UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

World Programming, Ltd.
Petitioner

v.

SAS Institute, Ltd.
Patent Owner

Case No. Unassigned
Patent 7,447,686

DECLARATION OF SYLVIA D. HALL-ELLIS, PH.D.

I. INTRODUCTION

1. My name is Sylvia D. Hall-Ellis. I have been retained as an expert by World Programming, Limited ("WPL").

2. I have written this report at the request of World Programming, Limited to provide my expert opinion regarding the authenticity and public availability of two books and a journal article. My report sets forth my opinions in detail and provides the basis for my opinions regarding the public availability of these publications.

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

4. 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 litigation. I have no other interests in this proceeding or with any of the parties.

5. All of the materials that I considered are discussed explicitly in this declaration.

II. QUALIFICATIONS

6. I am currently an Adjunct Professor in the School of Information at San José State University. I obtained a Masters 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 forty-five 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.

7. I am a member of the American Library Association (ALA) and its Association for Library Collections & Technical Services (ALCTS) Division, and I 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 and as a member of the Editorial Board for the ALCTS premier cataloging journal, *Library Resources and Technical Services*.

8. 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.

9. I have been deposed nine times: (1) *Symantec Corp. vs. Finjan, Inc.*, Petition for *Inter Partes Review* of U.S. Patent No. 7,613,926, May 26, 2016, on behalf of Symantec Corp.; (2) *Symantec Corp. vs. Finjan, Inc.*, 14-cv-299-HSG (N.D. Cal.), on behalf of Symantec Corp., September 14, 2017; (3) one deposition for ten matters: *Intellectual Ventures I LLC vs. AT&T Mobility LLC*; AT&T Mobility II LLC, New Cingular Wireless Services, Inc., SBC Internet Services, Inc., Wayport, Inc., and Cricket Wireless LLC, C.A. No. 12-193 (LPS); *Intellectual Ventures II LLC vs. AT&T Mobility LLC*; AT&T Mobility II LLC, New Cingular Wireless Services, Inc., SBC Internet Services, Inc., Wayport, Inc., and Cricket Wireless LLC, C.A. No. 13-1631 (LPS); *Intellectual Ventures I LLC vs. T-Mobile USA, Inc. and T-Mobile US, Inc.*, C.A. No. 13-1632 (LPS); *Intellectual Ventures II LLC vs. T-Mobile USA, Inc. and T-Mobile US, Inc.*, C.A. No. 13-1633 (LPS); *Intellectual Ventures I LLC, vs. Nextel Operations, Inc., Sprint Spectrum L.P., Boost Mobile, LLC and Virgin Mobile USA, L.P.*, C.A. No. 13-1634 (LPS); *Intellectual Ventures II LLC vs. Nextel Operations, Inc., Sprint Spectrum L.P., Boost Mobile, LLC and Virgin Mobile USA, L.P.*, C.A. No. 13-1635 (LPS); *Intellectual Ventures I LLC, vs. United States Cellular Corporation*, C.A. No. 13-1636 (LPS); *Intellectual Ventures I LLC vs. United States Cellular Corporation*,

C.A. No. 13-1637 (LPS); *Intellectual Ventures II LLC vs. AT&T Mobility LLC, AT&T Mobility II LLC, New Cingular Wireless Services, Inc.,* C.A. No. 15-799 (LPS); *Intellectual Ventures I LLC vs. T-Mobile USA, Inc. and T-Mobile US, Inc.,* C.A. No. 15-800 (LPS), on behalf of AT&T Mobility LLC; AT&T Mobility II LLC, Boost Mobile, LLC Cricket Wireless LLC, Nextel Operations, Inc., New Cingular Wireless Services, Inc., SBC Internet Services, Inc., Sprint Spectrum L.P., T-Mobile USA, Inc., T-Mobile US, Inc., United States Cellular Corporation Virgin Mobile USA, L.P., and Wayport, Inc., November 15, 2016; (4) *Hitachi Maxell, LTD., v. Top Victory Electronics (Taiwan) Co. Ltd., et al.,* 2:14-cv-1121 JRG-RSP (E.D. Texas), on behalf of Top Victory Electronics (Taiwan) Co. LTD, *et. al.,* January 20, 2016; (5) *Sprint Spectrum, L.P. vs. General Access Solutions, Ltd*., Petition for *Inter Partes Review* of U.S. Patent No. 7,173,916, on behalf of Sprint Spectrum L.P., July 13, 2018; (6) *Nichia Corporation vs. Vizio, Inc*., 8:16-cv-00545; on behalf of Vizio, Inc., October 12, 2018; (7) *Intellectual Ventures I LLC, vs. T-Mobile USA, Inc., T-Mobile US, Inc., Ericsson Inc., and Telefonaktiebolaget LM Ericsson,* 2:17-cv-00557 (JRG), on behalf of T-Mobile USA, Inc., T-Mobile US, Inc., Ericsson Inc., and Telefonaktiebolaget LM Ericsson, October 19, 2018; (8) Pfizer, Inc. vs. Biogen, Inc., Petition for *Inter Partes Review* of U.S. Patent No. 8,821,873, on behalf of Pfizer, November 3, 2018; and, (9) Finjan, Inc. vs. ESET, LLC and ESET SPOL. S.R.O., 3:17-cv-

00183-CAB-BGS, on behalf of ESET, January 15, 2019.

10. My full curriculum vitae is attached hereto as Attachment 8.

III. PRELIMINARIES

11. *Scope of this declaration.* I am not an attorney and will not offer opinions on the law. I am, 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 November 22, 2002.

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.

15. *Persons of ordinary skill in the art*. I am told by counsel that the subject matter of this proceeding relates to systems and methods for handling database statements.

16. I have been informed by counsel that a "person of ordinary skill in the art at the time of the inventions" is a hypothetical person who is presumed to be

familiar with the relevant field and its literature at the time of the inventions. This hypothetical person is also a person of ordinary creativity, capable of understanding the scientific principles applicable to the pertinent field.

17. I am told by counsel that persons of ordinary skill in this subject matter or art would have had at least an undergraduate degree or equivalent in computer science, or a master's degree in information science, and would also have two or three years of engineering, development, or systems administration experience related to database systems.

18. It is my opinion that such a person would have been engaged in research, learning, study, and practice in the field, and possibly formal instruction so that bibliographic resources relevant to his or her research would be familiar. In the 1980s and 1990s such a person would have had access to a vast array of long-established print resources in electrical engineering as well as to a rich set of online resources providing indexing information, abstracts, and full text services for computer science.

IV. LIBRARY PROFESSIONAL PRACTICES

19. In preparing this report, I used authoritative databases, such as the OCLC bibliographic database, the Library of Congress Online Catalog, and the Library of Congress Subject Authorities, to confirm citation details of the various publications 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 serial publication discussed below.

20. *U.S. Copyright Office*. Created by Congress in 1897, the Copyright Office is responsible for administering a complex and dynamic set of laws, which include registration, the recordation of title and licenses, a number of statutory licensing provisions, and other aspects of the *1976 Copyright Act* and the *1998 Digital Millennium Copyright Act*. The public catalog in the Copyright Office includes information filed since 1978 (https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?DB=local&PAGE=First). Individuals can search by title, personal or corporate name, key word, registration number, and document number. Works filed before 1978 can be located through the Copyright Public Records Reading Room (https://www.copyright.gov/circs/circ23.pdf). A researcher can find the date on which an item was published and deposited for copyright.

21. *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 indexing and library date information. However,

date information for indexing entries is often unavailable. This is especially true for online indices.

22. Indexing services use a wide variety of controlled vocabularies to provide subject access and other means of discovering the content of documents. The Library of Congress Subject Authorities includes standard forms of terms and cross references that are included in bibliographic records. The formats in which these access terms are presented vary from service to service.

23. Online indexing services 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. Prominent indexing services include *Scopus* and the *IEEE Xplore* database.

24. Prominent indexing services include the following:

a. *Scopus*. Produced by Elsevier, a major publisher, Scopus is the largest database of abstracts and citations of peer-reviewed literature. Its scope includes the social sciences, science, technology, medicine, and the arts. It includes 60 million records from more than 21,500 titles from some 5,000 international publishers. Coverage includes 360 trade publications, over 530 book series, more than 7.2 million conference papers, and 116,000 books. Records date from 1823.

b. *IEEE Xplore*. This scholarly research database includes indexes, abstracts, and full-text for articles and papers on computer science, electrical engineering, and electronics. The database mainly covers material from the Institute of Electrical and Electronics Engineers (IEEE) and the Institution of Engineering and Technology. The IEEE *Xplore* digital library provides Web access to more than 4.3-million full-text documents from some of the world's most highly cited publications. The content comprises over 180 journals, over 1,400 conference proceedings, more than 3,800 technical standards, over 1,800 eBooks and over 400 educational courses. Approximately 20,000 new documents are added to IEEE Xplore each month. Abstracts are free to access, but access to full text requires a subscription or institutional login.

c. *ACM Digital Library*. This index is produced by the Association for Computing Machinery, the world's largest scientific and educational computing society. AMC Digital Library contains the full text of all AMC publications, hosted full-text publications from selected publishers, and the ACM Guide to Computing Literature—a comprehensive bibliography of computing literature beginning in the 1950s with more than a million entries. All metadata in the database are freely available on the Web, including abstracts, linked references, citing work, and usage statistics. Full-text articles are available with subscription.

V. LIBRARY CATALOGING PRACTICES

25. 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.[1] MARC was first developed in the 1960's by the Library of Congress. A MARC-compatible library is one that has a catalog consisting of individual MARC records for each of its items. Today, MARC is the primary communications protocol for the transfer and storage of bibliographic metadata in libraries.[2]

---

[1] The full text of the standard is available from the Library of Congress at http://www.loc.gov/marc/bibliographic/ (last visited July 6, 2019).

[2] 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 July 6, 2019) ("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 forty 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

26. 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" 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

---

library catalogs used today."). MARC is the ANSI/NISO Z39.2-1994 (reaffirmed 2016) standard for Information Interchange Format.

[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 July 6, 2019).

> 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 July 6, 2019).

27. 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 July 6, 2019) ("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 (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 July 6, 2019).

28. 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. For example, a work's title is recorded in Field 245; the primary author of the work is transcribed in Field 100; an item's International Standard Book Number ("ISBN") consisting of ten or thirteen digits is transcribed in Field 020; an item's International Standard Serial Number ("ISSN") is transcribed in Field 022; the Library of Congress classification notation is recorded in Field 050; and, the publication date is recorded in Field 260 under the subfield "c." 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.[4]

29. The library that created the record is recorded in Field 040 in subfield "a" with a unique library code. When viewing the MARC record online via Online Computer Library Center's ("OCLC") bibliographic database, hovering over this code with the mouse reveals the full name of the library. I used this method of "mousing over" the library codes in the OCLC database to identify the originating library for the MARC records discussed in this report. Where this "mouse over" option was not available, I consulted the Directory of OCLC Libraries in order to identify the institution that created the MARC record.[5]

---

[4] http://www.loc.gov/marc/bibliographic/bd3xx.html .
[5] https://www.oclc.org/en/contacts/libraries.html .

30. MARC records also include several fields that include subject matter classification information. An overview of MARC record fields is available through the Library of Congress.[6] For example, 6XX fields are termed "Subject Access Fields."[7] Among these, for example, is the 650 field; this is the "Subject Added Entry – Topical Term" field.[8] 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.* Further, MARC records include call numbers, which themselves include a classification number. For example, the 050 field is the "Library of Congress Call Number."[9] 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, included in the 050 field is a subject matter classification. 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

---

[6] http://www.loc.gov/marc/bibliographic/.
[7] http://www.loc.gov/marc/bibliographic/bd6xx.html.
[8] http://www.loc.gov/marc/bibliographic/bd650.html.
[9] http://www.loc.gov/marc/bibliographic/bd050.html.

Decimal Classification scheme) and uses it consistently. When the Library of Congress assigns the classification number, it appears as part of the 050 field. If a local library assigns the classification number, it appears in a 090 field. In either scenario, the MARC record includes a classification number that represents a subject matter classification.

31. 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."[10] Among other services, OCLC and its members are responsible for maintaining the WorldCat database (http://www.worldcat.org/), used by independent and institutional libraries throughout the world.

32. OCLC also provides its members online access to MARC records through its OCLC bibliographic database. When an OCLC member institution

---

[10] 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).

acquires a work, it creates a MARC record for this work 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 institution or is tape-loaded from the Library of Congress, the MARC record is then made available to any other OCLC members online, and therefore made available to the public. Accordingly, once the MARC record is created by a cataloger at an OCLC member institution or is tape-loaded from the Library of Congress or another library anywhere in the world, 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.

33. When an OCLC member institution creates a new MARC record, OCLC automatically supplies the date of creation for that record. The date of creation for the MARC record appears in the fixed Field (008), characters 00 through 05. The MARC record creation date reflects the date on which, or shortly after which, the item was first acquired or cataloged. Initially, Field 005 of the MARC record is automatically populated with the date the MARC record was

created in year, month, day format (YYYYMMDD) (some of the newer library catalog systems also include hour, minute, second (HHMMSS)). 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). Field 005 is visible when viewing a MARC record via an appropriate computerized interface, but when a MARC record is printed to hardcopy, no "005" label appears. The initial Field 005 date (*i.e.*, the date the MARC record was created) does appear, however, next to the label "Entered."[11] The date upon which the most recent update to Field 005 occurred at a local library appears next to the label "Replaced." Thus, when an item's MARC record has been printed to hardcopy—as is the case with the exhibits to this report—the date reflected next to the label "Entered" is necessarily on or after the date the library first cataloged and indexed the underlying item.

34. Once one library has cataloged and indexed a publication by creating a MARC record for that publication, other libraries that receive the publication do not need to create additional MARC records—the other libraries rely on the original MARC record. They may update or revise the MARC record to ensure

---

[11] In this report, I sometimes refer to the "Entered" entry as Field 008, characters 00-05.

accuracy, but they do not replace or duplicate it. This practice does more than save libraries from duplicating labor. It also enhances the accuracy of MARC records. Further, it allows librarians around the world to know that a particular MARC record is authoritative (in contrast, a hypothetical system wherein duplicative records were created would result in confusion as to which record is authoritative).

35. In addition, a local library may use a field in the 900-999 range to record local information regarding the cataloging, binding, storage, and status of an item. The 9XX fields are not part of the standard MARC 21 format.[12] For example, the Library of Congress uses the Field 955 to track the movements of an item through the cataloging process. The date and initials of the staff member performing a task is recorded, including (but not limited to) transfer to the Dewey Decimal team for the assignment of a classification number, the verification of the CIP (cataloging in publication) data, and transfer to the BCCD (Binding and Collection Care Division) for binding and preservation services. Data transcribed at local libraries that are important to establish dates of public availability are described for each of the documents described in this report.

36. Catalogers can create MARC records for all types of print, online, and digital resources. The date of creation of the MARC record by a cataloger at an OCLC member institution generally reflects when the underlying item is accessible

---

12 https://www.oclc.org/bibformats/en/9xx.html.

to the public. 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.

37. MARC records cover serial publications, including both serially published monographs and journals. OCLC hosts MARC records for more than 320 million serial publications. Serial publications are those publications that have the same collective title but are intended to be continued indefinitely with enumeration such as a volume or issue number (*e.g.*, magazines, journals, etc.). In the OCLC bibliographic database, the first issue of the serial publication is typically cataloged (*i.e.*, a corresponding MARC record is created), but the date is left open-ended with the use of a punctuation mark such as a dash. OCLC serial publication MARC records represent the entire run of the serial title. With knowledge of the first issue published, future issues can be predicted based on the

information provided in the MARC record, for example in Field 362. In my extensive professional experience, is it highly unusual for a library to stop collecting and shelving a serial publication prior to the end of its publication run. If a subscription to a serial publication ends its run or is cancelled before the end of its run, the library will denote that it has stopped receiving new volumes by filling in the end date in the MARC record.

38. The handling of printed journal subscriptions is shown on the covers of individual issues. As was the best practice among libraries, issues arrived at a central facility and were immediately received, verified as part of a subscription, checked in, and stamped with the institution's name and date. Determining that the issue was part of the library subscription ensured that the entire set of publications for the year had been received so that they could be professionally bound and retained. This process also verified that all of the published issues arrived so that the library staff did not have to request or claim an issue that did not arrive as expected. In large public libraries with branches and multi-campus libraries within academic institutions, the journals were sorted and delivered to the subscribing unit. The issues were frequently stamped again to acknowledge receipt. The new issue was placed in the public area; the older issue was stored so that it remained available.

39. The foregoing process has been standard library practice longer than I

have been working in the profession. I first learned the steps in the process in the late 1970s and later supervised it. Although the checking in process has become automated and now links electronically to holdings records for the MARC record for each serial title, the manual stamping and placing the issue in a public area has not changed for 50 years. Unless I note otherwise below in reference to a specific serial publication, it is my expert opinion that this standard protocol was followed for each of the serial publications discussed below.

VI. PUBLICATION 1 -- EXHIBIT 1006 ("TEMPLETON")

40. Exhibit 1006 is a copy of Volume 4, Issue 2 of *The VLDB Journal* found in the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder. The article "InterViso: Dealing with the Complexity of Federated Database Access" by Marjorie Templeton, Herbert Henley, Edward Maros, and Darrel J. Van Buer (hereafter "Templeton") appears beginning on page 287 of this issue dated April 1995. Exhibit 1006 is a true and correct copy of the Templeton article that I understand is being submitted as an exhibit in this proceeding. Exhibit 1006 is a true and correct copy of the issue cover, table of contents, and the article (pages 287-317). I obtained this copy of the article from the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder. Specifically, the text of the article 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 1006 was found within the custody of a library – a place where, if authentic, a copy of this journal would likely be. Exhibit 1006 is a true and correct copy in a condition that creates no suspicion about authenticity.

41. The cover of the April 1995 issue of *The VLDB Journal* has a stamp affixed by or on behalf of the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder which shows that it was received, verified, and checked in on May 30, 1995. Therefore, in my experience, this issue of *The VLDB Journal* would have been available to users at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder on or shortly after that date.

42. Attached hereto as Attachment 1a is a true and correct copy of the MARC record for *The VLDB Journal* at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder. The library ownership is indicated by the presence of the library's code (DVD) in the 049 field. I personally identified and retrieved the MARC record that is Attachment 1a. Attachment 1a also shows that Exhibit 1006 was catalogued with descriptor terms "Database management ‡v Periodicals," "Databases ‡v Periodicals," and "Data structures (Computer science) ‡v Periodicals" in the 650 fields.

43. Based on finding a print copy of Exhibit 1006 in the Gemmill Library

of Engineering, Mathematics and Physics at the University of Colorado – Boulder and MARC record in its online library catalog attached as Attachment 1a, it is my opinion that the article "InterViso: Dealing with the Complexity of Federated Database Access" by Templeton, *et. al.* published in *The VLDB Journal* was available at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder on or shortly after May 30, 1995.

