003 Page 1479

- **fair**: punish attackers only
- **robust**: withstand attacks on the detection mechanism itself
- **inexpensive**: not induce an unacceptable per-packet overhead

TCP acknowledges the receipt of one or more packets by sending back a packet to the sender. The bandwidth attack detection mechanism proposed in this thesis exploits this property of TCP. In addition, TCP is an "**adaptive**" protocol: it reacts to the loss of acknowledgements by decreasing the sending rate. What follows is that the packet rate **to** a certain host or subnet should always be proportional to the packet rate **from** that host or subnet. A router observing a certain packet rate in one direction with a significantly lower packet rate in the opposite direction can suspect that the slower side is unable to cope with the traffic it is receiving and may, therefore, be under attack.

This thesis proposes a **MUlti-Level Tree for Online Packet Statistics (MULTOPS)**: an attack-resistant data structure enabling routers to detect ongoing bandwidth attacks by keeping track of packet rates to and from subnets. A MULTOPS can zoom in on subnets that behave conspicuously to gain more precise information and help determine the source(s) of the attack. The MULTOPS attack detection mechanism is a novel technique to detect bandwidth attacks that are mounted using unadaptive protocols such as UDP and ICMP. This covers all February 2000 attacks. In most cases, the MULTOPS bandwidth attack detection mechanism fails to detect attacks that are mounted using adaptive protocols such as TCP.

The rest of this thesis is organized as follows. Chapter 2 gives an overview of DoS attacks. Chapter 3 introduces the concepts behind MULTOPS bandwidth detection. Chapter 4 deals with the implementation of MULTOPS. Chapter 5 looks at how well a MULTOPS detects different bandwidth attacks. Chapter 6 gives recommendations for future research. Chapter 7 concludes this thesis.

Cloudflare, Inc. Exhibit 1003 Page 1480

Chapter 2

Overview of DoS attacks

DoS attacks can be divided into at least three types: (1) exploiting implementation bugs, (2) server resource attacks, and (3) server bandwidth attacks. DoS attacks that exploit implementation bugs are (in principle) relatively easy to solve by installing proper patches.

Attacks on server resources (memory, disk space, etc.) are more difficult to stop because they often exploit legitimate protocol features rather than simple bugs. Such attacks exploit features of applications or of protocols above the transport layer. Example attacks against pseudonym mail servers are generating exponential mail loops or maliciously creating a large number of pseudonym accounts in a short time [8]. These high-level DoS attacks must be handled by the application in question, since it is impossible for lower network or system layers to detect or counter the specific problem.

Attacks on server bandwidth are performed by congesting the victim's networks with (useless) traffic. Bugs in routers on the victim's network can cause the routers to crash, compounding the problem. Some attacks generate easily identifiable packets that can be filtered or rate-limited because they never occur in high volume during normal operations [9]. More subtly, bandwidth attacks may be caused by traffic that looks entirely normal except for its high volume [10]. Usually, bandwidth attacks require a group of attackers to cooperate in order to generate sufficient traffic.

The rest of this chapter is organized as follows. Sections 2.1 and 2.2 deal with some general DoS attack issues. Sections 2.3, 2.4, 2.5, and 2.6 describe specific DoS attacks. Section 2.7 classifies bandwidth attacks. Section 2.8 deals with some solutions that have been proposed so far.

2.1 IP Spoofing

IP spoofing is lying about one's own IP address. When writing to a raw socket, a program can fill the header fields of an IP packet with whatever it wants. This requires root permission which is always known to a user running Linux on a PC. Since routing is done based on the IP destination address only, the IP source address can be anything. In some cases, attackers use one specific IP source address on all outgoing IP packets to make all returning IP packets—and possibly ICMP messages—go to the unfortunate owner of that address.

Cloudflare, Inc. Exhibit 1003 Page 1481

Attackers also use IP spoofing to hide their location on the network. Section 2.3 looks at an attack that uses IP spoofing to flood a victim.