44. As noted in the holdings information (field 362), the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder has received *The VLDB Journal* since July 1992 and continues to receive the publication. In view of the MARC record for Exhibit 1006, the Templeton article was publicly available on or shortly after May 30, 1995, because the serial title had been cataloged and indexed at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder and made part of its catalog.

45. Attached hereto as Attachment 1b is a true and correct copy of the MARC record for *The VLDB Journal* 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 "NSD," which means that the MARC record for this serial was cataloged as part of the National Serials Data Program at the Library of Congress. As can be seen in

the "Entered" field in the MARC record for this exhibit, a cataloger at the Library of Congress created OCLC record number 27078023 on December 7, 1992. The "BLvl" entry in Attachment 1b is "s," which indicates that *The VLDB Journal* is a serial publication. Field 310 of Attachment 1b reads "Quarterly." Accordingly, the MARC record for Exhibit 1006 corresponds to those issues of *The VLDB Journal* from the time of the serial title began publication.

46. Attachment 1b further includes an entry in field 050 ("QA76.9.D3 ‡b V53")—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal System). Attachment 1b further includes an entry in field 082 ("004"), a subject matter consistent with the Dewey Decimal classification system. Attachment 1b further includes three English language field 650 entries reading "Database management ‡v Periodicals" (see Attachment 1c, Library of Congress subject heading sh2008102046), "Data structures (Computer science) ‡v Periodicals" (see Attachment 1d, Library of Congress subject heading sh85035862), and "Databases ‡v Periodicals" (see Attachment 1e, Library of Congress subject heading sh86007767). "Periodicals" is a form type subdivision that can be added to subject headings as appropriate (see Attachment 1f, Library of Congress subject heading sh85099890). 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 December 7, 1992, 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 subscribed to the serial, which means that the corresponding publication was publicly available on or shortly after that same date through any library with access to the OCLC bibliographic database or through an individual library.

47. Further supporting my opinion that Exhibit 1006 was publicly accessible to ordinarily skilled and interested researchers in the field is the fact that *The VLDB Journal* was included in several well-known indices, including *Science Citation Index, Science Citation Index Expanded (SciSearch), Journal Citation Reports/Science Edition, SCOPUS, Google Scholar, ACM Digital Library, Current Contents/Engineering, Computing and Technology, DBLP, EBSCO Applied Science & Technology Source, EBSCO Computers & Applied Sciences Complete, EBSCO Discovery Service, EBSCO STM Source, EBSCO TOC Premier, EI Compendex, Gale, Gale Academic OneFile, Gale InfoTrac, OCLC WorldCat Discovery Service, ProQuest Advanced Technologies & Aerospace Database, ProQuest SciTech Premium Collection, ProQuest Technology Collection,*

ProQuest-ExLibris Primo, and *ProQuest-ExLibris Summon*.[13] As noted above, this would have provided a person of ordinary skill in the art with keyword searching capabilities and other tools to locate this article.

48. Attachment 1b indicates that *The VLDB Journal* as cataloged at the Library of Congress is currently available from 126 libraries, as indicated by the number of other holdings stated at the top of the MARC record. In view of the above, this issue of *The VLDB Journal* was publicly available on or shortly after May 30, 1995, because by that date it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder. For these reasons, it is my opinion that Exhibit 1006 was published and accessible to the public on or shortly after May 30, 1995.

49. Not only was Exhibit 1006 accessible and available to others in the field as early as May 1995, researchers actually obtained and cited this article in other works, which is still further confirmation of its availability and accessibility. For example, Google Scholar indicates that this article has been cited 56 times (see Attachment 1g). Scopus reports that this article has been cited 19 times (see Attachment 1h). The following are exemplary articles that cite Exhibit 1006:

---

13 https://www.springer.com/computer/database+management+%26+information+retrieval/journal/778.

- Urhan, Tolga, Michael J. Franklin, and Laurent Amsaleg. "Cost-based Query Scrambling for Initial Delays." In *SIGMOD '98 Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data* (pp. 130-141). New York: ACM, c1998. (see Attachment 1i);
- Amsaleg, Laurent, Anthony Tomasic, Michael J. Franklin, and Tolga Urhan. "Scrambling Query Plans to Cope with Unexpected Delays." In *Proceedings of the Fourth International Conference on Parallel and Distributed Information Systems* (pp. 208-219). Los Alamitos, CA: IEEE Computer Society Press, c1996. (see Attachment 1j).

50. Therefore, Templeton (Exhibit 1006) was publicly accessible as early as May 30, 1995, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, received at the Gemmill Library of Engineering, Mathematics and Physics at the University of Colorado – Boulder, and found and cited by others.

VII. PUBLICATION 2 -- EXHIBIT 1007 ("NUTSHELL")

51. Exhibit 1007 is a copy of the book, *SQL In a Nutshell: A Desktop Quick Reference*, 1st edition, by Kevin Kline and Daniel T. Kline issued by O'Reilly & Associates, Inc. with a 2001 copyright. (hereinafter called "Nutshell"). The exhibit filed in this proceeding as Exhibit 1007 is a true and correct copy of

the title page, verso of the title page, table of contents, preface, Chapter 1 (pages 1–8), Chapter 2 (pages 9–26), a portion of Chapter 3 (pages 27–30, 46, 118-119, 132-139, 146-151), Chapter 4 (pages 163-193), Chapter 5 (pages 194-196), and the Appendix (pages 197-203). I obtained this copy of the book via counsel from the University of Texas at Austin Library. Specifically, the text of the book 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 1007 was found within the custody of a library – a place where, if authentic, a copy of this volume would likely be. Exhibit 1007 is a true and correct copy in a condition that creates no suspicion about its authenticity.

52. Attached hereto as Attachment 2a is a true and correct copy of the MARC record for this book from the University of Texas at Austin Library. The library ownership is indicated by the presence of the library's code (IXA) in the 049 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 2a.

53. Based on finding a copy of Exhibit 1007 in the University of Texas at Austin Library and a record in its online library catalog attached as Attachment 2a, it is my opinion that the book, *SQL In a Nutshell: A Desktop Quick Reference,* 1st edition, by Kline and Kline was cataloged on November 6, 2000, as shown in Field

008 ("001106"). The most recent modification of the MARC record occurred on February 27, 2001, as indicated in field 005 ("20010227"). Exhibit 1007 includes three pages of date due slips that provide information related to the Nutshell book which was checked out at least 16 times. The earliest date due shown for the Nutshell book was June 6, 2001, which indicates it was publicly available, found, and checked out before that date.

54. Attached hereto as Attachment 2b is a true and correct copy of the MARC record for the book, *SQL In a Nutshell: A Desktop Quick Reference,* 1st edition, by Kline and Kline obtained from the OCLC bibliographic database. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Exhibit 1007, that library code is "DLC," which means that the MARC record for this book was cataloged at the Library of Congress. As can be seen in the "Entered" field in MARC record for this exhibit, a cataloger created OCLC record number 45316889 on November 6, 2000. Attachment 2b further includes an entry in field 050 ("QA76.73.S67 ‡b K55 2001")[14] —as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system). Attachment 2b further includes an entry in field

[14] Mathematics—Instruments and machines—Calculating machines—Electronic computers. Computer science—Digital computers—Programming languages—Individual languages, A-Z—SQL. MySQL

082 (“005.75/85”), [15] a subject matter consistent with the Dewey Decimal classification system. Attachment 2b further includes two English language field 650 entries reading “SQL (Computer program language)” (see Attachment 2c, Library of Congress subject heading sh86006628) and “Client/server computing” (see Attachment 2d, Library of Congress subject heading sh93000502), in addition to a field 630 (uniform title) “SQL server” (see Attachment 2e, Library of Congress subject heading n90684343). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 2b was indexed according to its subject matter by virtue of at least four independently sufficient classifications: the field 050 entry, the field 082 entry, the field 630 entry, and the field 650 entries. Further, as of November 6, 2000, the MARC record attached hereto as Attachment 2b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the Nutshell book.

55. Attachment 2b indicates that the book, *SQL In a Nutshell: A Desktop Quick Reference,* 1st edition, by Kline and Kline, as cataloged at the Library of Congress is currently available from 335 libraries, as indicated by the number of other holdings stated at the top of the MARC record. The Nutshell book has been

---

[15] 005.75 Specific types of data files and databases

available at academic and public libraries across the United States since its publication. The Nutshell book was cataloged and available at the following representative institutions: Harvard University Library, University of Michigan Library, and, the University of California at San Diego Library.

56. Attached hereto as Attachment 2f is a true and correct copy of the MARC record for this monograph, which is Publication 2, at the Harvard University Library. The library ownership is indicated by the presence of the library's local code (HVD) in the INST field. I personally identified and retrieved the library catalog record which is Attachment 2f. The MARC record is based on OCLC record number 45316889 as shown in the second field 035. Attachment 2f indicates that the Nutshell book was cataloged at the Harvard University Library as of November 6, 2000, as shown in field 008 ("001106").

57. Attached hereto as Attachment 2g is a true and correct copy of the MARC record for this monograph, which is Publication 2, in the University of Michigan Library. The library ownership is indicated by the presence of the library's code (EYM) in subfield d of the 040 field. I personally identified and retrieved the library catalog record which is Attachment 2g. The MARC record is based on OCLC record number 45316889 as shown in the first field 035. Attachment 2g indicates that the Nutshell book was cataloged in the University of Michigan Library as of March 5, 2001, as shown in field 008 ("010305") and thus

was publicly available on or shortly after that date at the University of Michigan Library.

58. Attached hereto as Attachment 2h is a true and correct copy of the MARC record for this monograph, which is Publication 2, at the University of California San Diego Library. The library ownership is indicated by the presence of the library's code (CUS) in the second 910 field. I personally identified and retrieved the library catalog record which is Attachment 2h. The MARC record is based on OCLC record number 45316889 as shown in field 035. Attachment 2h indicates that the Nutshell book was cataloged by the University of California San Diego Library as of November 6, 2000, as shown in field 008 ("001106").

59. Attached hereto as Attachment 2i is a true and correct copy of the MARC record for this monograph, which is Publication 2, at the Library of Congress. I personally identified and retrieved the library catalog record which is Attachment 2i. The 955 field indicates that Nutshell was received and sent to Dewey on November 6, 2000, as indicated by those fields with "11-06-00" date codes. Those dates match the date that the Nutshell book was cataloged, as shown in field 008 ("001106"). Final CIP verification was performed on June 14, 2001, and Nutshell was sent to BCCD (Binding and Collection Care Division) on June 18, 2001. The data in the 955 fields for the Library of Congress MARC record indicates that Nutshell was cataloged on November 6, 2000 and that final

processing occurred on June 18, 2001, and the book was publicly available on or shortly after that date at the Library of Congress. Any person can enter the Library of Congress and request a book be brought to a reading room.

60. Not only was Exhibit 1007 accessible and available to others in the field as early as November 6, 2000, researchers actually obtained and cited this book in other works, which is still further confirmation of its availability and accessibility. For example, Google Scholar indicates that this book has been cited a number of times. Attachment 2j reports the number of citing references shown in Google Scholar for Exhibit 1007.

61. Therefore, the Nutshell book (Exhibit 1007) was publicly accessible as early as November 6, 2000, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received at and checked out from the University of Texas at Austin Library.

VIII. PUBLICATION 3 -- EXHIBIT 1010 ("GROFF")

62. Exhibit 1010 is a copy of the book, *SQL: The Complete Reference*, 1st edition, by James R. Groff and Paul N. Weinberg, and issued by Osborne/McGraw-Hill in 1999 (hereinafter called "Groff"). The exhibit filed in this proceeding as Exhibit 1010 is a true and correct copy of the title page, verso of the title page, table of contents, and Chapter 1 (pages 3–12), Chapter 2 (pages 13–

22), a portion of Chapter 3 (pages 23–40), a portion of Chapter 4 (pages 53–61), a portion of Chapter 5 (pages 68–88), a portion of Chapter 6 (pages 92–125), a portion of Chapter 7 (pages 164–178), a portion of Chapter 13 (pages 402–406), a portion of Chapter 21 (pages 736–740), a portion of Chapter 23 (page 830), and a portion of Chapter 24 (pages 831–845). I obtained this copy of the book via counsel from the University of Texas at Austin Library. Specifically, the text of the book 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 1010 was found within the custody of a library – a place where, if authentic, a copy of this volume would likely be. Exhibit 1010 is a true and correct copy in a condition that creates no suspicion about its authenticity.

63. Attached hereto as Attachment 3a is a true and correct copy of the MARC record for this book from the University of Texas at Austin Library. The library ownership is indicated by the presence of the library's code (IXA) in the 049 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 3a.

64. Based on finding a copy of Exhibit 1010 in the University of Texas at Austin Library and a record in its online library catalog attached as Attachment 3a, it is my opinion that the book, *SQL: The Complete Reference*, 1st edition, by Groff

and Weinberg was cataloged on May 3, 1999, as shown in Field 008 (“990503”). The most recent modification of the MARC record occurred on June 26, 2002, as indicated in field 005 (“20020626”). Exhibit 1010 includes a page of the date due slip that provides information related to the Groff book which was checked out at least 6 times. The earliest date due shown for the Groff book was August 28, 2002, which indicates it was publicly available, found, and checked out before that date.

65. Attached hereto as Attachment 3b is a true and correct copy of the MARC record for the book, *SQL: The Complete Reference*, 1st edition, by Groff and Weinberg obtained from the OCLC bibliographic database. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Exhibit 1010, that library code is “DLC,” which means that the MARC record for this book was cataloged at the Library of Congress. As can be seen in the “Entered” field in MARC record for this exhibit, a cataloger created OCLC record number 41291586 on May 3, 1999. Attachment 3b further includes an entry in field 050 (“QA76.73.S67 ‡b G758 1999”)[16]—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification

---

[16] Mathematics—Instruments and machines—Calculating machines—Electronic computers. Computer science—Digital computers—Programming languages—Individual languages, A-Z—SQL. MySQL

system). Attachment 3b further includes an entry in field 082 (“005.75/65”),[17] a subject matter consistent with the Dewey Decimal classification system. Attachment 3b further includes two English language field 650 entries reading “SQL (Computer program language)” (see Attachment 2c, Library of Congress subject heading sh86006628) and “Relational databases” (see Attachment 3c, Library of Congress subject heading sh86007768). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 3b 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 May 3, 1999, the MARC record attached hereto as Attachment 3b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the Groff and Weinberg book, which indicates that the corresponding publication was publicly available on or shortly after that same date through any library with access to the OCLC bibliographic database or through an individual library.

66. The U. S. Copyright Office Public Database record indicates that *SQL: The Complete Reference*, 1st edition, was published March 9, 1999, and registered as document TX0004891473 on May 5, 1999 (see Attachment 3d).

---

[17] 005.756 Relational databases

67. Attachment 3b indicates that the book, *SQL: The Complete Reference*, 1st edition, as cataloged at the Library of Congress is currently available from 193 libraries, as indicated by the number of other holdings stated at the top of the MARC record. The Groff and Weinberg book has been available at academic and public libraries across the United States since its publication. The Groff and Weinberg book was cataloged and available at the following representative institutions: Columbia University Libraries, the University of Illinois at Urbana-Champaign Library, and the University of California – Berkeley Library.

68. Attached hereto as Attachment 3e is a true and correct copy of the MARC record for this monograph, which is Publication 3, at the Columbia University Libraries. The library ownership is indicated by the presence of the library's local code (CU) in the 876 field. I personally identified and retrieved the library catalog record which is Attachment 3e. The MARC record is based on OCLC record number 41291586 as shown in the first field 035. Attachment 3e indicates that the Groff book was cataloged at the Columbia University Libraries as of May 3, 1999, as shown in field 008 ("990503"), which indicates that it was publicly available on or shortly after that date.

69. Attached hereto as Attachment 3f is a true and correct copy of the MARC record for this monograph, which is Publication 3, in the University of Illinois at Urbana-Champaign Library. The library ownership is indicated by the

presence of the library's code (UIU) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 3f. The MARC record is based on OCLC record number 41291586 as shown in the first field 035. Attachment 3f indicates that the Groff book was cataloged in the University of Illinois at Urbana-Champaign Library as of May 3, 1999, as shown in field 008 ("990503"), which indicates that it was publicly available on or shortly after that date. The most recent modification of the MARC record occurred on April 15, 2002, as indicated in field 005 ("20020415").

70. Attached hereto as Attachment 3g is a true and correct copy of the MARC record for this monograph, which is Publication 3, in the University of California Berkeley Library. The library ownership is indicated by the presence of the library's code (CUY) in the 994 field. I personally identified and retrieved the library catalog record which is Attachment 3g. The MARC record is based on OCLC record number 41291586 as shown in field 001. Attachment 3g indicates that the Groff book was cataloged in the University of California Berkeley Library as of November 24, 1999, as shown in field 008 ("990024"), which indicates that it was publicly available on or shortly after that date.

71. Attached hereto as Attachment 3h is a true and correct copy of the MARC record for this monograph, which is Publication 3, at the Library of Congress. I personally identified and retrieved the library catalog record which is

Attachment 3h. The 955 field indicates that Groff was sent to cataloging on July 1, 1999 (indicated by “07-01-99 to cat.”), to the shelflist on July 9, 1999 (indicated by “to sl 07-09-99”), and to Dewey Decimal Classification on July 15, 1999 (indicated by “07-15-99 to ddc”). The data in the 955 fields for the Library of Congress MARC record indicates that the processing of Groff was completed on July 15, 1999, and publicly available shortly after that date. Any person can enter the Library of Congress and request a book be brought to a reading room.

72. Not only was Exhibit 1010 accessible and available to others in the field as early as May 3, 1999, researchers actually obtained and cited this book in other works, which is still further confirmation of its availability and accessibility. For example, Google Scholar indicates that this book has been cited 206 times. Attachment 3i reports the number of citing references shown in Google Scholar for Exhibit 1010.

73. Therefore, Groff (Exhibit 1010) was publicly accessible as early as May 3, 1999, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been published, cataloged and indexed at the Library of Congress, and made part of the OCLC bibliographic database. Additionally, Groff had been checked out from the University of Texas at Austin Library before the priority date, and the MARC records for both the University of Texas at Austin Library and the University of Illinois at Urbana-Champaign

Library were last update before the priority date, which further indicates that Groff was publicly available at those libraries before the priority date.

IX. PUBLICATION 4 -- EXHIBIT 1008 ("INFORMIX")

74. Exhibit 1008 is a copy of the book, *Informix Guide to SQL: Tutorial*, 2nd edition published by Informix Press/Prentice Hall PTR in 2000 (hereinafter called "Informix"). The exhibit filed in this proceeding as Exhibit 1008 is a true and correct copy of the title page, verso of the title page, table of contents, a portion of Chapter 1 (pages 1-14 through 1-17), a portion of Chapter 2 (pages 2-56 through 2-68), a portion of Chapter 4 (pages 4-21 through 4-36), and a portion of Chapter 5 (pages 5-3 through 5-4 and 5-11 through 5-20). I obtained this copy of the book via counsel from the University of Texas at Austin Library. Specifically, the text of the book 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 1008 was found within the custody of a library – a place where, if authentic, a copy of this volume would likely be. Exhibit 1008 is a true and correct copy in a condition that creates no suspicion about its authenticity.

75. Attached hereto as Attachment 4a is a true and correct copy of the MARC record for this book from the University of Texas at Austin Library. The library ownership is indicated by the presence of the library's code (IXA) in the

049 field. I personally identified and retrieved the library catalog record which is Attachment 4a.

76. Based on finding a copy of Exhibit 1008 in the University of Texas at Austin Library and a record in its online library catalog attached as Attachment 4a, it is my opinion that the book, *Informix Guide to SQL: Tutorial*, 2nd edition, was available as early as January 24, 2000, as shown in Field 008 ("000124"). The most recent modification of the MARC record occurred on March 2, 2000, as indicated in field 005 ("20000302"). Exhibit 1008 includes a page of the date due slip that provides information related to the Informix book which was checked out at least 6 times. The earliest date due shown for the Informix book was April 10, 2000, which indicates it was publicly available, found, and checked out before that date.

77. Attached hereto as Attachment 4b is a true and correct copy of the MARC record for the book, *Informix Guide to SQL: Tutorial*, 2nd edition, obtained from the OCLC bibliographic database. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Exhibit 1008, that library code is "DLC," which means that the MARC record for this book was cataloged at the Library of Congress. As can be seen in the "Entered" field in MARC record for this exhibit, a cataloger created OCLC record number 43333718 on August 2, 2000. Attachment 4b further includes an entry in field 050

(“QA76.73.S67 ‡b I54 1999”)[18]—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal System). Attachment 4b further includes an English language field 650 entry reading “SQL (Computer program language)” (see Attachment 2c, Library of Congress subject heading sh86006628). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 4b was indexed according to its subject matter by virtue of at least two independently sufficient classifications: the field 050 entry and the field 650 entry. Further, as of August 8, 2000, the MARC record attached hereto as Attachment 4b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the book, which indicates that the corresponding publication was publicly available on or shortly after that same date through any library with access to the OCLC bibliographic database or through an individual library.

78. The U. S. Copyright Office Public Database record indicates that *Informix Guide to SQL: Tutorial*, 2nd edition, was published November 18, 1999, and registered as document TX0005125481 on January 19, 2000 (see Attachment 4c).

---

[18] Mathematics—Instruments and machines—Calculating machines—Electronic computers. Computer science—Digital computers—Programming languages—Individual languages, A-Z—SQL. MySQL

79. Attachment 4b indicates that the book, *Informix Guide to SQL: Tutorial*, 2nd edition, as cataloged at the Library of Congress is currently available from 59 libraries, as indicated by the number of other holdings stated at the top of the MARC record. The *Informix Guide to SQL: Tutorial*, 2nd edition, book has been available at academic and public libraries across the United States since its publication. The book was cataloged and available at the following representative institutions: Texas Tech University Libraries, the University of California – Los Angeles Library, and, the Simon Fraser University.

80. Attached hereto as Attachment 4d is a true and correct copy of the MARC record for this monograph, which is Publication 4, at the Texas Tech University Libraries. The library ownership is indicated by the presence of the library's code (ILU) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 4d. The MARC record is based on OCLC record number 43333718 as shown in field 035. Attachment 4d indicates that the Informix book was cataloged at the Texas Tech University Libraries as of January 24, 2000, as shown in field 008 ("000124"), which indicates that it was publicly available on or shortly after that date.

81. Attached hereto as Attachment 4e is a true and correct copy of the MARC record for this monograph, which is Publication 4, in the University of California at Los Angeles Library. The library ownership is indicated by the

presence of the library's code (CLU) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 4e. The MARC record is based on OCLC record number 43333718 as shown in the first field 035. Attachment 4e indicates that the Informix book was cataloged in the University of California at Los Angeles Library as of August 2, 2000, as shown in field 008 ("000802"), which indicates that it was publicly available on or shortly after that date.

82. Attached hereto as Attachment 4f is a true and correct copy of the MARC record for this monograph, which is Publication 4, in the Simon Fraser University Library. The library ownership is indicated by the presence of the library's local code (SFUL) in the INST field. I personally identified and retrieved the library catalog record which is Attachment 4f. The MARC record is based on OCLC record number 43333718 as shown in the second field 035. Attachment 4f indicates that the Informix book was cataloged in the Simon Fraser University Library as of January 24, 2000, as shown in field 008 ("000124"), which indicates that it was publicly available on or shortly after that date.

83. Attached hereto as Attachment 4g is a true and correct copy of the MARC record for this monograph, which is Publication 4, at the Library of Congress. I personally identified and retrieved the library catalog record which is Attachment 4g. The 955 field indicates that Informix was received on August 2,

2000 (indicated by "to ASCD pb11 08-02-00"), sent to the shelflist on March 27, 2001 (indicated by "to SL 03-27-01"), and that two additional copies were sent to Dewey on May 3, 2001 (as indicated by "2 copies to Dewey 05-03-01"). The date Informix was sent to ASCD matches the date the book was cataloged, as shown in field 008 ("000802"). The most recent modification of the MARC record occurred on November 16, 2001, as indicated in field 005 ("20011116"). The data in the 955 fields for the Library of Congress MARC record indicates that Informix was cataloged on August 2, 2000 and processing of the first copy was completed on March 27, 2001, and it was publicly available shortly after that date. Any person can enter the Library of Congress and request a book be brought to a reading room.

84. Therefore, *Informix Guide to SQL: Tutorial*, 2nd edition (Exhibit 1008) was publicly accessible as early as January 24, 2000, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received at and checked out from the University of Texas at Austin Library.

X. PUBLICATION 5 -- EXHIBIT 1014 ("SELVARAJ")

85. Exhibit 1014 is a copy of Volume 32, Issue 6 of the journal *ACM SIGPLAN Notices* found in the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder. The article "Implementation

of a Database Factory" by Asokan R. Selvaraj and Debasish Ghosh (hereafter "Selvaraj") appears beginning on page 14 of this issue dated June 1997. Exhibit 1014 is a true and correct copy of the Selvaraj article that I understand is being submitted as an exhibit in this proceeding. Exhibit 1014 is a true and correct copy of the issue cover, table of contents, and the article (pages 14-18). I obtained this copy of the article from the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder. Specifically, the text of the article 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 1014 was found within the custody of a library – a place where, if authentic, a copy of this journal would likely be. Exhibit 1014 is a true and correct copy in a condition that creates no suspicion about authenticity.

86. The cover of the June 1997 issue of the journal *ACM SIGPLAN Notices* has ownership markings affixed by or on behalf of the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder. The *ACM SIGPLAN Notices* is published by the Association of Computing Machinery (ACM) and is indexed in the *ACM Digital Library* (see paragraph 24c). Therefore, in my experience, this issue of the journal *ACM SIGPLAN Notices* would have been available to users at the Gemmill Library on July 14, 1997, or shortly after that date.

87. Attached hereto as Attachment 5a is a true and correct copy of the MARC record for the journal *ACM SIGPLAN Notices* in the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder. The library ownership is indicated by the presence of the library's code (DVD) in the 049 field. I personally identified and retrieved the MARC record that is Attachment 5a. Attachment 5a also shows that Exhibit 1014 was catalogued with descriptor terms "Programming languages (Electronic computers) ‡v Periodicals" in the 650 field.

88. Based on finding a print copy of Exhibit 1014 in the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder and MARC record in its online library catalog attached as Attachment 5a, it is my opinion that the article "Implementation of a Database Factory" by Selvaraj and Ghosh published in the journal *ACM SIGPLAN Notices* was available in the Gemmill Library on or shortly after July 14, 1997.

89. As noted in the holdings information (field 362), the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder has received the journal *ACM SIGPLAN Notices* since 1967 and continues to receive the publication.[19] In view of the MARC record for Exhibit 1014, the

---

[19] The publication was initially known as *SIGPLAN Notices* (1966-September 1991). The title was changed to *ACM SIGPLAN Notices* in October 1991 and

Selvaraj article was publicly available on or shortly after July 14, 1997, because the serial title had been cataloged and indexed at the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder and made part of its online catalog database.

90. Attached hereto as Attachment 5b is a true and correct copy of the MARC record for the journal *ACM SIGPLAN Notices* obtained from the OCLC bibliographic database. I personally identified and retrieved the MARC record that is Attachment 5b. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Attachment 5b, that library code is "TJC," which means that the MARC record for this serial was cataloged at the Vanderbilt University Library. As can be seen in the "Entered" field in the MARC record for this exhibit, a cataloger at the Vanderbilt University Library created OCLC record number 25073822 on January 9, 1992. The "BLvl" entry in Attachment 5b is "s," which indicates that the journal *ACM SIGPLAN Notices* is a serial publication. Field 310 of Attachment 5b reads "Monthly." Accordingly, the MARC record for Exhibit 1014 corresponds to those issues of the journal *ACM SIGPLAN Notices* from the time of the serial title began publication.

---

retained that title through February 2007. The publication adopted its current title, *SIGPLAN Notices: A Monthly Publication of the ACM Special Interest Group on Programming Languages*, in March 2007 and is published in print and digital versions.

91. Attachment 5b further includes an entry in field 050 ("QA76.7 $b .S54")—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system). Attachment 5b further includes an entry in field 082 ("001.64/24/05"), a subject matter consistent with the Dewey Decimal classification system. Attachment 5b further includes an English language field 650 entry reading "Programming languages (Electronic computers) ‡v Periodicals" (see Attachment 5c, Library of Congress subject heading sh2010108608). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 5b 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 entry. Further, as of January 9, 1992, the MARC record attached hereto as Attachment 5b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that subscribed to the serial, which indicates that the corresponding publication was publicly available on or shortly after that same date through any library with access to the OCLC bibliographic database or through an individual library.

92. Attachment 5b indicates that the journal *ACM SIGPLAN Notices* as cataloged at the Vanderbilt University Library is currently available from 348 libraries, as indicated by the number of other holdings stated at the top of the

MARC record. In view of the above, this issue of the journal *ACM SIGPLAN Notices* was publicly available on or shortly after July 14, 1997, because by that date it had been cataloged and indexed at the Vanderbilt University Library, made part of the OCLC bibliographic database, and received at the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder. For these reasons, it is my opinion that Exhibit 1014 was published and accessible to the public on or shortly after July 14, 1997.

93. Not only was Exhibit 1014 accessible and available to others in the field as early as July 1997, researchers actually obtained and cited this article in other works, which is still further confirmation of its availability and accessibility. For example, Google Scholar indicates that this article has been cited 7 times (see Attachment 5d). Scopus reports that this article has been cited 3 times (see Attachment 5e).

94. Therefore, Selvaraj and Ghosh (Exhibit 1014) was publicly accessible as early as July 14, 1997, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Vanderbilt University Library, made part of the OCLC bibliographic database, and received at the Gemmill Library for Engineering, Mathematics, and Physics at the University of Colorado – Boulder.

XI. PUBLICATION 6 -- EXHIBIT 1009 ("TRANSACT-SQL")

95. Exhibit 1009 is a copy of the book, *Transact-SQL Programming*, 1st edition, by Kevin Kline, Lee Gould, and Andrew Zanevsky and published by O'Reilly & Associates in 1999 (hereinafter called "Transact-SQL"). The exhibit filed in this proceeding as Exhibit 1009 is a true and correct copy of the title page, verso of the title page, table of contents, foreword, preface, a portion of Chapter 1 (pages 3-16), a portion of Chapter 3 (pages 76-100), and a portion of Chapter 12 (pages 366-380). I obtained this copy of the book via counsel from the University of Texas at Austin Library. Specifically, the text of the book 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 1009 was found within the custody of a library – a place where, if authentic, a copy of this volume would likely be. Exhibit 1009 is a true and correct copy in a condition that creates no suspicion about its authenticity.

96. Attached hereto as Attachment 6a is a true and correct copy of the MARC record for this book from the University of Texas at Austin Library. The library ownership is indicated by the presence of the library's code (IXA) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 6a.

97. Based on finding a copy of Exhibit 1009 in the University of Texas at

Austin Library and a record in its online library catalog attached as Attachment 6a, it is my opinion that the book, *Transact-SQL Programming*, 1st edition, was cataloged on April 19, 1999, as shown in Field 008 ("990419"). The last modification to the MARC record occurred on December 2, 1999, as shown in field 005 ("19991202"). Exhibit 1009 includes two pages of the date due slips that provide information related to the Transact-SQL book which was checked out at least 12 times. The earliest date due shown for the Transact-SQL book was January 18, 2000, which indicates it was publicly available, found, and checked out before that date.

98. Attached hereto as Attachment 6b is a true and correct copy of the MARC record for the book, *Transact-SQL Programming*, 1st edition, obtained from the OCLC bibliographic database. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Exhibit 1009, that library code is "DLC," which means that the MARC record for this book was cataloged at the Library of Congress. As can be seen in the "Entered" field in MARC record for this exhibit, a cataloger created OCLC record number 41388930 on May 26, 2000. Attachment 6b further includes an entry in field 050

("QA76.73.S67 ‡b K555 1999")[20]—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal System). Attachment 6b further includes an entry in field 082 ("005.75/6")[21]—as described above, this includes a subject matter classification number consistent with the Dewey Decimal classification system. Attachment 6b further includes an English language field 650 entry reading "SQL (Computer program language)" (see Attachment 2c, Library of Congress subject heading sh86006628). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 6b 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 entry. Further, as of May 26, 2000, the MARC record attached hereto as Attachment 6b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the book, which means that the corresponding publication was publicly available on or shortly after that same date

---

[20] Mathematics—Instruments and machines—Calculating machines—Electronic computers. Computer science—Digital computers—Programming languages—Individual languages, A-Z—SQL. MySQL

[21] Computer science, information & general works – Specific types of data files and databases – Databases based on specific data models – Relational databases

through any library with access to the OCLC bibliographic database or through an individual library.

99. The U. S. Copyright Office Public Database record indicates that *Transact-SQL Programming*, 1st edition, was published March 31, 1999, and registered as document TX0005106977 on December 9, 1999 (see Attachment 6c).

100. Attachment 6b indicates that the book, *Transact-SQL Programming*, 1st edition, as cataloged at the Library of Congress is currently available from 135 libraries, as indicated by the number of other holdings stated at the top of the MARC record. The *Transact-SQL Programming*, 1st edition, book has been available at academic and public libraries across the United States since its publication. The book was cataloged and available at the following representative institutions: the University of Denver, Carnegie Mellon University, and the University of California – San Diego.

101. Attached hereto as Attachment 6d is a true and correct copy of the MARC record for this monograph, which is Publication 6, at the University of Denver Libraries. The library ownership is indicated by the presence of the library's code (DVP) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 6d. The MARC record is based on OCLC record number 41388930 as shown in the second field 035. Attachment 6d indicates that the Transact-SQL book was cataloged at the University of Denver

Libraries as of May 14, 1999, as shown in field 008 ("990514"), which indicates that it was publicly available on or shortly after that date.

102. Attached hereto as Attachment 6e is a true and correct copy of the MARC record for this monograph, which is Publication 6, in the Hunt Library at Carnegie Mellon University. The library ownership is indicated by the presence of the library's code (PMC) in the 948 field. I personally identified and retrieved the library catalog record which is Attachment 6e. The MARC record is based on OCLC record number 41388930 as shown in the field 903. Attachment 6e indicates that the Transact-SQL book was cataloged in the Hunt Library at Carnegie Mellon University as of May 26, 2000, as shown in field 008 ("000526"), which indicates that it was publicly available on or shortly after that date. There is no indication that the MARC record has been modified since the original cataloging.

103. Attached hereto as Attachment 6f is a true and correct copy of the MARC record for this monograph, which is Publication 6, in the University of California – San Diego Library. The library ownership is indicated by the presence of the library's local code (CUS) in the 994 field. I personally identified and retrieved the library catalog record which is Attachment 6f. The MARC record is based on OCLC record number 41388930 as shown in the 001 field. Attachment 6f indicates that the Informix book was cataloged in the University of

California – San Diego Library as of May 26, 2000, as shown in field 008 ("000526"), which indicates that it was publicly available on or shortly after that date.

104. Attached hereto as Attachment 6g is a true and correct copy of the MARC record for this monograph, which is Publication 6, at the Library of Congress. I personally identified and retrieved the library catalog record which is Attachment 6g. The 955 field indicates that Transact-SQL was received) on May 26, 2000 (indicated by "to ASCD pb08 05-26-00"), sent to the shelflist on October 2, 2000 (indicated by "to SL 10-2-00"), and sent to Dewey on October 5, 2000 (as indicated by "to Dewey 10-05-00"). Also, a second copy was added and sent to BCCD (Binding and Collection Care Division) on December 4, 2000 (indicated by "copy 2 added jf16 to BCCD 12-04-00"). The date Transact-SQL was sent to ASCD matches the date the book was cataloged, as shown in field 008 ("000526"). The data in the 955 fields for the Library of Congress MARC record indicates that Transact-SQL was cataloged on May 26, 2000, and that final processing occurred on October 5, 2000, and was publicly available on or shortly after that date. Any person can enter the Library of Congress and request a book be brought to a reading room.

105. Not only was Exhibit 1009 accessible and available to others in the field as early as May 26, 2000, researchers actually obtained and cited this book in

other works, which is still further confirmation of its availability and accessibility. For example, Google Scholar indicates that this book has been cited 18 times. Attachment 6h reports the number of citing references shown in Google Scholar for Exhibit 1009.

106. Therefore, *Transact-SQL Programming*, 1st edition, (Exhibit 1009) was publicly accessible as early as April 19, 1999, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received at and checked out from the University of Texas at Austin Library.

XII. PUBLICATION 7 -- EXHIBIT 1011 ("KERNIGHAN")

107. Exhibit 1011 is a copy of the book, *The C Programming Language*, 2nd edition, by Brian W. Kernighan and Dennis M. Ritchie and published by Prentice Hall in 1988 (hereinafter called "Kernighan"). The exhibit filed in this proceeding as Exhibit 1011 is a true and correct copy of the title page, verso of the title page, table of contents, a portion of Chapter 4 (pages 91-92), and a portion of Appendix A (pages 231-232). I obtained this copy of the book from the Dayton Memorial Library at Regis University and made the copies which comprise Exhibit 1011. Specifically, the text of the book 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 1011 was found within the custody of a library – a place where, if authentic, a copy of this volume would likely be. Exhibit 1011 is a true and correct copy in a condition that creates no suspicion about its authenticity.

108. Attached hereto as Attachment 7a is a true and correct copy of the MARC record for this book from the Dayton Memorial Library at Regis University. The library ownership is indicated by the presence of the library's code (COR) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 7a.