Ingress/Egress filtering ([11], [12]) is performed by routers to effectively eliminate IP spoofing. Routers match the IP source address of each outgoing packet against a fixed set of IP addresses. If no match is found, the packet is dropped. For example, a router at MIT will only route outgoing packets that have an IP source address from subnet 18.0.0.0/8 (see Appendix B for an explanation of this notation). Although IP spoofing is a nice weapon for attackers to wield, it is in many cases no more than just that. As more attackers are involved in an attack, each operating from different networks, the need for IP spoofing becomes less. Ingress and egress filtering will make the life of an attacker more difficult, but they are far from being a panacea.

Stefan Savage et al. have devised a scheme called **IP Traceback** [13] that assists in tracking down attackers post-mortem. Their technique requires routers to probabilistically mark packets such that the receiving end can reconstruct the route that packets followed, provided that enough packets were sent. This technique looks very promising, but its main weakness is that it only assists in finding attackers; it provides no protection against bandwidth attacks. IP Traceback is primarily effective by being a deterrent.

2.2 Distributed DoS attacks

One single non spoofing attacker that generates more traffic than the victim can handle is easily identifiable. The defense is to deploy a filter in a router on the victim's or—preferably—the attacker's network that blocks all packets from the attacker. If the attacker randomly spoofs IP addresses, then an ingress/egress filtering edge router on the attacker's network will stop most packets. That forces the attacker to only use IP addresses from her own network for spoofing. Confronted with this attacker, the victim is faced with a bigger problem. As soon as the victim knows from which network the attack is coming, a filter could be deployed that drops all traffic from that specific network. Even though this stops the attack, it also denies service to all other clients on the attacker's network. This solution is a little blunt, but still constitutes some progress because the rest of the world remains unaffected by this filter.

The attacker can get around this filter by launching an attack from different networks. The victim is now faced with a distributed DoS (DDoS) attack. If the attacker compromises 1 machine on N networks, then each machine needs to generate $1/N$ of the required bandwidth to flood the victim. If N is large (and, consequently, the amount of traffic generated from each network small), the victim can no longer label traffic from one network as malicious or benign based on its relative volume. Installing a filter that drops packets from all suspected networks means denying service to all users on all those networks and, therefore, defeats its own purpose.

The conclusion is that sophisticated attackers will try to use IP source addresses from many different subnets. To do that, the attacker has to either find networks without ingress/egress filtering edge routers, or launch her attack from many different sites so that each attacking client can generate an inconspicuous amount of traffic. So far, attackers have not done any sophisticated IP spoofing. In many cases they use

Cloudflare, Inc. Exhibit 1003 Page 1482

0.0.0.0 as IP source address. The sad conclusion is that many attacks work only because ingress/egress filtering is inactive in many routers. However, it seems only a matter of time before more sophisticated attacks as described above will occur.

2.3 Smurf: ICMP flood

A "Smurf" attack [9] is carried out using ICMP. ICMP is a protocol used for sending control messages. ECHO REQUEST and ECHO REPLY are two such messages. UNIX's "ping" program, for example, uses ICMP to measure round-trip delays between two machines. Figure 2-1 exemplifies this. A sends an ECHO REQUEST message to B. B replies with an ECHO REPLY message. The notation "A → B" represents the IP source address and IP destination address in the IP packet that carries the ICMP message. B knows where to send its ECHO REPLY by looking at the IP source address in the arriving IP packet.