109. Based on finding a copy of Exhibit 1011 in the Dayton Memorial Library at Regis University and a record in its online library catalog attached as Attachment 7a, it is my opinion that the book, *The C Programming Language*, 2nd edition, by Kernighan and Ritchie was cataloged on February 16, 1989, as shown in field 998 ("02-16-89"). The last modification of the MARC record occurred on February 16, 1989, as indicated in field 005 ("19890216"). A library employee placed an ownership stamp with the receipt date of April 1, 1998, on the verso of the title page for the copy of the book that is Exhibit 1011

110. Attached hereto as Attachment 7b is a true and correct copy of the MARC record for the book, *The C Programming Language*, 2nd edition, by Kernighan and Ritchie obtained from the OCLC bibliographic database. As

previously noted, the library that created the record is recorded in field 040 with a unique library code. For Exhibit 1011, that library code is "DLC," which means that the MARC record for this book was cataloged at the Library of Congress. As can be seen in the "Entered" field in MARC record for this exhibit, a cataloger created OCLC record number 17650642 on March 3, 1988. Attachment 7b further includes an entry in field 050 ("QA76.73.C15 ‡b K47 1988")[22]—as described above, this includes a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system). Attachment 7b further includes an entry in field 082 ("005.13/3")[23]—as described above, this includes a subject matter classification number consistent with the Dewey Decimal classification system). Attachment 7b further includes an English language field 650 entry reading "C (Computer program language)" (see Attachment 7c, Library of Congress subject heading sh85018532). Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 7b was indexed according to its subject matter by virtue of at least three independently sufficient classifications:

---

[22] Mathematics—Instruments and machines—Calculating machines—Electronic computers. Computer science—Digital computers—Programming languages—Individual languages, A-Z—C

[23] Computer science, information & general works – Computer programming, programs, data, security – Programming languages – Specific programming languages

the field 050 entry, the field 082 entry, and the field 650 entry. Further, as of March 3, 1988, the MARC record attached hereto as Attachment 7b was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that acquired the book, which indicates that the corresponding publication was publicly available on or shortly after that same date through any library with access to the OCLC bibliographic database or through an individual library.

111. The U. S. Copyright Office Public Database record indicates that *The C Programming Language*, 2nd edition, by Kernighan and Ritchie was published March 10, 1988, and registered as document TX0002341624 on June 17, 1988 (see Attachment 7d).

112. Attachment 7b indicates that the book, *The C Programming Language*, 2nd edition, by Kernighan and Ritchie as cataloged at the Library of Congress is currently available from 1,375 libraries, as indicated by the number of other holdings stated at the top of the MARC record. *The C Programming Language*, 2nd edition, by Kernighan and Ritchie has been available at academic and public libraries across the United States since its publication. The book was cataloged and available at the following representative institutions: the University of Texas at Austin Library, the University of Colorado – Boulder Libraries, and, San José State University Library.

113. Attached hereto as Attachment 7e is a true and correct copy of the MARC record for this monograph, which is Publication 7, at the University of Texas at Austin Library. The library ownership is indicated by the presence of the library's code (IXA) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 7e. The MARC record is based on OCLC record number 17650642 as shown in the first field 001. Attachment 7e indicates that the Kernighan book was cataloged for the University of Texas at Austin Library as of March 3, 1988, as shown in field 008 ("880303"), which indicates that it was publicly available on or shortly after that date. The most recent modification of the MARC record occurred on November 8, 1988, as indicated in field 005 ("19881108").

114. Attached hereto as Attachment 7f is a true and correct copy of the MARC record for this monograph, which is Publication 7, in the University of Colorado – Boulder Libraries. The library ownership is indicated by the presence of the library's code (COD) in the 049 field. I personally identified and retrieved the library catalog record which is Attachment 7f. The MARC record is based on OCLC record number 17650642 as shown in the field 001. Attachment 7f indicates that the Kernighan book was cataloged in the University of Colorado – Boulder Libraries as of March 3, 1988, as shown in field 008 ("880303"), which indicates that it was publicly available on or shortly after that date. There is no

indication that the MARC record has been modified.

115. Attached hereto as Attachment 7g is a true and correct copy of the MARC record for this monograph, which is Publication 7, in the King Library at San José State University. The library ownership is indicated by the presence of the library's local code (SJO) in the corresponding INST field. I personally identified and retrieved the library catalog record which is Attachment 7g. The MARC record is based on OCLC record number 17650642 as shown in the first field 035. Attachment 7g indicates that the Kernighan book was cataloged in the King Library at San José State University as of March 3, 1988, as shown in field 008 ("880303"), which indicates that it was publicly available on or shortly after that date.

116. Therefore, *The C Programming Language*, 2nd edition (Exhibit 1011) was publicly accessible as early as March 3, 1988, and in any event more than a year before the November 22, 2002, priority date, because by that time it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received at the Dayton Memorial Library at Regis University and University of Texas at Austin Library, which have not updated their MARC records since February 16, 1989, and November 8, 1988, respectively.

XIII. SUMMARY OF OPINIONS

117. In view of the foregoing, it is my opinion that the publications

described above were publicly available on or shortly after the corresponding date listed in the table below, and in any event more than a year before the November 22, 2002, priority date.

Exhibit	Publication	Publicly Available As Of
1006	Templeton, Marjorie, Herbert Henley, Edward Maros, and Darrel J. Van Buer. “InterViso: Dealing with the Complexity of Federated Database Access.” *The VLDB Journal*, vol. 4, issue 2 (April 1995): 287–317.	May 30, 1995
1007	Kline, Kevin, and Daniel T. Kline. *SQL In a Nutshell: A Desktop Quick Reference.* 1st ed. Cambridge, MA: O’Reilly & Associates, 2001.	November 6, 2000
1010	Groff, James R., and Paul N. Weinberg. *SQL: The Complete Reference.* 1st ed. Berkeley, CA: Osborne/McGraw-Hill, 1999.	May 3, 1999
1008	*Informix Guide to SQL: Tutorial*. 2nd ed. Menlo Park, CA: Informix Press ; Upper Saddle River, NJ: Prentice Hall PTR, ©2000.	January 24, 2000
1014	Selvaraj, Asokan R., and Debasish Ghosh. “Implementation of a Database Factory.” *ACM SIGPLAN Notices*, vol. 32, issue 6 (June 1997): 14-18.	July 14, 1997
1009	Kline, Kevin, Lee Gould, and Andrew Zanevsky. *Transact-SQL Programming*. 1st ed. Cambridge, MA: O'Reilly & Associates, 1999.	April 19, 1999
1011	Kernighan, Brian W., and Dennis M. Ritchie. *The C Programming Language*. 2nd ed. New York: Prentice Hall, 1988.	March 3, 1988

118. In signing this Declaration, I recognize that the Declaration will be filed as evidence in a case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the

United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

119. 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: August 2, 2019

By: Sylvia D. Hall-Ellis

Dr. Sylvia D. Hall-Ellis

Attachment 1a

(Search History) ▼

```
LEADER 00000cas a22004937a 4500
001    27078023
003    OCoLC
005    19940715101612.0
008    921207c19929999cauqr p       0   a0eng d
010    sn 92006846
022 0  1066-8888
022    |y0949-877x
040    NSD|cNSD|dCOD
042    nsdp|alcd
049    DVDI
090    QA76.9.D3|bV25
090    QA76.9.D3|bV25
210 0  VLDB j
222  0 Very large data bases journal
222  4 The VLDB journal
245 04 The VLDB journal
246 2  Very large data bases journal
246 13 Very large data bases journal
246 30 Very large data journal
260    Pacific Grove, CA :|bBoxwood Press
300    volumes :|billustrations ;|c25cm
310    Quarterly
336    text|btxt|2rdacontent
337    unmediated|bn|2rdamedia
338    volume|bnc|2rdacarrier
362 0  Vol. 1, no. 1 (July 1992)-
500    Title from cover
500    "The international journal on very large data bases."
550    "A publication of the VLDB Endowment."
583 1  |31-13 (1992-2004) : |acommitted to retain|c20180101
       |d20430101|5CoU|fAlliance Shared Trust|uhttps://
       www.coalliance.org/shared-print-archiving-policies
650  0 Database management|vPeriodicals
650  0 Databases|vPeriodicals
650  0 Data structures (Computer science)|vPeriodicals
710 2  VLDB Endowment
907    .b19104236|b06-10-10|c11-30-94
907    .b19104236|b04-25-08|c11-30-94
944    RDA ENRICHED
944    MARS
947    04/05aw cg|bno
947    cgs 3/98 af www access
```

This title is available online:

Full-text in ACM digital library	**July 01, 1992-**
Full-text in CASC	**Feb. 01, 2010 - Feb. 14, 2018**

Location	**Internet Access Online**
Holdings	**Springer has: 1996-, ACM has: 1992-**

Location	**PASCAL Offsite QA76.9.D3 V25**
Holdings	**v.1 (1992) - v.13 (2004)**
	Request a PDF (UCB only)

LOCATION	CALL #	STATUS
PASCAL Offsite	QA76.9.D3 V25 v.1 1992	AVAILABLE

PASCAL Offsite	QA76.9.D3 V25 v.2 1993	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.3 1994	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.4 1995	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.5-6 1996-97	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.7-8 1998/2000	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.9-10 2000/01	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.11 2002	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.12 2003	AVAILABLE
PASCAL Offsite	QA76.9.D3 V25 v.13 2004	AVAILABLE

Attachment 1b

OCLC 27078023 **No holdings in XXX - 126 other holdings**

Continuing Resources						**Rec Stat**	c	**Entered**	19921207	**Replaced**	20181120205007.7
Type	a	**ELvl**		**Srce**	c	**GPub**		**Ctrl**		**Lang**	eng
BLvl	s	**Form**		**Conf**	0	**Freq**	q	**MRec**		**Ctry**	gw
S/L	0	**Orig**		**EntW**		**Regl**	r	**Alph**	a		
Desc	a	**SrTp**	p	**Cont**		**DtSt**	c	**Dates**	1992 , 9999		

010 sn 92006846
040 NSD $b eng $c NSD $d CAS $d CUS $d OCL $d MUQ $d ORZ $d ORU $d MYG $d DLC $d OCLCQ $d UKMGB $d OCLCQ $d OCLCF $d IUL
012 $l a
016 7 011772855 $2 Uk
019 1063579263
022 0 1066-8888 $l 1066-8888 $2 6
030 VLDBFR
042 nsdp $a pcc
050 14 QA76.9.D3 $b V53
082 10 004 $2 12
090 $b
210 0 VLDB j.
222 4 The VLDB journal
245 04 The VLDB journal : $b very large data bases : a publication of the VLDB Endowment.
246 3 Very large data bases journal
246 30 Very large data bases
260 Pacific Grove, CA : $b Boxwood Press, $c c1992-
260 3 $3 <2004-> : $a Heidelberg, Germany : $b Springer ; $a New York, NY : $b Association for Computing Machinery
300 v. : $b ill. ; $c 28 cm.
310 Quarterly
362 0 Vol. 1, no. 1 (July 1992)-
500 Subtitle varies.
588 Description based on surrogate of Vol. 1, no. 1 (July 1992); title from cover.
588 Latest issue consulted: Vol. 13, no. 3 (Sept. 2004).
650 0 Database management $v Periodicals.
650 0 Data structures (Computer science) $v Periodicals.
650 0 Databases $v Periodicals.
650 6 Bases de données $v Périodiques.
650 7 Data structures (Computer science) $2 fast $0 (OCoLC)fst00887978
650 7 Database management. $2 fast $0 (OCoLC)fst00888037
650 7 Databases. $2 fast $0 (OCoLC)fst00888065
655 7 Periodicals. $2 fast $0 (OCoLC)fst01411641
710 2 VLDB Endowment.
776 08 $i Online version: $t VLDB journal (Online) $x 0949-877X $w (DLC) 2006242190 $w (OCoLC)43476370

776 08 $i Online version: $t VLDB journal $w (OCoLC)607455343
856 41 $u http://www.springerlink.com/openurl.asp?genre=journal&issn=1066-8888
856 41 $u
http://portal.acm.org/affiliated/browse_dl.cfm?linked=1&part=affil&idx=J869&coll=portal&dl=ACM
891 20 $9 853 $8 1 $a v. $b no. $u 4 $v r $i (year) $j (month) $w q
891 41 $9 863 $8 1.1 $a 1 $b 1 $i 1992 $j 07

Attachment 1c

LC control no.: sh2008102046
LCCN Permalink: https://lccn.loc.gov/sh2008102046
HEADING: Database management Periodicals
000 00478cz a2200169n 450
001 7439992
005 20110729160139.0
008 080207|| anannbabn |n ana
010 __ **|a** sh2008102046
035 __ **|a** (DLC)377612
035 __ **|a** (DLC)sh2008102046
040 __ **|a** DLC **|b** eng **|c** DLC
150 __ **|a** Database management **|v** Periodicals
667 __ **|a** Record generated for validation purposes.
670 __ **|a** Work cat.: Access office, VB advisor magazine, -[1999]
906 __ **|t** 8888 **|u** te00 **|v** 0
953 __ **|a** te00

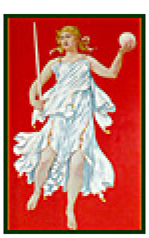

Library of Congress
URL: https://www.loc.gov/

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

Attachment 1d

LC control no.: sh 85035862
LCCN Permalink: https://lccn.loc.gov/sh85035862
HEADING: Data structures (Computer science)
000 00681cz a2200217n 450
001 4687969
005 20120323071915.0
008 860211i| anannbabn |b ana
010 __ |**a** sh 85035862
035 __ |**a** (DLC)sh 85035862
035 __ |**a** (DLC)34654
040 __ |**a** DLC |**c** DLC |**d** DLC
150 __ |**a** Data structures (Computer science)
450 __ |**a** Information structures (Computer science)
450 __ |**a** Structures, Data (Computer science)
450 __ |**a** Structures, Information (Computer science)
550 __ |**w** g |**a** Electronic data processing
550 __ |**w** g |**a** File organization (Computer science)
550 __ |**a** Abstract data types (Computer science)
906 __ |**t** 9831 |**u** te07 |**v** 1
953 __ |**a** xx00 |**b** ta25

Previous Next

Attachment 1e

LC control no.: sh 86007767

LCCN Permalink: https://lccn.loc.gov/sh86007767

HEADING: Databases

000 02796cz a2200541n 450

001 4805239

005 20120329071820.0

008 861205i| anannbabn |a ana

010 __ **|a** sh 86007767

035 __ **|a** (DLC)sh 86007767

035 __ **|a** (DLC)151924

035 __ **|a** (DLC)7680749

035 __ **|a** (DLC)sp 86007767

035 __ **|a** (DLC)410912

040 __ **|a** DLC **|c** DLC **|d** DLC **|d** ViU

053 _0 **|a** QA76.9.D32 **|c** Computer science

053 _0 **|a** ZA4450 **|b** ZA4460 **|c** Information resources

150 __ **|a** Databases

360 __ **|i** subdivision **|a** Databases **|i** under subjects for actual databases and works on databases in those subjects

450 __ **|a** Data banks

450 __ **|w** nne **|a** Data bases

450 __ **|a** Databanks

450 __ **|a** Database systems

550 __ **|w** g **|a** Computer files

550 __ **|w** g **|a** Electronic information resources

670 __ **|a** Work cat.: Lesko, M. The computer data and data base source book, 1984.

670 __ **|a** ASTI **|b** (Indexed under Data base theory)

670 __ **|a** Business Periodicals Index **|b** (Data base systems; Data base management systems; x Data bases, Computer)

670 __ **|a** Encyclopedia of Computer Science **|b** (Data base and data base mgmt. is an article caption)

670 __ **|a** Engr. index **|b** (Indexed under Data base systems)

670 __ **|a** Maynard, J. Dictionary of data processing **|b** (Data base; a file so structured that all logical relationships can be expressed between records, and that is independent of any particular application)

670 __ **|a** Encyclopedia of computer science **|b** ("The term data base has yet to achieve a widely accepted meaning")

670 __ **|a** Webster’s new world compact dictionary of computer terms **|b** (data base: a set of inter-related data records stored on a direct access storage device)

670 __ **|a** LC database, Nov. 26, 1986.

670 __ **|a** National Computing Centre, Ltd. Thesaurus of computing terms **|b** (data bases)

670 __ **|a** Computer dict. **|b** (database; a typical database is a vast and continually updated file of information, abstracts, or references on a particular subject or subjects)

670 __ **|a** Amer. Heritage dict. **|b** (database; also data base)

670 __ **|a** Random House **|b** (database; also data-base, data base; data bank; also databank)

670 __ **|a** MESH **|b** (database; databases)

670 __ **|a** Hennepin **|b** (database; databases)

670 __ **|a** LC database, Nov. 6, 1993 **|b** (data base (607); database (1029); data bases (162); databases (331))

670 __ **|a** UMI business vocab. **|b** (Data banks use Data bases)

670 __ **|a** LIV **|b** (Data banks)

670 __ **|a** Daintith, J. Dictionary of computing, 2008: **|b** (database system; UF: database)

675 __ **|a** Am. National Dict. for Information Processing Systems; **|a** Jordain. Condensed computer encyclopedia; **|a** Words bus.; **|a** Web. 3

906 __ **|t** 0846 **|u** te04 **|v** 0

952 __ **|a** LC-BRF

953 __ **|a** fg01 **|b** ym04

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

~~Help~~ - Search - ~~Search History~~ - ~~Headings List~~ - Start Over

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

Attachment 1f

LC control no.: sh 85099890
LCCN Permalink: https://lccn.loc.gov/sh85099890
HEADING: Periodicals

000 01324cz a2200337n 450
001 4749694
005 20120717060125.0
008 860211 | anannbabn |b ana
010 __ |**a** sh 85099890
035 __ |**a** (DLC)sh 85099890
035 __ |**a** (DLC)6226910
035 __ |**a** (DLC)sp 85099890
035 __ |**a** (DLC)314952
040 __ |**a** DLC |**c** DLC |**d** DLC
053 _0 |**a** AP |**c** General periodicals
053 _0 |**a** PN4700 |**b** PN5650 |**c** History
150 __ |**a** Periodicals
360 __ |**i** subdivision |**a** Periodicals |**i** under specific subjects, e.g. |**a** Engineering--Periodicals; United States--History--Periodicals
450 __ |**a** Journals (Periodicals)
450 __ |**a** Magazines
550 __ |**w** g |**a** Library materials
550 __ |**w** g |**a** Mass media
550 __ |**w** g |**a** Serial publications
550 __ |**a** Newspapers
550 __ |**a** Press
670 __ |**a** UMI business vocab. |**b** (Journals)
680 __ |**i** Here are entered works on the periodicals of the world and, with appropriate subdivisions, works on special aspects or sections of periodicals. Works on periodicals in a specific language, or in a specific country or larger area, are entered under the adjectival form of the language or area, e.g. |**a** Arabic periodicals; Canadian periodicals.
906 __ |**t** 0420 |**u** te03 |**v** 0
910 __ |**a** Proposal saved by tc03 on 05/14/2012 at 17:15:30
953 __ |**a** xx00 |**b** td14
906 __ |**t** 1206 |**v** 0

Attachment 1g

InterViso: Dealing with the complexity of federated database access

M Templeton, H Henley, E Maros… - The VLDB Journal—The …, 1995 - dl.acm.org

Connectivity products are finally available to provide the" highways" between computers containing data. IBM has provided strong validation of the concept with their" Information Warehouse." DBMS vendors are providing gateways into their products, and SQL is being retrofitted on many older DBMSs to make it easier to access data from standard 4GL products and application development systems. The next step needed for data integration is to provide (1) a common data dictionary with a conceptual schema across the data to mask …

Attachment 1h

Scopus

1 document result

Search within results...

Refine results

Limit to Exclude

Access type

- Other (1)

Year

- 1995 (1)

Author name

- Henley, H. (1)
- Maros, E. (1)
- Templeton, M. (1)
- Van Buer, D.J. (1)

Subject area

- Computer Science (1)

Document type

Publication stage

Source title

Keyword

Affiliation

Funding sponsor

Country/territory

Source type

Language

Limit to Exclude

Export refine

Analyze search results | Show all abstracts | Sort on: Date (newest)

All | Export | Download | View citation overview | View cited by | Add to List

	Document title	Authors	Year	Source	Cited by
1	InterViso: Dealing with the complexity of federated database access	Templeton, M., Henley, H., Maros, E., Van Buer, D.J.	1995	The VLDB Journal 4(2), pp. 287-317	19

View abstract | Find It | View at Publisher | Related documents

Display: 20 results per page

1

Top of page

About Scopus

What is Scopus

Language

日本語に切り替える

Customer Service

Help

Attachment 1i

Cost-based Query Scrambling for Initial Delays[*] [†]

Tolga Urhan[*]
University of Maryland
urhan@cs.umd.edu

Michael J. Franklin[*]
University of Maryland
franklin@cs.umd.edu

Laurent Amsaleg[‡]
IRISA
Laurent.Amsaleg@irisa.fr

Abstract

Remote data access from disparate sources across a wide-area network such as the Internet is problematic due to the unpredictable nature of the communications medium and the lack of knowledge about the load and potential delays at remote sites. Traditional, static, query processing approaches break down in this environment because they are unable to adapt in response to unexpected delays. Query scrambling has been proposed to address this problem. Scrambling modifies query execution plans on-the-fly when delays are encountered during runtime. In its original formulation, scrambling was based on simple heuristics, which although providing good performance in many cases, were also shown to be susceptible to problems resulting from bad scrambling decisions. In this paper we address these shortcomings by investigating ways to exploit query optimization technology to aid in making intelligent scrambling choices. We propose three different approaches to using query optimization for scrambling. These approaches vary, for example, in whether they optimize for total work or response-time, and whether they construct partial or complete alternative plans. Using a two-phase randomized query optimizer, a distributed query processing simulator, and a workload derived from queries of the TPC-D benchmark, we evaluate these different approaches and compare their ability to cope with initial delays in accessing remote sources. The results show that cost-based scrambling can effectively hide initial delays, but that in the absence of good predictions of expected delay durations, there are fundamental tradeoffs between risk aversion and effectiveness.

[*]Appears in ACM SIGMOD International Conference on the Management of Data, Seattle, WA, June 1998.

[†]This work was partially supported by the NSF under grant IRI-94-09575, by the Office of Naval Research under contract number N66001-97-C8539 (DARPA order number F475), by Bellcore, and by an IBM Shared University Research award.

[‡]Laurent Amsaleg's work was partially supported by an INRIA fellowship. He performed this work while he was with the University of Maryland.

1 Introduction

The ubiquity of wide-area connectivity has led to tremendous increases in the number, variety, and distribution of data sources that can be accessed from one's desktop. At present, most such access is performed in a browsing mode, either by navigating through hyperlinks or by performing word-based searches via search engines. Distributed query processing, as developed for relational and object relational database systems, currently plays little role in wide-area data access. As a result, the benefits of declarative query processing, such as the expressive power of query languages and automated optimization of query plans, are largely unavailable when accessing wide-spread data sources across the internet.

The absence of declarative query processing places unnecessary restrictions on the types of applications that can exploit the increased interconnectivity of data sources. The severity of such a limitation has been demonstrated before, most recently in the battle between object-oriented database (OODB) and object-relational database (ORDB) systems. To date, navigation-oriented OODB approaches have remained largely niche solutions, while the query-oriented ORDB approach has been embraced by most of the major database vendors [CD96]. Given the importance of declarative query processing for many applications, it is natural to investigate ways to provide such functionality over the wealth of data that is available across current wide-area networks.

Query processing in wide-area distributed environments poses a number of difficult technical challenges. Issues such as semantic heterogeneity, manipulation of semi-structured data, and resource discovery (i.e., locating relevant sources) have been the subject of much research in recent years [Kim95, SAD+95, TRV96]. While these problems are daunting in their most general forms, pragmatic approaches that provide useful functionality for many typical situations are starting to appear. In particular, solutions based on the wrapper-mediator model and other non-traditional techniques (eg. [MMM97]), provide the abstractions necessary to implement applications that utilize multiple sources on the Internet.

While significant effort has been placed on addressing the semantic issues of wide-area data access, relatively little effort has been put into solving the performance problems that are inherent in such access. A key performance issue that arises in wide-area distributed information systems is *response-time unpredictability*. Data access over wide-area networks involves a large number of remote data source, intermediate sites, and communications links, all of which are vulnerable to congestion and failures. Such problems can cause significant and unpredictable *delays* in the access of information from remote sources.

Traditional distributed query processing strategies break down in the wide-area environment because they are unable to adapt in response to unexpected delays. Query execution plans are typically generated statically, based on a set of assumptions about the costs of performing various operations and the costs of obtaining data. Due to the apparent randomness of delays when accessing remote data, it is not possible to optimize for such delays *a priori*. Thus, the execution of any statically optimized query plan is likely to be sub-optimal in the presence of the response time problems that will inevitably arise during the query run-time.

1.1 Query Scrambling

To address the issue of unpredictable delays in the wide-area environment, we proposed a dynamic approach to query execution called *query scrambling* [AFTU96]. Query scrambling reacts to unexpected delays by rescheduling, *on-the-fly*, the operations of a query during its execution. In a remote access setting, query scrambling hides delays encountered when obtaining data from the remote sources by performing other useful work, such as transferring other needed data or performing query operations, such as joins, that would normally be scheduled for a later point in the execution.

Query scrambling as defined in [AFTU96] consists of two different phases: a *rescheduling* phase, in which the scheduling of the operators of an active query plan is changed when a delay is detected, and an *operator synthesis* phase in which the query plan is *restructured*, typically by creating new operators that are not in the current query plan. In the original algorithm, both of these phases were *heuristic-driven*. That is, the algorithm was specified as a set of heuristic rules that were activated as delays in obtaining remote data were detected. The heuristics described in that paper were shown to be very effective at hiding delays in some situations, but they were also shown to be prone to making poor scrambling decisions in other cases. In some cases, the proposed heuristics could result in performance that is worse than simply waiting for the delayed data to arrive.

In this paper, we address the shortcomings of the heuristic-based approach by investigating ways to introduce query optimization into the scrambling decision making process. For simplicity, we focus on the problem of *initial delay*, in which delays are manifested as problems in receiving the first tuple from a particular remote source. Initial delays typically arise when there is difficulty in establishing a connection to a remote source or the source is heavily loaded. Also, in the absence of global query optimization (i.e., optimization that considers costs at both the query site and the remote sources such as [RAH+96, TRV96]), initial delays can arise if a remote source must perform a significant amount of work before it can return the first tuple.

1.2 Making Cost-based Decisions

There are a number of fundamental issues that arise when trying to exploit database query optimization technology for scrambling. A basic question is whether the objective function of optimization should be based on *total work or response time*.

Relational optimizers traditionally aim to reduce total work (or “cost”). For example, the cost model of the classic System R-type optimizer includes terms for cpu and disk usage, but does not model the possible overlap of cpu and disk processing [SAC+79]. Likewise, the distributed extensions to this optimizer for the R* system added additional terms for message costs, but also did not model the overlap of such costs [ML86]. In contrast, a response time-based optimizer predicts the overlap of work in addition to the total amount of work [GHK92]. Thus, a response time optimizer might choose a plan with higher total work but more parallelism over a plan with less work but higher sequentiality.

The notion of *delay*, as arises in wide-area remote access is inherently a response time issue. Delays incur no work but still postpone the completion of a given query. We therefore investigate the use of response time-based optimization for query scrambling. In fact, a major result of our work is that if a response time-based optimizer is given an estimate of an expected delay, it can place the access to the delayed data at the proper point in the query execution plan. The quality of such placement of course, depends upon the accuracy of the delay prediction. The current state-of-the-art in delay prediction on the Internet is quite primitive. We therefore investigate two approaches for integrating a response time-based optimizer into the scrambling process. One approach is very aggressive in its scrambling (i.e., it assumes that delays will be long), the other approach is more conservative. In addition, we also develop an algorithm for performing scrambling using an optimizer that is based on total work.

The rest of the paper is organized as follows. Section 2 gives an overview of the cost based query scrambling and describes three approaches for integrating query optimization with scrambling decisions. Section 3 describes the environment used in the experiments. Section 4 analyzes the cost based scrambling algorithms using a two-phase randomized query optimizer, a distributed query processing simulator, and a workload derived from queries of the TPC-D benchmark. Related work is discussed in Section 5. Finally Section 6 presents our conclusions.

2 Cost-based Query Scrambling

2.1 Query Scrambling Overview

In this paper we assume a query execution environment consisting of query sites and a number of remote data sources. The processing work for a given query is split between the query source and the remote sites.[1] Query plans are produced by a query optimizer based on its cost model, statistics, and objective functions. This arrangement is typical of mediated database systems that integrate data from distributed, heterogeneous sources.

An example query execution plan for such an environment is shown in Figure 1. The query involves five different base relations stored at four different sites. In the example, relations A and B reside at separate remote sites, relation C resides locally at the query site, and relations D and E are co-located at a fourth site.

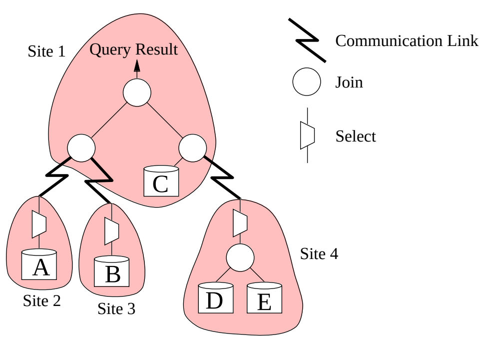

Figure 1: Example of a Complex Query Tree

Using a static scheduling policy, a remote access query plan such as this is susceptible to delays that arise when accessing the remotely-stored data. For example, using an iterator approach [Gra93], the first data access would be to request a tuple of Relation A (from site 2). If there is a delay in accessing that site then the scan of A, and hence the entire query execution, is blocked until the site recovers.

Query Scrambling reacts to such delays in two ways (referred to as Phase 1 and Phase 2 respectively):

- *Rescheduling* - the execution plan of a query can be dynamically rescheduled when a delay is detected. In this case, the basic shape of the query plan remains unchanged (although some additional "materialization" operators may be added as discussed in Section 2.2).
- *Operator Synthesis* - new operators (e.g., a join between two relations that were not directly joined in the original plan) can be created when there are no other operators that can execute. In this case, the shape of the query plan can be significantly modified through the addition, removal and/or re-arrangement of query operators.

[1] As currently specified, query scrambling treats remote sources as black boxes, regardless of how the remote data is computed. Thus, it operates solely at the query sites.

Query scrambling is an iterative process; it works by repeatedly (if necessary) applying these two techniques to a query plan. For example, assume that the query shown in Figure 1 stalls while retrieving tuples of A. Instead of waiting for the remote site to recover, Query Scrambling could perform rescheduling, and retrieve the tuples of B while A is unavailable. These tuples would need to be stored temporarily at the query site. If, after obtaining B, A is still unavailable, then rescheduling could be invoked again, for example, to trigger the execution of (D⋈E) at site 4, and to join this result with C. If at this point, A is still unavailable, then Operator Synthesis can be invoked to create a new join between B and (D⋈E)⋈C. Operators initiated by Query Scrambling may as well experience delays, which may cause Scrambling to be invoked further.

In this paper, we assume that once a scrambling step (i.e., the rescheduling of a query sub-tree, or the execution of a synthesized plan) has been started, the system does not check for the availability of delayed data unless the delayed data is accessed during the step. Once the step has been completed, the arrival of the delayed data is checked. If the delayed data has still not arrived, another iteration of the scrambling algorithm is begun. Likewise, if during scrambling, a delay arises when accessing a remote source, the current scrambling step is abandoned, and a new one is started.

As discussed in Section 1.1, the original formulation of scrambling was heuristic-based [AFTU96]. In this paper, we address the shortcomings of that earlier approach by incorporating query optimization into the scrambling process. The focus of the paper is on Phase 2 of scrambling, but we also apply cost-based decision making in Phase 1. We therefore first briefly describe how cost information is used during Phase 1, and then describe our three approaches for integrating cost information into Phase 2. The comparison of these latter three approaches is then addressed in detail in the remainder of the paper.

2.2 Cost-Based Rescheduling (Phase 1)

Phase 1 starts by identifying *runnable subtrees*, i.e., sub-trees of the plan that are made up entirely of operators that are not currently blocked. A runnable subtree can be scheduled out of order by inserting a "materialization" operator between its root and the root's parent. These materialization operators issue Open, Next, and Close calls to the root of the subtree and save the result in a temporary relation to be used when the result is needed later.

The original scrambling algorithm rescheduled subtrees simply by traversing the query plan from left to right, and choosing to run each "maximal" runnable subtree (i.e., a runnable subtree whose parent operator is blocked) it encountered. This approach was taken because the original algorithm did not use query optimization for scrambling.

In this paper, we use a query optimizer to compute the expected cost (in terms of total work) of each runnable subtree. Three costs are associated with each runnable subtree: 1) MW, the cost of writing the materialized temporary result produced by the subtree to disk; 2) MR, the cost of reading the temporary from disk when it is to be used; and 3) P, the cost of executing the subtree itself. Note that MW and MR represent the *additional* cost incurred by running the subtree out of order.[2] Also note that MW and MR can differ depending on the relative costs of disk writes and reads in a system.

The *efficiency* of each runnable subtree is then computed as $\frac{P-MR}{P+MW}$. Intuitively $P - MR$ is how much work will be saved in the future by scheduling this subtree, and $P + MW$ is the duration of the scrambled operation. Thus, the ratio gives the improvement in response time per unit of scrambled execution. Phase 1 then chooses to re-schedule runnable subtrees in decreasing order of efficiency. Subtrees with efficiency below a certain threshold (currently set at 0.75) are not considered for execution during Phase 1. This approach to using cost favors runnable subtrees that are materialized by the original plan (they have efficiency = 1) and ones that produce small results.

Each iteration of Phase 1 chooses a runnable subtree, runs it, and then checks to see if the delayed data has started to arrive. If so, then scrambling is terminated at this point. If the delayed data is still missing, then another iteration of scrambling is initiated.

2.3 Cost-Based Operator Synthesis (Phase 2)

Phase 2 of query scrambling is invoked after no more progress can be made in Phase 1. Unlike Phase 1, which simply changes the scheduling of existing operators, Phase 2 actually creates a new plan, which typically contains new operators. In this paper, we study three approaches to using a query optimizer during Phase 2.

2.3.1 Optimization Strategies

As stated in Section 1.2, an optimizer can be integrated with scrambling using an objective function based on total work or response time. In scrambling, we deal with these two types of optimization differently.

Response time-based optimizers are naturally suited for query scrambling because they are able to estimate not only the total work to be done for a query, but also how that work can be overlapped. The ability to consider such overlap can be exploited for query scrambling. Simply by telling the optimizer how long a particular data source will be delayed, the optimizer can be coerced into finding plans that perform other useful work that is overlapped with the delay. Of course, the work that the optimizer schedules to overlap the delay can not be in any way dependent on the delayed data. Fortunately, query optimizers must normally deal with such dependencies, in order to generate valid query plans.

There is a problem with the above approach, however. It requires that the duration of the delay of a source to be known *a priori*. Of course, if such knowledge exists, then there is no need for scrambling. When dealing with delays of remote sources on the Internet the current state of prediction is quite primitive (e.g., note the "time remaining" line on the bottom of your browser window). The approach that we take in this paper is to "lie" to the optimizer by providing it a fixed estimate of the expected delay duration when scrambling is invoked. We propose two such approaches. One approach, called *Include Delay*(IN), simply chooses a very long (relative to the query response time) delay duration so that the optimizer will push any accesses to the delayed data as far back in the plan as possible. The other approach, called *Estimated Delay* (ED) initially assumes that delays will be brief, and subsequently increases its estimate if a delay turns out to be longer than the earlier guess.

An alternative to using a response time-based optimizer is to use a more traditional objective function based on total work (i.e., one that ignores potential overlapping). Such an optimizer, however, can not adequately cope with the notion of delay since delay information is not taken into account in the objective function. Our solution to this problem is to simply remove the accesses to the delayed data from the query. In this paper we explore an approach called Pair, which has the optimizer generate plans for individual joins one at-a-time. This policy is a cost-based analogue to the heuristic-based Phase 2 of the original scrambling algorithm. In the following sections we briefly describe these three approaches: Pair, IN, and ED.

2.3.2 Pair

The Pair approach uses a total work-based optimizer to construct a query plan containing only a single join using two relations that are not currently known to be blocked. The optimizer analyzes each pair of non-blocked relations that share a join predicate (i.e., it avoids Cartesian products) and calculates the cost of the best way of joining them. It then chooses the join with the least total cost and executes it. The cost

[2] Thus, MR and MW are zero for subtrees rooted at operators that write their results to disk under the regular query schedule.

computation includes costs associated with materializing the result. The result of this join is materialized to disk, and is available for use in later scrambling iterations. In addition to avoiding Cartesian products, the pair policy also avoids joins that produce a result that takes longer to read from disk than it would take to compute from scratch later. At the end of each join, the policy checks for the arrival of delayed data. If the data has not yet arrived, another iteration is begun. If Pair runs out of qualified joins before the delay terminates, scrambling simply halts and waits for the delayed data to arrive. When all the blocked relations become available the Pair policy constructs a single query tree that computes the result before the normal execution resumes. This is necessary because during scrambling the Pair policy does not try to maintain a complete query plan that represents the final result, but rather works on pairs of relations. This phase, which is called *reconstruction phase*, utilizes the optimizer to find the optimal query plan.

2.3.3 Include Delayed (IN)

In contrast to the Pair approach, IN uses a response time-based optimizer, and thus, each iteration of scrambling generates a *complete* alternative plan. For all delayed data sources, the optimizer is told that the delay duration will be very large (i.e., many times longer than the expected response time of the non-delayed plan). This approach results in delayed accesses being pushed far back in the plan schedule. It is interesting to note however, that the optimizer does not always push such accesses to the very end of the schedule. In some cases, doing so would incur excessive work after the delayed data has arrived, which would result in even worse performance. The optimizer is naturally able to recognize such situations and places delayed accesses in the "right place" in the plan.

One issue that arises when using a response time-based optimizer in this manner is that the optimizer is geared towards choosing the plan that ultimately results in the best response time for the delay value that we give it. Since we really do not know what the delay will be however, this single-mindedness can sometimes be harmful. In general, when the scrambler decides to initiate work in order to hide delay, it commits to performing an entire step. We refer to the duration of the step as *risk* and to the potential improvement in response time as *benefit*. The optimizer chooses the plan with the greatest benefit whose risk can be overlapped with the expected delay duration. This can cause problems when the delay turns out to be relatively short.