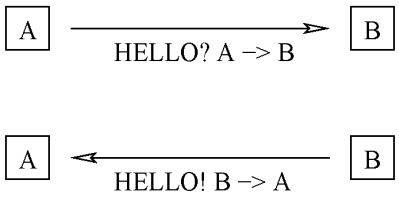


Figure 2-1: ICMP without IP spoofing

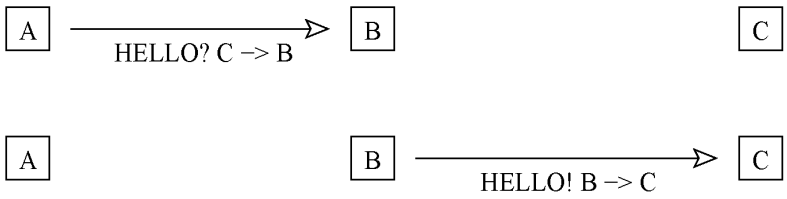


Figure 2-2: ICMP with IP spoofing

If A spoofs the IP source address, the situation as shown in Figure 2-2 occurs. A sends an ECHO REQUEST message to B, but spoofs the IP source address by using C, not A. As a result, C receives an ICMP ECHO REPLY from B, seemingly out of the blue. This does not pose an immediate threat to C because C can simply discard the message. The situation becomes harmful when A sends an ECHO REQUEST to a broadcast address. Each machine on the receiving network gets the ECHO REQUEST and each machine responds. If A spoofs the IP source address, an innocent party will receive all the ECHO REPLYs (see Figure 2-3). Consequently, links and routers to C might get clogged by all the traffic.

Cloudflare, Inc. Exhibit 1003 Page 1483

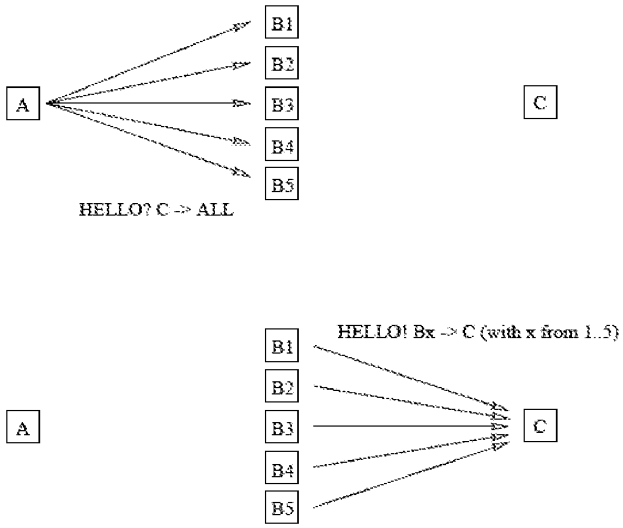


Figure 2-3: "Smurf" attack

"Smurf" attacks work because of a bug in many ICMP implementations. A host may never send an ECHO REPLY in response to a broadcasted ECHO REQUEST. Unfortunately, many ICMP implementations fail to check the IP destination address of the incoming ECHO REQUEST.

IP spoofing is an absolute necessity for this attack to succeed. If the attacker fails to spoof the IP source address, the ECHO REPLY flood will come back to himself which is similar to shooting into one's own foot.

2.4 Trinoo: UDP flood

Trinoo attacks are far more sophisticated than "Smurf" attacks. After compromising a machine—a whole science of its own [6]—a small daemon is installed which waits for commands from a master: the compromised machine is turned into a zombie. Communication between the master and the zombies is often encrypted so as to complicate matters for network intrusion detectors. At any point in time, the master can instruct all the zombies to start sending UDP packets to one destination.

2.5 SYN flooding

Normal establishment of a TCP connection requires three packets to be sent between the client and the server: (1) a client sends a SYN packet, (2) the server allocates a TCP control block and sends back a

Cloudflare, Inc. Exhibit 1003 Page 1484

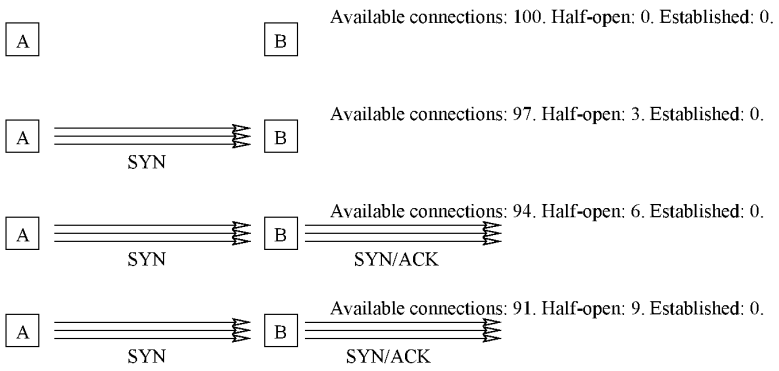


Figure 2-4: SYN flood

SYN/ACK packet, and (3) the server waits for an ACK to come back from the client. This is called a 3-way handshake. As long as the ACK from step 3 has not come back to the server, the connection is in half-open state. When no ACK comes back from the client at all, the connection remains in half-open state until TCP times out after a few minutes. When TCP times out, the allocated control block becomes available again.

A simple attack (see Figure 2-4) is for A to continuously send SYN packets in spoofed IP packets. The SYN/ACK packets will go to an innocent third party (who will drop them) and the required ACK packets never get sent by anyone. This will cause B to run out of TCP control blocks very fast with all available connections in half-open state. A server without any available TCP control blocks is unable to accept any more incoming TCP connections.

Several solutions have been proposed for solving SYN floods: lowering the TCP timeout, increasing the number of TCP control blocks, SYN cookies [14] that eliminate the need to store information on half-open connections, and special firewalls that buffer SYN packets.

2.6 Stealth bombs

To remain filter-proof, the next generation of bandwidth attacks might simply generate a huge amount of normal TCP/IP traffic. A JavaScript running in a browser that pops up a few dozen windows each fetching a Web page from one server means certain death for that server if a few thousand people are willing to run this script in their browser simultaneously [10]. Such a script could easily spread by means of self-replicating e-mail viruses.

We make a distinction between a Stealth bomb and a **flashcrowd**. A flashcrowd is a group of benign clients that overload one Web server or link close to the Web server by causing too much traffic. A Web server displaying the results of the Super Bowl finals on a Web page is very likely to be subjected to such

Cloudflare, Inc. Exhibit 1003 Page 1485

	not distributed		distributed
	no spoofing	**spoofing**	
adaptive	Stealth	*n.a.*	D-Stealth/Flashcrowd
unadaptive	Flood	Flood+	D-Flood

Figure 2-5: Bandwidth attack classification

a scenario. A flashcrowd is similar to a SYN flood in the sense that it can exhaust all the available TCP control blocks on the server. It differs from a SYN flood in that the required 3-way handshake to establish a TCP connection progresses normally.

2.7 Bandwidth attack classification

Many different types of bandwidth attacks exist. For later use, a convenient classification of bandwidth attacks is chosen; see the table in Figure 2-5. "+" indicates the use of spoofing; "D-" indicates a distributed attack. This classification uses three properties: protocol-type, distribution, and whether or not IP spoofing is involved. An adaptive protocol is one that adjusts its rate when packets get lost. TCP is an adaptive protocol. Examples of non adaptive protocols are UDP and ICMP.

No distinction is made between a distributed attack with spoofing and a distributed attack without spoofing. One could also argue that a non distributed attack with spoofing and a distributed attack are of the same type, too. From the victim's point of view this is true, but in chapter 5 we will show that, from a MULTOPS point of view, they are different.

SYN floods are not bandwidth attacks, but rather server resource attacks. SYN floods do not fit in this classification.

2.8 Solutions

Ingress/egress filtering and IP Traceback are methods to deal with attackers that use IP spoofing. Both are described in section 2.1.

A Cisco white paper [15] explains how to use a Cisco router to characterize an attack and how to reconfigure the router to minimize the negative effects. Unfortunately, in some February 2000 attacks, it was not a server that crashed under the load, but rather an upstream router. In some cases, the machine that was supposed to stop bandwidth attacks crashed itself [16].