In order to address this problem, we introduce a parameter called the "risk/benefit knob" (*RBknob*), that prevents the optimizer from choosing very high-risk plans for relatively small potential gains over lower-risk plans. The *RBknob* is expressed as the ratio of the amount of benefit the optimizer is willing to give up for a given savings in risk. Increasing the *RBknob* has the effect of making the policy more conservative. The performance of this knob is studied in Section 4.3.

2.3.4 Estimated Delay (ED)

The Estimated Delay (ED) approach works similarly to IN except that rather than starting by assuming a huge delay, it first tries relatively short delays, and successively increases its delay estimates if necessary. The motivation behind this approach is that assuming a large delay initially may cause the optimizer to pick a risky plan that has high payoff for long delays, but hurts performance for short delays. Likewise, if the delay estimate is too small, scrambling may be rendered ineffective for larger delays because the optimizer will refuse to run high risk/high pay off plans.

ED works as follows: It starts by picking an estimated delay value equal to the 25 % of the original query response time. Until the delayed data arrives, iterations are repeated with this estimated delay value as long as some progress is being made by each iteration. When this value becomes too small to allow any progress, it is increased to 50 % of the original query response time. Finally when this becomes insufficient, we use a value of 100 % of the original response time. This scheme allows scrambling to first perform iterations with low risk, but still make progress. Thus, in the event that a delay turns out to be short, scrambling has helped rather than hurt. In the event of longer delays, ED becomes more aggressive, which allows it to attempt higher-risk plans. Note that the *RBknob* described for the IN policy is also used for ED, but in general it has less impact here, because ED is already a more conservative policy than IN.

3 Experimental Environment

Our experiments are performed using a detailed simulator of a distributed query processing environment, a two-phase randomized query optimizer, and a workload based on queries from the TPC-D benchmark. We describe each of these in the following sections.

3.1 Simulation Environment

To study the performance of the cost-based scrambling approaches we implemented them on top of a simulator that models a distributed, peer-to-peer database environment and that is capable of realizing iterator-based or process-based scheduling of query operators. In this study, we used only a single join method, namely, hybrid hash [Sha86].

Table 1 shows the main parameters for configuring the simulator and the main settings used for this study. There are two types of sites. Data sources, which store base data that will be used in queries, and Query sites, which execute queries. Every site has a CPU whose speed is specified by the *Mips* parameter, *NumDisks* disks, and a main-memory buffer pool of size *SourceMem* or *QSiteMem*. For the current study, the simulator was configured to have a single query site and six remote data source sites. In all the experiments described in this paper, we placed no additional load on the source sites beyond what was generated by the query requests.

Parameter	Value	Description
NumSources	6	number of data source sites
Mips	200	CPU speed (10^6 instr/sec)
NumDisks	1	number of disks per site
DskPageSize	4096	size of a disk page (bytes)
WriteBufSize	4	size of disk write buffer (pages)
RequestSize	40	size of a data request (bytes)
NetPageSize	4096	size of a data transfer (bytes)
Compare	4	instr. to apply a predicate
HashInst	25	instr. to hash a tuple
Move	2	instr. to copy 4 bytes
QSiteMem	300, 1,000 or 10,000	Query site memory size (pages)
SourceMem	10,000	Data source memory size (pages)
NetBw	12	network bandwidth (Mbits/sec)
MsgInst	20000	instructions to send or receive a message
PerSizeMI	3	instructions per byte sent
DiskInst	5000	instructions to read a page from disk

Table 1: Simulation Parameters and Main Settings

The simulator charges for all the functions performed by query operators like hashing, comparing, and moving tuples in memory, as well as for system costs such as disk I/O processing and network protocol overhead as described below.

Disks are modeled using a detailed characterization and settings adapted from the ZetaSim model [Bro92]. The disk model includes costs for random and sequential physical accesses and also charges for software operations implementing I/Os. The unit of disk I/O for the database is pages of size *DskPageSize*. The disks prefetch pages when reads are performed. In the current version of the simulator, 4 pages are obtained for each read access request made to the disk. In addition to the disk costs, there is a charge of *DiskInst* instructions for each disk access. In our experiments, disks were seen to deliver data at an average rate of approximately 10 Mbits/sec with sequential I/Os, and a rate of approximately 3 Mbits/sec with random I/Os.

In this study, the disk at the query site is used mostly to temporarily store intermediate query results and data obtained from remote sources during a query execution. In addition, some small base relations can also be permanently stored at the query site. The other base relations are stored on disk at the sources as described in Section 3.3. Although sources are configured with memory, the workload used in the experiments here is performed such that the server memory is not useful (i.e., there is no caching across queries and relations at the remote sites are accessed once per query). Thus, in the experiments that follow, base relations are always read (sequentially) from the sources' disks for each query execution. In addition, any selections or projections on base data that are required by a query are performed at the *sources* before the data is shipped to the query site.

At the query site, when scrambling is being used, the buffer manager uses a special policy that pins any memory-resident data that will be used in the next scrambling iteration. Any other cached data is unpinned, and is managed using an LRU policy. In the case that the amount of pinned data prevents the query from allocating the additional buffer space it needs, some of the pinned data is released. In particular, memory-resident relations, which do not have an image on the local disk and are accessed early in the plan are given preference to remain in memory over other relations. Disk writes for intermediate results are buffered in groups of *WriteBufSize* in order to increase the amount of sequential I/O. For join partitions, where memory is at a premium, writes are done one page at-a-time.

The network is modeled as point to point connections between each source and the query site. As such, link failures are independent of each other. The speed of each link (*NetBw* Mbits/sec) is set to be slightly higher than speed of sequential disk access at the data sources, in order to make sure that network speed is not the bottleneck in these experiments. The details of a particular networking technology (e.g., Ethernet, ATM) are not modeled. The cost of sending messages, however, is modeled as follows: the simulator charges for the time-on-the-wire (depending on the message size and the network bandwidth) as well as CPU instructions for networking protocol operations, which consist of a fixed cost per message (*MsgInst*) and a per-byte cost based on the size of the message (*PerSizeMI*). The CPU costs for messages are paid both at the sender and the receiver.

The query execution model uses a page-at-a-time (i.e., non-streaming) approach to remote data access. That is, when an operator running at the query site needs data from a remote source, it sends a request (of *RequestSize* bytes) to that source and waits for the reply. A source responds with a block of *TransferSize* bytes of data. The query site employs prefetching (of one page) to reduce network latency.

In the experiments, delays are modeled by simply blocking the link between a remote source and the query site. Such delays could also be modeled by suspending processing at the source. Since we use point-to-point connections between the sources and the query site, these two methods are equivalent.

3.2 The Query Optimizer

In the study, we use a two-phase randomized optimizer similar to the one described in [IK90, IW87]. The optimizer first runs *Iterative Improvement* (II) algorithm for some time, followed by *Simulated Annealing* (SA). It is possible to trade off the quality of the chosen query plan vs. the optimization time by changing two parameters, *OptTries* and *OptMoves*, which control the number of starting points in the search space and the number of iterations performed on each starting point during the II phase. These variables (shown in Table 2) are scaled with the number of relations used in the query. Thus the number of plans generated by the optimizer increases quadratically with the number of relations. The search space for the optimizer includes left deep, right deep and bushy plans.

Parameter	Value	Description
OptTries	$10 \times NR$	# of starting points in search space
OptMoves	$1 \times NR$	# of iterations performed on a starting point
NR	variable	Number of relations in the query being optimized
RBKnob	0.01	Risk/Benefit knob

Table 2: Optimizer Parameters and Main Settings

The optimizer can use any provided objective function to rate alternative query plans. We use two such functions in this study: one based on total work, and one based on response time. Our response time model is derived from the one defined in [GHK92]. It calculates expected response times by considering potential parallelism in addition to the work done by query operators.

The optimizer takes the following input: 1) information about the relations that participate in the query including cardinalities, fields, and the data source, etc.; 2) the join predicates between the relations; and 3) the selection predicates on the relations together with a selectivity factor for the predicate. Depending on the scrambling approach used, information on all or only a subset of the relations may be given to the optimizer. The Estimated Delay and Include Delayed approaches also provide information about which relations are delayed and an estimate of how long the delay is expected to last. Recall that, as described in Section 2.3, the response time-based approaches also use a special "knob" to control the risk/benefit tradeoffs made by the optimizer. The default value of this knob is 0.01, which is a very aggressive setting. The effect of more conservative settings is investigated in Section 4.3.

3.3 Workload

As stated previously, we examine the performance of the cost-based approaches using queries and a database derived from the TPC-D benchmark. The database is based on a TPC-D Scaling Factor (SF) of 1, and is described in Table 3. The table shows several different data sizes for each relation. In the experiments, we model the effect of projections on the tuples processed at the query site by reducing the size of all projected tuples sent to that site to a fixed amount (40 bytes). Selections and projections are pushed to the data sources where possible. As a result, remote sources read *TuplePages* pages from disk when scanning a relation but transmit only *ProjectedPages* pages to the query source when a projection is applied at the source (and possibly fewer if a selection predicate is also applied). We also use 40 bytes as the size of the tuples produced by a join.

Table	Tuples	Tuple Size	Pages	Projected Pages	Primary Key
Region	5	120	1	1	`regionkey`
Nation	25	120	1	1	`nationkey`
Supplier	10 K	160	400	100	`suppkey`
Customer	150 K	180	6818	1486	`custkey`
Order	1,500 K	100	37500	14852	`orderkey`
Part	200 K	160	8000	1981	`partkey`
Lineitem	6,000 K	120	181818	59406	`orderkey+ linenumber`
PartSupp	800 K	140	28571	7921	`suppkey+ partkey`

Table 3: Database schema and data sizes

Two of the relations (`REGION` and `NATION`) are very small "detail tables" that change very infrequently, so copies of these are maintained and accessed at the query site. The remaining base tables are each placed at a separate remote server.

In terms of queries, we have chosen three of the TPC-D queries (Q5, Q8, and Q9) for our experiments. These queries were chosen because they are fairly complex (6 to 8-way joins) so they provide significant opportunities for interesting scrambling behavior. Because our simulator does not model aggregate functions, GROUP BY and ORDER BY clauses, or sub-queries, we have modified the original queries slightly. The modified versions are described in Section 4. It should be noted that our goal in using TPC-D queries as a starting point is to allow us to examine the approaches using realistic join graphs, cardinalities, and selectivities; we do not claim to draw any conclusions about performance on an actual TPC-D benchmark.

3.4 Experimental Methodology

All the graphs shown in the following section plot the duration of an initial delay of a remote source vs. the response time achieved with each of the scrambling approaches. The results for all of the approaches tend to exhibit a step behavior due to the iterative nature of the scrambling process. The graphs were generated as follows.

First, for each combination of query, memory allocation, and delayed relation, we ran each scrambling approach with a very long delay to find the delay duration (i.e., the point on the x-axis) where each iteration would occur for that approach. For all the delay durations in the interval between two such points, the query response time will be the same. This run, however does not show what the value of that response time will be. We therefore pick one delay value within each interval and run the scrambling approach to obtain the response time with a delay of that duration. This response time is the response time for the entire interval.

Because we are using a randomized optimizer, we needed to be careful that both the initial plans and the scrambled plans that were generated were good plans. Otherwise, particularly bad plans could result in spurious effects that were not due to the scrambling approaches. To ensure that we had good plans we did the following: First, we generated the intervals using higher values for the II parameters to increase the thoroughness of the search. These runs were repeated three times to ensure the repeatability of the scrambling iteration intervals. Then, we also ran each data point (i.e., combination of delay interval, scrambling approach, query, memory allocation, and delayed relation) at least three times (using the normal optimizer parameter settings), and checked that the plans generated at each point conformed to the intervals found initially, and that the generated response times were accurate to within plus or minus 2%.

Num.Relations	Optim.Time
4	1.290 secs
5	2.528 secs
6	5.109 secs
7	8.308 secs
8	11.657 secs

Table 4: Optimization times for various numbers of base relations.

The results that we report in Section 4 do not include the time required for running the query optimizer. The goal was to avoid mixing numbers from the *real* query optimizer with those from a *simulated* system. As can be seen in Table 4, the optimization times obtained with our optimizer (on an IBM RS/6000 42T PowerPC) are quite small compared to the 750+ second-response times of the queries. These times were obtained using an optimizer that was not tuned to reduce optimization time. Thus, we would expect to be able to lower them even further if necessary, for example, in order to handle queries larger than 8-way joins.

4 Experiments and Results

4.1 Experiment 1 - National Market Share

We begin by studying the performance of the three cost-based approaches when delays are encountered during the execution of a modified version of TPC-D Query Q8, the National Market Share Query (referred to as MQ8). The SQL statement for MQ8 is shown in Figure 2.

MQ8 is an 8-way join query, with selections on the `REGION`, `ORDER`, and `PART` relations. Figure 3 shows the *query graph* corresponding to this query. In the query graphs, we abbreviate relation names using their first letters.[3] An edge between two relations indicates a join predicate between those relations in the query; the edge is labeled with the join attribute(s). Selection predicates are indicated by boxes con-

```
SELECT  O.ORDERDATE, L.EXTENDEDPRICE, N2.NAME
FROM    PART, CUSTOMER, ORDER, LINEITEM,
        SUPPLIER, NATION N1, NATION N2, REGION
WHERE   P.PARTKEY        = L.PARTKEY
        AND L.SUPPKEY    = S.SUPPKEY
        AND O.ORDERKEY   = L.ORDERKEY
        AND C.CUSTKEY    = O.CUSTKEY
        AND C.NATIONKEY  = N1.NATIONKEY
        AND N1.REGIONKEY = R.REGIONKEY
        AND R.NAME       = 'EUROPE'
        AND S.NATIONKEY  = N2.NATIONKEY
        AND O.ORDERDATE BETWEEN '94-01-01'
                            AND '95-12-31'
        AND P.TYPE       = 'SMALL PLATED STEEL'
```

Figure 2: Modified National Market Share Query (MQ8)

taining the selection attribute(s) and the selectivity of the predicate is listed as a fraction below the selection box.

We run the query using each of the three different memory allocations identified in table 1. Because we assume that selections and projections (where appropriate) are applied at the remote data sources, the amount of data that must be processed at the query site is significantly less than the sum of the raw relation sizes. In this experiment, a memory allocation of 1000 pages is more than sufficient to run MQ8 with no hash partitioning. When the smaller memory allocation (300 pages) is used, the two largest (intermediate) relations (the result of `PART`⋈`LINEITEM` and `ORDER`) must be partitioned in order to be joined.

The initial query plans generated by the optimizer for the 1000 page and 300 page allocations are shown in Figures 4(a) and 4(b), respectively (the initial plan for a memory size of 10,000 is identical to that of Figure 4(a)). All binary operations shown in the figures are *hash joins*, and the bold edges indicate joins which require the relations to be partitioned.

We now turn to the results of the experiment. Figures 5 and 6 show the results for MQ8 with 1000 page and 300 page memory allocation, respectively. In all of the graphs shown in this paper the x-axis indicates the initial delay (in seconds) of a remote relation (in this case, the `PART` relation) and the y-axis indicates the query response time (in seconds). In addition to the curves for each of the scrambling approaches, the graphs also contain two parallel diagonal lines. The lower line simply indicates the magnitude of the delay. Since a query cannot complete until all of the relevant data has been accessed, this "delay" line represents a lower-bound on response time. The higher diagonal line, labeled "No Scr" represents the response time that would be obtained if scrambling is not used.

In both cases, we delay `PART`, which is the left-most relation of the optimized query plan. `PART` is a very valuable relation in this plan for two reasons. First, since the iterator execution model activates operators in a pre-order manner, delaying the left-most relation leaves the most possible remaining work to be done in the absence of scrambling. Second, in

[3] In this query the detail relation `NATION`, which is kept at the query site is used twice. We refer to these uses as `N1` and `N2`.

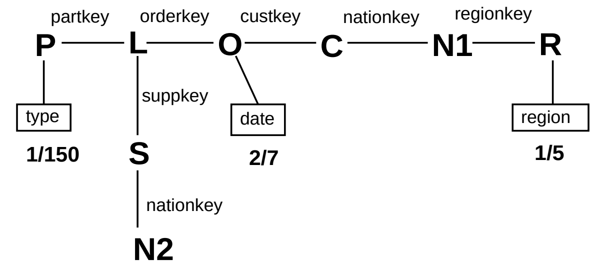

Figure 3: Query graph for query MQ8

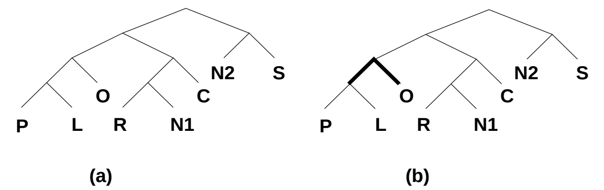

Figure 4: Query plans for (a) memory≥1000 and (b) memory=300

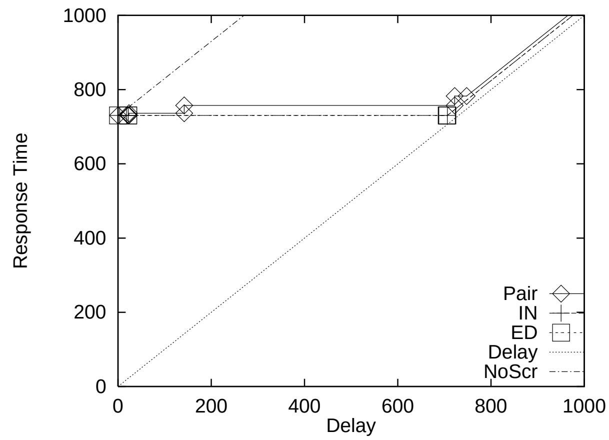

Figure 5: MQ8, PART delayed, Memory=1000

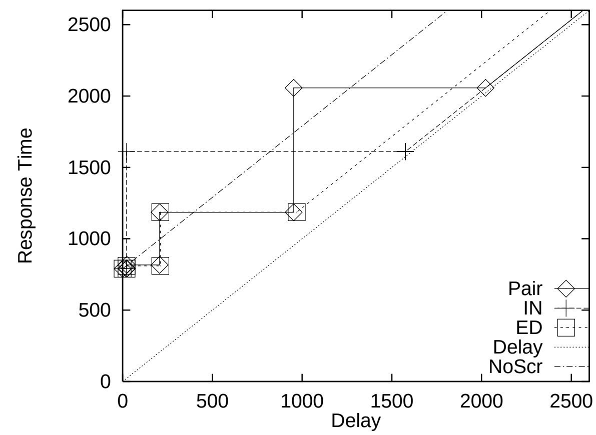

Figure 6: MQ8, PART delayed, Memory=300

query MQ8, the `PART` relation plays the role of a *reducer* for `LINEITEM`, the largest relation in the schema. That is, because of its selection predicate and the fact that it participates in a functional join with `LINEITEM` (and assuming uniform distribution and independence of the join attribute values), the selection on `PART` reduces the size of the intermediate result, `PART`⋈`LINEITEM` by the selectivity of its selection predicate. Thus, the presence of `PART` is important here, because it significantly reduces the number of tuples that must be processed later in the query.

4.1.1 *Query MQ8 - Large Memory Allocations*

Turning to the 1000 page case (Figure 5), it can be seen that all three of the cost-based scrambling approaches are very effective at hiding the delay of `PART`. In fact, IN and ED are able to effectively hide nearly 100% of the delay here; when the delay is 706 seconds or less (i.e., up to the knee in the curve) the response time is virtually unchanged from the non-delayed value of 730.5 seconds. In other words, the two approaches are able to effectively hide a delay that is nearly equal to the original response time of the entire query in this case. Beyond a delay of 706 seconds, scrambling has run out of additional work to perform, so the response time increases parallel to the delay. The difference between the response time lines and the delay line represents the amount of work that must be done after the delayed tuples begin to arrive. Note that the Pair approach also does well here; it performs slightly worse than IN and ED because it materializes some intermediate results to disk.[4]

In this case the first phase is run when scrambling starts, materializing the subtree that contains relation `REGION`, `NATION1` and `CUSTOMER`. After this iteration the second phase begins. In this case, all of the approaches are able to find alternative plans that perform well. As shown in Figures 3 and 4, the initial plan basically traverses the graph from left to right. When `PART` is delayed this traversal becomes impossible. It is, however, possible to start at the other end of the query graph and traverse from right to left. This traversal picks up the other reducers in the query (the result computed in the first phase which contains the reducers `REGION`, and `ORDER`) before accessing the large `LINEITEM` relation.

In contrast to the cost-based approaches, the original heuristic-based scrambling algorithm [AFTU96] follows the policy of executing the left-most runnable sub-tree of the query plan, which in this case, results in joining `LINEITEM` and `ORDER` in the absence of the other reducer, `REGION` (the heuristic-based algorithm is not shown in the figure). With 1000 pages of memory, this join requires partitioning, which results in a large performance hit. For this experiment, the response time obtained with the heuristic-based approach jumps to 1621 seconds, for delays between 23 and 1542 sec-

[4] The extra materializations are due to our particular implementation of Pair on our simulator, and could be avoided by using a more sophisticated memory management approach in the simulator.

onds long. Thus, for many delay values, that algorithm performs significantly worse than simply waiting for the delayed relation to arrive.

We also ran this experiment for a memory allocation of 10,000 pages (not shown). In this case all of the cost-based approaches performed identically to the 1000 page case (because 1000 pages is sufficient to run the scrambled plans without partitioning). The main difference was that with this large memory (i.e., approximately 33 times more than what was allocated to the original query plan), the heuristic-based algorithm performed roughly as well as the cost-based approaches.[5] Its better performance here is due to the fact that the extra memory allows even the inefficient joins that it picks to run without partitioning. In this environment, the CPU costs are a negligible portion of the query execution time so avoiding partitioning leads to reasonable performance. The tradeoffs for the heuristic-based algorithm are similar for the rest of the experiments in this study, so we do not show any further results for that algorithm. Rather, we focus on the tradeoffs among the three cost-based approaches.

4.1.2 Query MQ8 - *Small Memory Allocation*

With the smaller memory allocation (300 pages, shown in Figure 6), the story changes significantly. 300 pages is sufficient to run the query when there are no delays with reasonable efficiency (the response time here is 790.5 seconds, only slightly higher than in the larger memory case). The smaller memory, however, causes problems when scrambling is required and results in different performance for the various approaches. Query scrambling starts by running the first phase on the subtree containing `REGION`, `NATION1`, and `CUSTOMER`. This result will be used by the subsequent steps of the different scrambling algorithms.

The IN approach is the most aggressive — it assumes that the delay of `PART` will be long, so it is willing to initiate a lot of scrambling work in order to be able to hide more delay. In this case, the IN approach simply pushes `PART` to the far right of the query plan, and joins the remaining relations in the same order as in the previous case. These joins are more expensive here, however, because the lack of memory results in more partitioning and thus, more local I/O. For delays less than 970 seconds, IN has the worst performance of the scrambling approaches, even performing worse than not scrambling for much of that range. Ultimately, however, at a delay of 1576 seconds, IN manages to perform nearly all of the work of the query during the delay, so its performance becomes nearly the same as the delay.

The ED approach is more conservative here. It begins by joining the result computed during the first phase (i.e. `REGION` $\bowtie$ `NATION1` $\bowtie$ `CUSTOMER`) with `ORDER`, and `NATION2` with `SUPPLIER` (this is the second step in the curve). It then brings `LINEITEM` over, writes it to the local disk, and waits for `PART` to arrive (at this point, its curve goes diagonal). This more conservative behavior results in better performance than IN for shorter delays, but ultimately worse performance for longer delays, since more work remains to be done when the tuples of `PART` eventually begin to arrive. Finally, Pair initially performs the same steps as ED, but at 953 seconds (roughly when ED stops scrambling) it chooses to perform a join that includes `LINEITEM`, which requires partitioning, and hence, is quite expensive. Given a long enough delay, this additional join will eventually pay off, with Pair having similar performance to IN after a delay of 2022 seconds. It is interesting to note that while Pair performs the same joins as IN, it has worse performance than IN in the delay range of 953 to 2022 seconds. This worse performance arises because Pair does not generate a complete plan, but rather, makes local decisions one pair at a time. It is not able to make intelligent decisions on how to partition the results of its joins, because it does not know if or how those results will be used in a subsequent operation. In other words Pair policy does not recognize interesting orders and therefore cannot asses the future savings due to executing a slightly more expensive plan which generates an interesting order. As a result, Pair simply materializes its intermediate results, and re-reads them to partition them later if necessary. This re-partitioning is expensive, because it generates significant amounts of *random* I/O.

[5] Actually, in this case the original heuristic based algorithm works well with as few as 5000 pages, or 16.67 times more memory than allocated to the original query.

4.2 Experiment 2 - Local Supplier Volume

We now turn to our second set of experiments, which uses a modified version of TPC-D Query 5 that we call MQ5. The SQL for this query is shown in Figure 7. This query is a 6-way join with two selection predicates. As shown in figure 8, the query graph of MQ5 contains a cycle, unlike the "chain" graph of MQ8 in the previous experiments.

```
SELECT N.NAME,L.EXTENDEDPRICE*(1-L.DISCOUNT)
FROM   CUSTOMER, ORDER, LINEITEM,
       SUPPLIER, NATION, REGION
WHERE  C.CUSTKEY       = O.CUSTKEY
       AND O.ORDERKEY  = L.ORDERKEY
       AND L.SUPPKEY   = S.SUPPKEY
       AND C.NATIONKEY = S.NATIONKEY
       AND S.NATIONKEY = N.NATIONKEY
       AND N.REGIONKEY = R.REGIONKEY
       AND R.NAME      = "AMERICA"
       AND O.ORDERDATE BETWEEN '95-01-01'
                           AND '95-12-31'
```

Figure 7: Modified Local Supplier Volume Query (MQ5)

Figures 9(a) and 9(b) show the initial optimized query plans for memory allocation of 1000 pages or greater and a memory allocation of 300 pages, respectively. Notice that

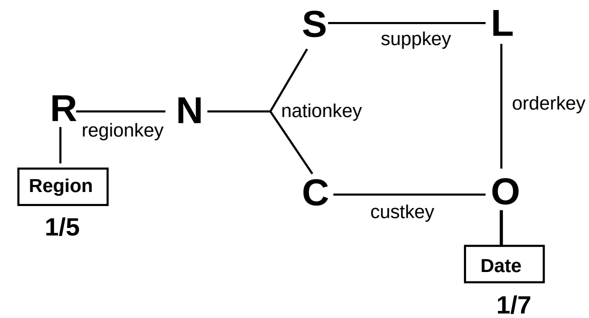

Figure 8: Query graph for query MQ5

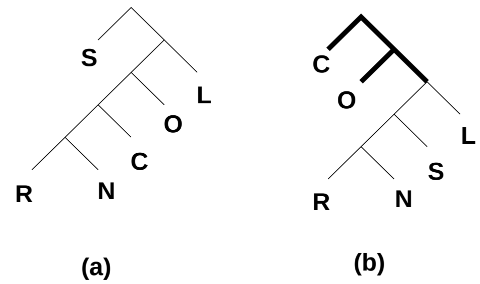

Figure 9: Query plans generated for (a) memory≥1000 and (b) memory=300

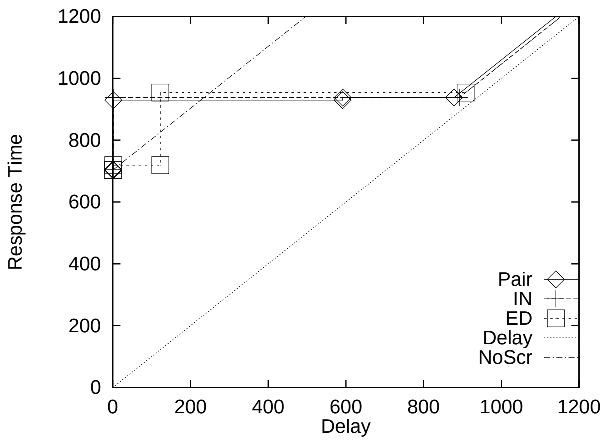

Figure 10: MQ5, CUSTOMER delayed, Mem=1000

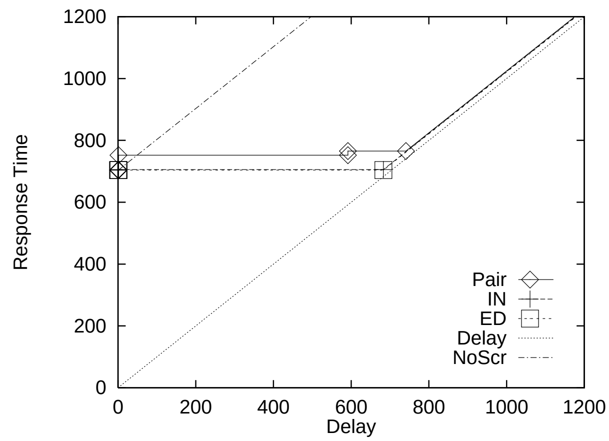

Figure 11: MQ5, CUSTOMER delayed, Mem=10000

with the larger memory allocation, the bulk of the query execution proceeds in a counter-clockwise direction around the join cycle. For the smaller allocation the execution proceeds in the opposite direction, and two of the joins require partitioning. In both memory cases we delay `CUSTOMER`. In the large memory case, the hash table for `SUPPLIER` will be built in the memory and `REGION` ⋈ `NATION` will be computed before `CUSTOMER` is determined to be blocked. In the small memory case the delayed relation is encountered immediately.

4.2.1 *Query MQ5 - Large Memory Allocations*

Figure 10 shows the results for an initial delay on `CUSTOMER` with a memory allocation of 1000 pages. `CUSTOMER` is an important relation in the query plan because it helps transmit the selection predicates on `REGION` and `ORDER` to the large `LINEITEM` relation. In this case, all of the algorithms provide benefits over not scrambling beyond a delay of approximately 250 seconds, and hide nearly all of the delay when the duration is about 950 seconds. The Pair and IN approaches perform similarly here because they basically execute the same operations, even though Pair produces its plan one join at-a-time. Both of the approaches perform a join that requires partitioning (because it involves `ORDER` without first reducing it by joining it with `REGION` as is done in the initial plan). As a result, the response time of the scrambled plan is approximately 234 seconds longer than that of the initial plan (with no delay). In contrast, ED performs better for short delays (up to 122 seconds) due to its conservative approach. It first brings `ORDER` from the remote site and stores it on the local disk before committing to any other scrambling moves. Fetching `ORDER` has little risk for short delays; if `CUSTOMER` arrives during this time, `ORDER` can be used later in the query. ED's slight performance penalty between delays of 122 seconds and 908 seconds results from the fact that `ORDER` needs to be repartitioned if `CUSTOMER` does not arrive in time and ED re-scrambles. Thus, ED pays a small penalty for its conservative approach in this delay range in order to win its advantage for small delays.

The results for the 10,000 memory allocation are shown in Figure 11. In this case the same initial plan is used, but in the presence of delay, different scrambling plans are produced. Because of the larger memory, these scrambled plans can be executed with no partitioning. Thus, they all provide excellent protection from delays of `CUSTOMER`, up to approximately 683 seconds (or 97% of the non-delayed query response time). As was seen in the experiments with MQ8, Pair pays a slight penalty due to unnecessary materializations of temporary results to the local disk.

4.2.2 Query MQ5 - Small Memory Allocation

We also experimented with query MQ5 using the small (300 pages) memory allocation when CUSTOMER is delayed (not shown). In this case all of the scrambling approaches perform well (i.e., the response time is basically flat for delay durations up to the response time of the non-delayed query). This is because the first phase of scrambling (i.e., rescheduling) is able to perform all of the joins except the one involving CUSTOMER (see Figure 9(b)) without needing to create any new operators. That is, the entire sub-tree which computes ORDER $\bowtie$ REGION $\bowtie$ NATION $\bowtie$ SUPPLIER $\bowtie$ LINEITEM is simply executed by rescheduling. Thus, the different cost-based approaches do not come into play here.

4.3 Experiment 3 - Product Type Profit Measure

The third (and final) set of experiments we describe were performed using a modified version of TPC-D Q9, shown in Figure 12. The query graph for this query is shown in Figure 13. In this case, we obtained similar results for all three memory sizes, so we show results only for the 10,000 page allocation. The initial query plan in this case is shown in Figure 14. For this experiment, we delay PART, the only reducer in the query.

```
SELECT N.NAME, O.ORDERDATE.YEAR,
       L.EXTENDEDPRICE*(1-L.DISCOUNT) -
          (PS.SUPPLYCOST * L.QUANTITY)
FROM   PART, SUPPLIER, LINEITEM, PARTSUPP,
       ORDER, NATION
WHERE  S.SUPPKEY       = L.SUPPKEY
       AND PS.SUPPKEY  = L.SUPPKEY
       AND PS.PARTKEY  = L.PARTKEY
       AND P.PARTKEY   = L.PARTKEY
       AND O.ORDERKEY  = L.ORDERKEY
       AND S.NATIONKEY = N.NATIONKEY
       AND P.NAME LIKE '%magenta%'
```

Figure 12: Modified Product Type Profit Measure Query (MQ9)

The performance of the scrambling approaches for varying delays of PART is shown in Figure 15. The IN approach joins all relations other than PART, which in query MQ9, results in intermediate results that are 20 times larger than if PART had been used. These intermediate results propagate through the entire plan, resulting in a very high-risk move. The Pair approach performs the same joins as IN, but performs them one-at-a-time. This has two effects: it reduces risk for short delays, but also incurs additional overhead for long delays, due to the need to partition intermediate results that it has saved to the local disk (a similar effect was seen in query MQ8). The more conservative ED approach performs much better than IN and as good as the Pair policy for delays up to about 2000 seconds (i.e., more than twice the non-delayed response time). ED avoids joining large relations, choosing rather to simply wait for the delayed relation beyond a certain point. As usual this conservatism results in a penalty for longer delays, but in this case, the penalty is quite small, and is more than outweighed by the advantages for shorter delays.

The preceding experiment demonstrated clearly the potential benefits of making conservative scrambling decisions. Recall that the ED and IN approaches both incorporate a "risk/benefit" knob (which is used by the response time based optimizer), that prevents the policies from choosing very high-risk plans for relatively small potential gains over lower-risk plans. In all of the experiments described so far, this knob was set at 0.01 (as described in Table 2), which means that the optimizer is willing to give up 0.01 units (e.g., seconds) of potential benefit for long delays to get a plan whose total work is 1 unit less, which results in less risky behavior for short delays.

To study the effect of the setting of this knob, we repeated this experiment using several different values (0, 0.10, 0.20, and 0.30). The results using 0 and 0.10 were similar to the results using the default setting (0.01). The more conservative settings did have an impact however. The results with a setting of 0.30 are shown in Figure 16 (the ones using 0.20 are similar). First, the knob has no effect on the Pair approach, because that approach is based only on total work; it has no notion of risk vs. benefit. In contrast, both the ED and IN approaches are effected, but the more conservative setting has a greater impact on IN. Overall, the conservative setting results in substantially better performance for IN with short delays. For example, for delays between 1 second and 837 seconds, the response time using IN is 1184 seconds, and for delays between 837 and 1073 seconds the response time using IN is 1300 seconds, compared to 2263 seconds for these ranges with the aggressive setting. For these additional benefits, IN pays only a 46 second cost in terms of the amount of delay that it can hide for long delays, and this cost only arises for delays over 2000 seconds. Thus, the advantages for short delays clearly outweigh the costs at higher delays. It is interesting to note that with the conservative knob setting, IN requires a second scrambling iteration (for delays greater than 837 seconds), because its first iteration produces a plan that leaves work remaining to be done, even without the delayed relation.

The more conservative setting has a lesser effect on the ED approach. This is because ED already favors conservative decisions for small delays. In this case, the higher knob value prevents ED from performing its last iteration, which has an ultimate benefit of 5 seconds, at a risk of 472 seconds. This results in ED having better performance here for delays between 1072 and 1539 seconds than it did with the more aggressive knob setting.

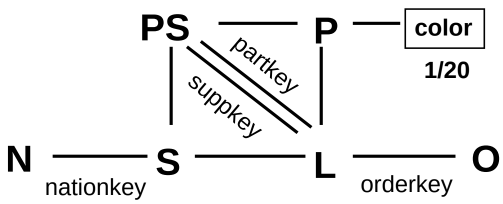

Figure 13: Query graph for query MQ9

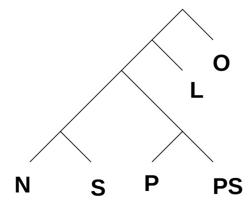

Figure 14: Query plan generated for memory=10000

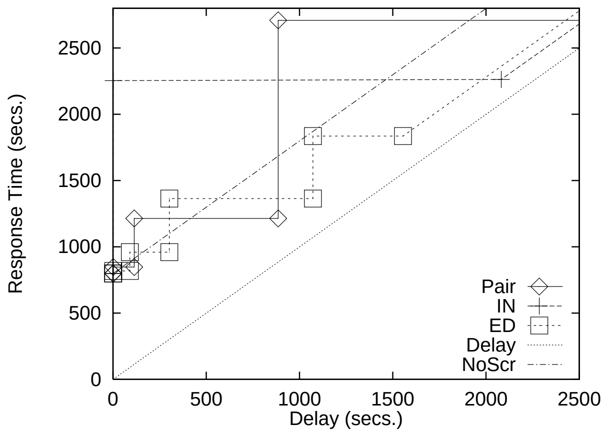

Figure 15: MQ9, PART delayed, Mem=10000, RBknob = 0.01

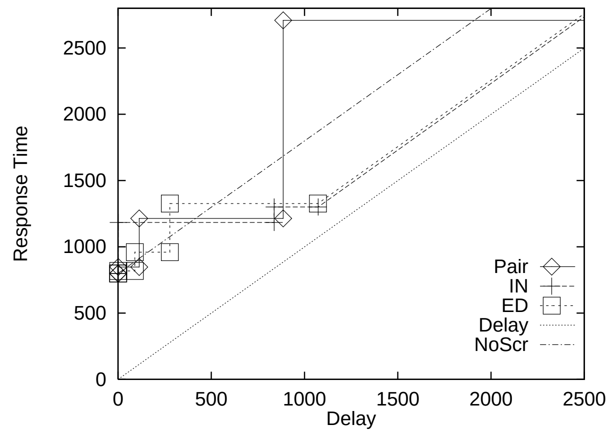

Figure 16: MQ9, PART delayed, Mem=10000, RBknob = 0.30

4.3.1 Summary of Results

The experiments we have described in this section demonstrate several important results for cost-based query scrambling. We briefly summarize those results here. First, the experiments showed that with sufficient memory, all of the cost-based approaches are able to effectively hide initial delays for realistic data processing queries. When the delayed relation is encountered early in the query execution, a delay as long as the normal (non-delayed) response time of the query can be almost completely absorbed. In contrast, the original heuristic-based algorithm can actually perform significantly worse than simply waiting for the delay to end unless substantial extra memory is dedicated to scrambling.