The IETF is currently writing a draft that proposes a new ICMP message type: ICMP Traceback. These traceback messages should help solve bandwidth attacks. When forwarding packets, routers can, with a low probability (1/20000), generate a traceback message that is sent along to the destination. With enough

Cloudflare, Inc. Exhibit 1003 Page 1486

traceback messages from enough routers along the path, the traffic source and path can be determined [17]. Unfortunately, this idea has several problems:

- Routers that support these traceback messages need to be widely deployed before the victim can actually trace back packets to their source.
- Attackers can generate traceback messages, too. The IETF fails to describe a proper authentication mechanism for traceback messages. Authentication implies encryption. Encryption implies computation, which leads to a vulnerability for resource (CPU) attacks. Attackers could send many bogus traceback messages to a host that will, consequently, spend most of its time decrypting these messages.
- Routers assist the attacker by adding more packets to the flood.
- The victim of a bandwidth attack might not receive enough traceback messages because they might get dropped by overloaded routers.

Cloudflare, Inc. Exhibit 1003 Page 1487

Cloudflare, Inc. Exhibit 1003 Page 1488

Chapter 3

Design of MULTOPS

This chapter deals with the design of MULTOPS. Chapter 4 deals with its implementation.

3.1 Overview

A MULTOPS is a data structure designed to keep track of aggregate packet rates to and from different subnets. MULTOPS should be deployed in routers (see Figure 3-1). A MULTOPS is a tree where nodes in higher layers in the tree contain aggregate packet rates to and from subnets of increasing size. Nodes in the deepest level of the tree contain packet rates to and from single hosts. More precisely: the root node contains packet rates to and from subnets with netmask 8. Children of the root node contain packet rates to and from subnets with netmask 16. Nodes in the bottom layer contain packet rates to and from single hosts. See Appendix B for an explanation on netmasks and their notation.

A MULTOPS is suitable for detecting bandwidth attacks by using a significant asymmetry in packet rates to and from a certain subnet as an indication for an attack on or from that subnet. This is based on the assumption that the TCP packet rate to a host or subnet is always proportional to the packet rate from that host or subnet.

The shape of the tree is determined by packet rates. If the packet rate to or from a subnet exceeds a certain threshold, a node is created to keep track of packet rates for subnets within that subnet. Similarly, a node is destroyed if the aggregate packet rate to and from the associated subnet falls below a certain threshold. This mechanism enables a MULTOPS to zoom in and out on subnets dynamically.

A MULTOPS is primarily effective as part of a mechanism to detect Flood and D-Flood attacks, i.e., attacks using an unadaptive protocol without IP spoofing. To successfully detect other attacks, collaboration between routers is required.

Cloudflare, Inc. Exhibit 1003 Page 1489

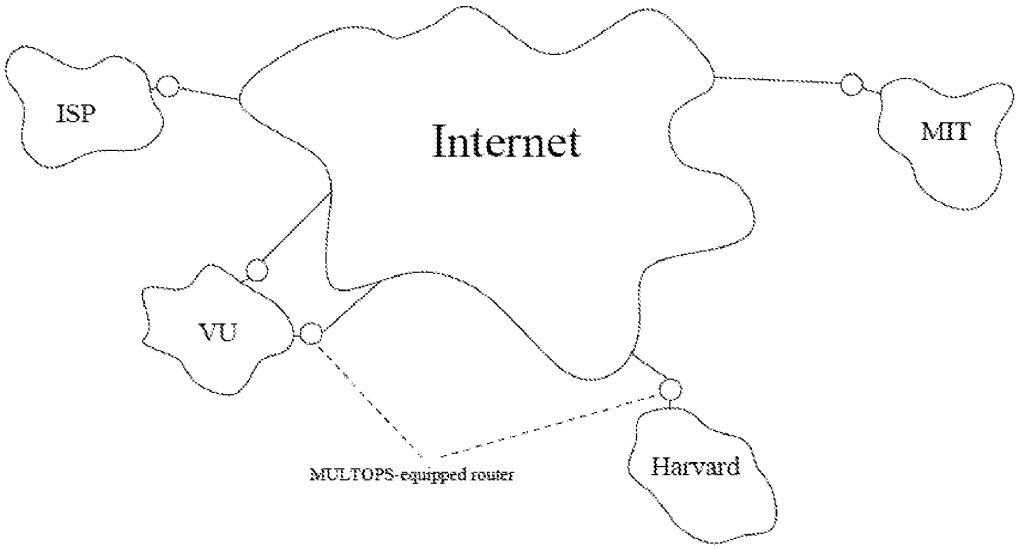


Figure 3-1: Location of MULTOPS-equipped routers

3.2 Assumptions

A few assumptions about bandwidth attacks are made in this thesis. These assumptions apply to the February 2000 attacks.

The attacker(s) and the victim are separated by at least one router. This assumption covers all February and May 2000 attacks.

In the normal case, packet rates between two hosts are symmetric. We say that packet rates between hosts A and B are symmetric when the packet rate from A to B is equal to the packet rate from B to A times some constant k. In other words: the packet rate to a host is proportional to the packet rate from that host. What follows is that packet rates between two subnets are also symmetric.

Although this assumption does not fully cover reality, it does apply to hosts that communicate with commonly used TCP implementations. TCP acknowledges the receipt of one or more packets by sending back an ACK packet. For example, a client that receives a Web page from a server will send back ACK packets for TCP packets that come from the server. Some TCP implementations send an ACK packet for every received packet, others send one ACK packet for k received packets. Neither violates this assumption.

This assumption also applies to hosts that are sending back and forth ICMP ECHO REQUEST/REPLY messages.

Cloudflare, Inc. Exhibit 1003 Page 1490

What follows from this assumption is that asymmetric packet rates are an indication of unusual behavior.

Routes through MULTOPS-equipped routers are symmetric and stable In this context, a symmetric route means that if packets going from A to B go through router R, then packets from B to A also go through R. If this is not true, then—obviously—a router cannot detect asymmetries in packet rates. Besides being symmetric, routes are assumed to be stable, i.e., routes do not change more than once every few minutes.

The details of keeping track of packet rates will be discussed shortly.

3.3 Data structure

A MULTOPS is a tree of nodes. Every node has 256 records with 2 fields to keep track of aggregate packet rates to and from subnets within a certain subnet. The tree is at most 4 levels deep. Level 0 contains the root node only. The root node contains aggregate packet rates to and from 8-bit prefix subnets (0.0.0.0/8, 1.0.0.0/8, ..., 255.0.0.0/8). A node on level 1 keeps track of aggregate packet rates to and from 16-bit prefix subnets. A node in level 2 contains aggregate packet rates to and from 24-bit prefix subnets. Level 3 contains packet rates to and from single hosts. Figure 3-2 shows a MULTOPS. The node labeled "N130" keeps track of aggregate packet rates to and from subnets 130.0.0.0/16, 130.1.0.0/16, ..., 130.255.0.0/16. Node "N130.37.24" keeps track of packets rates to and from single hosts within subnet 130.37.24.0/24. See Appendix B for an explanation of this notation.

Nodes are created and destroyed in real time to reflect changes in packet rates. When the aggregate packet rate to or from subnet S exceeds a certain threshold, a child node is created under the relevant record to keep track of packet rates to and from all subnets within S. This "**unfolding**" can continue down to level 3 in the tree. When packet rates go down, relevant parts of the tree are destroyed ("**folding**") to minimize memory consumption and to avoid the overhead of updating packet rates that are not of interest. Folding and unfolding are the essential mechanisms in a MULTOPS.

3.4 Algorithm

When a router gets a packet, it extracts the first byte of the source address of the packet and uses this value as an index to find the appropriate record in which it updates the aggregate packet rate **from** all addresses with the same prefix. When a packet passes through the router in opposite direction, the router extracts the first byte of the **destination** address to find the appropriate record in which it updates the aggregate packet rate **to** all addresses with the same prefix. If either rate exceeds a certain threshold, a pointer in that record points to a child (or, if not, a child is created) that keeps track of all subnets that share the same prefix. The second byte from the address is used as an index in this child node to find and update the aggregate packet rate. This process can continue down to level 3 in the tree (see Figure 3-2).

Cloudflare, Inc. Exhibit 1003 Page 1491

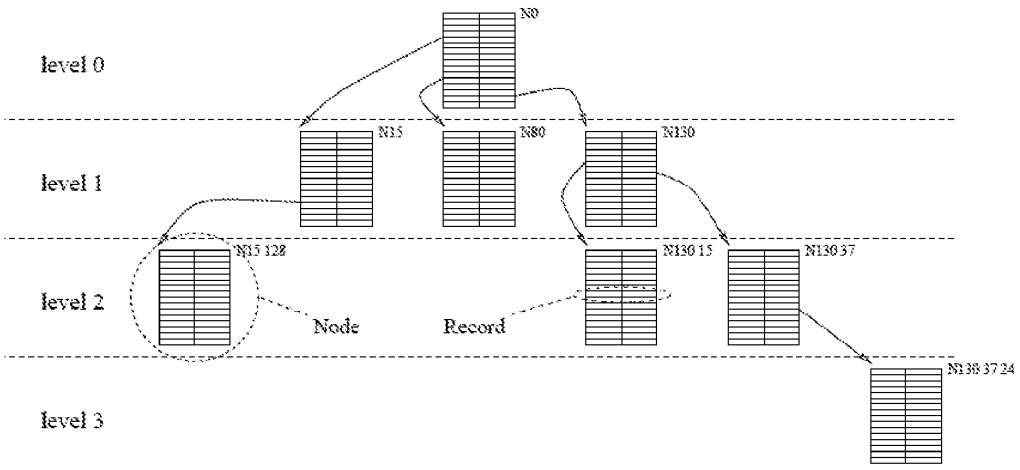


Figure 3-2: MULTOPS

A MULTOPS keeps track of packet rates to **and** from subnets. This requires a notion of direction. A router on the edge of subnet D can act in two different ways for each packet it is forwarding **to** D: (1) it always updates the numbers that keep track of packet rates **to** D, or (2) it always updates the numbers that keep track of packet rates **from** the source host (or source subnet). A router that uses method 1 is in **outside protection** mode. A router that uses method 2 is in **inside protection** mode.

The difference between inside and outside protection mode is important when the source(s) of the attack need to be determined. This is best explained using the following example. If MIT's router is running in outside protection mode, the router uses the destination address of outgoing (i.e., coming from MIT's network) packets and the source address for incoming (i.e., going to MIT's network) packets to search and update packet rates in the MULTOPS. In this case, packet rates for subnets with prefix 18.0.0.0/8 are all zero because no machine with this prefix exists outside of MIT's network. If MIT's edge router is running in inside protection mode, the router uses the destination address of incoming packets and the source address of outgoing packets. Packet rates for subnets with prefix 18.0.0.0/8 will be the only non zero values in the MULTOPS. In this example, when a bandwidth attack is launched **from** MIT's network, a MULTOPS running in outside protection mode is unable to determine the source(s) of the attack because it only knows the address(es) of the victim(s). Similarly, when a bandwidth attack **on** MIT is launched, a MULTOPS running in inside protection mode is unable to determine the source(s) of the attack.

In either case, each packet causes at least 1 update (in the root node) and at most 4 updates (when the tree is unfolded to level 3 for that IP address). The complete algorithm is given in Figure 3-3. A router runs `per_packet_execute` for each packet it receives on one of its inputs. Unless specified otherwise, we assume that MULTOPS-equipped routers operate in inside protection mode.

Cloudflare, Inc. Exhibit 1003 Page 1492

```
#define INSIDE_PROTECTION  0
#define OUTSIDE_PROTECTION 1
#define TO   0
#define FROM 1
extern int threshold;   /* Threshold for unfolding                    */
extern int mode;        /* INSIDE_PROTECTION or OUTSIDE_PROTECTION  */
extern Subnet S;        /* S is some subnet close to this router    */
extern Node *root;      /* root of MULTOPS                            */

per_packet_execute(IPpacket *packet)
{
  if(mode == INSIDE_PROTECTION) {
    if(packet is heading for S)
      update_rate(FROM, packet->source_addr);
    else
      update_rate(TO, packet->dest_addr);
  } else { /* i.e., mode == OUTSIDE_PROTECTION */
    if(packet is heading for S)
      update_rate(TO, packet->dest_addr);
    else
      update_rate(FROM, packet->source_addr);
  }
  if(folding necessary)
    fold();
}

update_rate(int direction, IPaddress ip_address)
{
  Node *node = root;
  Record *record;
  for(i = 0; i <= 3; i++) {
    int n = ip_address[i];      /* i.e., i-th byte from ip_address */
    record = node->record[n];   /* i.e., n-th record from node     */
    update rate[direction] in record /* i.e., either TO or FROM rate up-
dated */
    if(rate[direction] > threshold)
      create child node;
    if(no child node) /* descend? */
      break;
    else
      node = record->child_node;
  }
}
```

Figure 3-3: Algorithm for MULTOPS

Cloudflare, Inc. Exhibit 1003 Page 1493

3.5 Rates

Let $F(S)$ be the aggregate packet rate from subnet S and let $T(S)$ be the aggregate packet rate to S. R is defined as:

$$R(S) = \frac{T(S)}{F(S)}$$

In words: $R(S)$ is the ratio between the aggregate packet rate to subnet S and the aggregate packet rate from S, as measured by a router. Because TCP acknowledges the receipt of every k packets by sending back an ACK packet, R will be close to k for all subnets, i.e., packet rates are symmetric. A router in inner-protection mode can suspect a bandwidth attack **from** one or more machines within subnet S if R is too small. It can suspect a bandwidth attack **on** one or more machines within subnet S if R is too big. These thresholds, which we refer to as R_{min} and R_{max} respectively, influence the router's sensitivity to detect bandwidth attacks: if R_{min} is too small or if R_{max} is too big, attacks might remain undetected. If R_{min} is too big or if R_{max} is too small, false positives might result.

3.6 Unfolding

A MULTOPS keeps track of packet rates on different aggregation levels. If $T(S)$ or $F(S)$ in a certain record r exceed threshold Q, the MULTOPS unfolds r. Record r is unfolded by creating a child node that keeps track of aggregate packet rates to and from subnets within S. In Figure 3-2, for example, the record with index 130 in the root node is unfolded into node N130. A record in a node in level 3 of the tree cannot be unfolded.

3.7 Example

Suppose MIT's edge router has two interfaces: `eth0` for all packets coming from the outside world to MIT's network, and `eth1` for packets in the opposite direction. Assume that this router is running in outside protection mode. Initially, the MULTOPS contains one node only: the root node.

A packet with source address 18.24.16.27 and destination address 130.37.24.4 arrives on `eth1`. The first byte of the destination address is 130, so record 130 is fetched from the root node and the aggregate packet rate **to** all subnets with prefix 130.0.0.0/8 is updated. 200ms later, a packet with source address 130.37.24.4 and destination address 18.24.16.27 arrives on `eth0`. The first byte of the source address is 130, so record 130 is (again) fetched from the root node and the aggregate packet rate **from** all subnets with prefix 130.0.0.0/8 is updated. If communication proceeds normally, the aggregate packet rates to and from subnet 130.0.0.0/8 will be roughly equal.

A few minutes later, a bandwidth attack is launched on 130.37.24.1, 130.37.24.2, and 130.37.24.3 from a dozen or so machines on MIT's network. The observed packet rate to subnet 130.0.0.0/8 quickly exceeds

Cloudflare, Inc. Exhibit 1003 Page 1494