Second, for the cost-based approaches, in the absence of a reasonable prediction of delay duration there is a tradeoff between conservative approaches, which are safer for short delays, and more aggressive approaches which lead to bigger savings in the event of long delays. In general, the amount of delay that can be hidden by scrambling (in the absence of creating additional parallelism, as is discussed in [AFT98]) is limited by the normal response time of the query. This is because scrambling hides delays by performing other useful work, so its ability to hide delay is limited by the amount of useful work that can be done. Thus, as the delay increases beyond the normal response time of the original query, the benefits of scrambling *as a percentage of total execution time* begin to decrease. This argument would lead towards favoring more conservative policies rather than taking larger risks.

Third, as the memory available for scrambling is reduced, scrambled plans in general become more expensive and hence, a longer delay duration may be required in order for scrambling to pay off. Thus, in a low-memory situation scrambling becomes less conservative, and therefore, in the absence of predictions of delay durations, more dangerous.

Fourth, we showed how the aggressiveness of the IN and ED policies can be adjusted through the use of a parameter that tells the optimizer to give up potential gains for long delays in order to reduce risk for short delays. As stated above, this tradeoff makes sense in the absence of reasonably accurate predictions of delay durations.

A final important result from these cases, is that approaches that lack a global view of the scrambled plan (e.g., Pair) may perform unnecessary work. By considering only pairs of relations the Pair policy, is unable to pick slightly suboptimal plans that generate interesting orders. In order to have a complete (and reasonable) scrambled plan, however, one must use an optimizer that uses *response time* as its objective function. A response time-based optimizer allows the delayed relation to be placed at its proper point in the plan (for a given predicted delay), which allows a complete alternative query plan to be generated.

We have also conducted experiments using more than one delayed relation on synthetically generated queries. We have found that when more relations are delayed, the risk associated with each scrambling decision is increased, favoring more conservative algorithms.

5 Related Work

We now briefly discuss related work. The Volcano optimizer [CG94, Gra93] provides dynamic query scheduling by introducing choose-plan operators into a query plan above a set of alternative subplans in order to compensate for the lack of information about system parameters at compile time. At query startup time the appropriate subplan is chosen depending on the current value of the parameters. [INSS92] proposes a related approach that generates multiple alternative plans, and chooses among them when the query is initialized. Neither of these approaches, however, can adapt to changes in the system parameters that occur *during* the query execution.

Rdb/VMS uses a different approach as described in [Ant93]. In this approach, multiple different executions of the same logical operator are started at the same time. When one execution of an operator is determined to be better, the other execution is terminated, and the winner is executed to completion.

The work most closely related to ours is the MIND heterogeneous database project [ONK+97], which performs optimization during the query execution. A query is divided into subqueries and each subquery is sent to a participating site for execution. The results are then composed incrementally by dynamically introducing operators that process them as the results arrive. As such, their algorithm resembles our Pair algorithm with a different set of heuristics that rely on statistical techniques in order to avoid bad decisions.

In [DSD95] the response time of queries is improved by reordering left-deep join trees into bushy join trees and creating subtrees without increasing the cost. Several reordering algorithms are presented. Although this work is limited to left-deep queries and assumes that reordering is done entirely at compile time, one can still use it to *bushify* the plans during run time, possibly at the expense of a slight increase in total work. Bushy plans are generally less vulnerable to delays since different branches of tree can be found that are not directly affected by the delayed relations; such subtrees can be executed independently.

The research prototype Mermaid [CBTY89] and its commercial successor InterViso [THMB95] are heterogeneous distributed databases that perform dynamic query optimization. Mermaid constructs its query plan entirely at run-time, thus each step in query optimization is based on dynamic information such as the intermediate relation cardinalities and system performance. Mermaid neither takes advantage of a statically generated plan nor does it dynamically account for a source which does not respond at run-time.

6 Conclusions and Future Work

In this paper, we proposed and investigated three different approaches to using a query optimizer to help make intelligent choices during query scrambling. Two of the approaches used an optimizer with an objective function based on response time, while the other approach used a more traditional optimizer based on total work. In general, the use of a response time optimizer has the advantage of being able to construct complete query execution plans that include access to delayed data. Based on an estimate of the expected delay duration, the optimizer places the accesses to delayed data to the proper place in the plan.

Given the poor state of current estimation techniques for wide-area data access, we proposed two different ways of using a response-time optimizer. We demonstrated that these approaches exhibit fundamental tradeoffs between risk aversion (for short delays) and the ability to hide large delays. However, we also showed that in many cases the algorithms were very effective at hiding delays over a wide range. In the best cases, the approaches were able to hide delays of a duration equal to the response time of the query in a non-delayed situation.

Due to the growing importance of wide-area data access, particularly in chaotic environments such as the Internet, there is much future work that can be done on scrambling and related dynamic techniques. First, although not discussed in this paper, the scrambling techniques we have described here can be adapted for use with other types of delay, such as bursty arrival, in which sites repeatedly stall and recover. As described in [AFT98] such delays introduce a number of scheduling and memory management issues that must be addressed by scrambling. In addition, we would like to investigate the use of delay prediction techniques in the scrambling approaches. Finally, as described in [ABF+97], additional techniques are required for dealing with very long periods of outage. Unlike scrambling, these techniques necessarily change the answer that is returned to the user, and thus, raise a number of interesting semantic questions in addition to the performance-oriented questions that we have addressed here.

Our current focus is on incorporating the cost based query scrambling into the query engine of PREDATOR [SLR97] and extending it by adding remote access capability. We plan to use this system as a test bed for query scrambling over the Internet.

References

[ABF+97] L. Amsaleg, P. Bonnet, M. Franklin, A. Tomasic, and T. Urhan Improving Responsiveness for Wide-Area Data Access. *IEEE Data Engineering Bulletin*, Vol. 20, No. 3.

[AFT98] L. Amsaleg, M. J. .Franklin, and A. Tomasic. Dynamic Query Operator Scheduling for Wide-Area Remote Access. *Journal of Distributed and Parallel Databases*, Vol. 6, No. 3, July 1998.

[AFTU96] L. Amsaleg, M. J. .Franklin, A. Tomasic, and T. Urhan. Scrambling Query Plans to Cope With Unexpected Delays. *PDIS Conf.*, Miami, USA, 1996.

[Ant93] G. Antoshenkov. Dynamic Query Optimization in Rdb/VMS. *ICDE Conf.*, Vienna, Austria, 1993.

[Bro92] K. Brown. Prpl: A database workload specification language. Master's thesis, University of Winsconsin, Madison, Winsconsin, 1992.

[CBTY89] A. Chen, D. Brill, M. Templeton, and C. Yu. Distributed Query Processing in a Multiple Database System. *IEEE Journal on Selected Areas in Communications*, 7(3), 1989.

[CD96] M. J. Carey, D. J. DeWitt. Of Objects and Databases: A Decade of Turmoil. *22nd VLDB Conf.*, Bombay, India, 1996.

[CG94] R. Cole and G. Graefe. Optimization of Dynamic Query Execution Plans. *ACM SIGMOD Conf.*, Minneapolis, MN, 1994.

[DSD95] W. Du, M. Shan, and U. Dayal. Reducing Multidatabase Query Response Time by Tree Balancing. *ACM SIGMOD Conf.*, San Jose, CA, 1995.

[GHK92] S. Ganguly, W. Hasan, R. Krishnamurthy. Query Optimization for Parallel Execution. *ACM SIGMOD Conf.*, San Diego, CA, 1992.

[Gra93] G. Graefe. Query Evaluation Techniques for Large Databases. *ACM Computing Surveys*, 25(2), 1993.

[IK90] Y. E. Ioannidis, and Y. Kang. Randomized algorithms for optimizing large join queries. *ACM SIGCOM Conf.*, Atlantic City, NJ, 1990.

[INSS92] Y. E. Ioannidis, R. T. Ng, K. Shim, and T. K. Sellis. Parametric Query Optimization. *18th VLDB Conf.*, Vancouver, BC, Canada, 1992.

[IW87] Y. E. Wong, and E. Wong. Query optimization by simulated annealing. *ACM SIGMOD Conf.*, San Fransisco, CA, 1987.

[Kim95] M. Kim. *Modern Database Systems: the Object Model, Interoperability, and beyond*. ACM Press, New York, NY, 1995.

[MMM97] A. O. Mendelzon, G. A. Mihaila, T. Milo. Querying the World Wide Web. *PDIS Conf.*, Miami, USA, 1996.

[ML86] L. F. Mackert, G. M. Lohman R* Optimizer validation and Performance Evaluation for Distributed Queries *12th VLDB Conf.*, Kyoto, Japan, 1986.

[ONK+97] F. Ozcan, S. Nural, P. Koksal, C. Evrendilek, A. Dogac. Dynamic Query optimization in Multidatabases. *Data Engineering Bulletin*, Vol. 20, No. 3, September, 1997.

[RAH+96] M. T. Roth, M. Arya, L. M. Haas, M. J. Carey, W. F. Cody, R. Fagin, P. M. Schwarz, J. Thomas, E. L. Wimmers. The Garlic Project. *ACM SIGMOD Conf.*, Montréal, Canada, 1996.

[SAC+79] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access Path Selection in a Relational Database Management System. *ACM SIGMOD Conf.*, Boston, USA, 1979.

[SAD+95] M. Shan, R. Ahmen, J. Davis, W. Du, and W. Kent. *Modern Database Systems: The Object Model, Interoperability, and Beyond*, chapter Pegasus: A Heterogeneous Information Management System. ACM Press, 1995.

[Sha86] L. D. Shapiro. Join Processing in Database Systems with Large Main Memories. *ACM Trans. on Database Systems*, Vol. 11, No. 3, 1986.

[SLR97] P. Seshadri, M. Livny, R. Ramakrishnan. The Case for Enhanced Abstract Data Types. *23rd VLDB Conf.*, Athens, Greece, 1997.

[THMB95] M. Templeton, H. Henley, E. Maros, and D. Van Buer. InterViso: Dealing with the Complexity of Federated Database Access. *VLDB Journal*, 4(2), 1995.

[Tra97] Transaction Processing Council. TPC Benchmark D Standard Specification, Rev. 1.2.3.

[TRV96] A. Tomasic, L. Raschid, and P. Valduriez. Scaling Heterogeneous Databases and the Design of DISCO. *ICDCS Conf.*, Hong Kong, 1996.

Attachment 1j

Scrambling Query Plans to Cope With Unexpected Delays

Laurent Amsaleg[*†]
University of Maryland
amsaleg@cs.umd.edu

Michael J. Franklin[†]
University of Maryland
franklin@cs.umd.edu

Anthony Tomasic
INRIA
Anthony.Tomasic@inria.fr

Tolga Urhan[†]
University of Maryland
urhan@cs.umd.edu

Abstract

Accessing data from numerous widely-distributed sources poses significant new challenges for query optimization and execution. Congestion and failures in the network can introduce highly-variable response times for wide-area data access. This paper is an initial exploration of solutions to this variability. We introduce a class of dynamic, run-time query plan modification techniques that we call query plan scrambling. We present an algorithm that modifies execution plans on-the-fly in response to unexpected delays in obtaining initial requested tuples from remote sources. The algorithm both reschedules operators and introduces new operators into the query plan. We present simulation results that demonstrate how the technique effectively hides delays by performing other useful work while waiting for missing data to arrive.

1 Introduction

Ongoing improvements in networking technology and infrastructure have resulted in a dramatic increase in the demand for accessing and collating data from disparate, remote data sources over wide-area networks such as the Internet and intranets. Query optimization and execution strategies have long been studied in centralized, parallel, and tightly-coupled distributed environments. Data access across widely-distributed sources, however, imposes significant new challenges for query optimization and execution for two reasons: First, there are semantic and performance problems that arise due to the *heterogeneous* nature of the data sources in a loosely-coupled environment. Second, data access over wide-area networks involves a large number of remote data sources, intermediate sites, and communications links, all of which are vulnerable to congestion and failures. From the end user's point of view, congestion or failure in any of the components of the network are manifested as *highly-variable response time* — that is, the time required for obtaining data from remote sources can vary greatly depending on the specific data sources accessed and the current state of the network at the time that such access is attempted.

The query processing problems resulting from heterogeneity have been the subject of much attention in recent years (e.g., [SAD+95, BE96, TRV96]). In contrast, the impact of unpredictable response time on wide-area query processing has received relatively little attention. The work presented here is an initial exploration into addressing problems of response-time variability for wide-area data access.

1.1 Response Time Variability

High variability makes efficient query processing difficult because query execution plans are typically generated statically, based on a set of assumptions about the costs of performing various operations and the costs of obtaining data (i.e., disk and/or network accesses). The causes of high-variability are typically failures and congestion, which are inherently runtime issues; they cannot be reliably predicted at query optimization time or even at query start-up time. As a result, the execution of a statically optimized query plan is likely to be sub-optimal in the presence of unexpected response time problems. In the worst case, a query execution may be blocked for an arbitrarily long time if needed data fail to arrive from remote data sources.

The different types of response time problems that can be experienced in a loosely-coupled, wide-area environment can be categorized as follows:

- **Initial Delay** - There is an unexpected delay in the arrival of the *first* tuple from a particular remote source. This type of delay typically appears when there is difficulty connecting to a remote source, due to a failure or congestion at that source or along the path between the source and the destination.
- **Slow Delivery** - Data is arriving at a regular rate, but this rate is much slower than the

*Laurent Amsaleg is supported by a post-doctoral fellowship from INRIA Rocquencourt, France.

†Supported in part by NSF Grant IRI-94-09575, an IBM SUR award, and a grant from Bellcore.

0-8186-7475-X/96 $5.00 © 1996 IEEE

expected rate. This problem can be the result, for example, of network congestion, resource contention at the source, or because a different (slower) communication path is being used (e.g., due to a failure).

- **Bursty Arrival** - Data is arriving at an unpredictable rate, typically with bursts of data followed by long periods of no arrivals. This problem can arise from fluctuating resource demands and the lack of a global scheduling mechanism in the wide-area environment.

Because these problems can arise unpredictably at runtime, they cannot be effectively addressed by static query optimization techniques. As a result, we have been investigating a class of dynamic, runtime query plan modification techniques that we call *query plan scrambling*. In this approach, a query is initially executed according to the original plan and associated schedule generated by the query optimizer. If however, a significant performance problem arises during the execution, then query plan scrambling is invoked to modify the execution *on-the-fly*, so that progress can be made on other parts of the plan. In other words, rather than simply stalling for slowly arriving data, query plan scrambling attempts to *hide* unexpected delays by performing other useful work.

There are three ways that query plan scrambling can be used to help mask response time problems. First, scrambling allows useful work to be done in the hope that the cause of the problem is resolved in the meantime. This approach is useful for all three classes of problems described above. Second, if data are arriving, but at a rate that hampers query processing performance (e.g., in the Slow Delivery or Bursty Arrival cases), then scrambling allows useful work to be performed while the problematic data are obtained in a background fashion. Finally, in cases where data are simply not arriving, or are arriving far too slowly, then scrambling can be used to produce partial results that can then be returned to users and/or used in query processing at a later time [TRV96].

1.2 Tolerating Initial Delays

In this work, we present an initial approach to query plan scrambling that specifically addresses the problem of Initial Delay (i.e., delay in receiving the initial requested tuples from a remote data source). We describe and analyze a query plan scrambling algorithm that follows the first approach outlined above; namely, other useful work is performed in the hope that the problem will eventually be resolved, and the requested data will arrive at or near the expected rate from then on. The algorithm exploits, where possible, decisions made by the static query optimizer and imposes no optimization or execution performance overhead in the absence of unexpected delays.

In order to allow us to clearly define the algorithm and to study its performance, this work assumes an execution environment with several properties:

- The algorithm addresses only response time delays in receiving the initial requested tuples from remote data sources. Once the initial delay is over, tuples are assumed to arrive at or near the originally expected rate. As stated previously, this type of delay models problems in connecting to remote data sources, as it is often experienced in the Internet.
- We focus on query processing using a data-shipping or hybrid-shipping approach [FJK96], where data is collected from remote sources and integrated at the query source. Only query processing that is performed at the query source is subject to scrambling. This approach is typical of mediated database systems that integrate data from distributed, heterogeneous sources, e.g., [TRV96].
- Query execution is scheduled using an iterator model [Gra93]. In this model every run-time operator supports an *open()* call and a *get-next()* call. Query execution starts by calling open() on the topmost operator of the query execution plan and proceeds by iteratively calling get-next() on the topmost operator. These calls are propagated down the tree; each time an operator needs to consume data, it calls get-next() on its child (or children) operator(s). This model imposes a schedule on the operators in the query plan.

The reminder of the paper is organized as follows. Section 2 describes the algorithm and gives an extended example. Section 3 presents results from a simulation study that demonstrate the properties of the algorithm. Section 4 describes related work. Section 5 concludes with a summary of the results and a discussion of future work.

2 Scrambling Query Plans

This section describes the algorithm for scrambling queries to cope with initial delays in obtaining data from remote data sources. The algorithm consists of two phases: one that *changes the execution order of operations* in order to avoid idling, and one that *synthesizes new operations* to execute in the absence of other work to perform. We first provide a brief overview of the algorithm and then describe the two phases in detail using a running example. The algorithm is then summarized at the end of the section.

2.1 Algorithm Overview

Figure 1 shows an operator tree for a complex query plan. Typically, such a complicated plan would be generated by a static query optimizer according to its cost model, statistics, and objective functions. At the leaves of the tree are base relations stored at remote sites. The nodes of the tree are binary operators (we focus our study on hash-based joins) that are executed at the query source site.[1]

As discussed previously, we describe the scrambling algorithm in the context of an iterator-based execution

[1] Unary operators, such as selections, sorting, and partitioning are not shown in the figure.

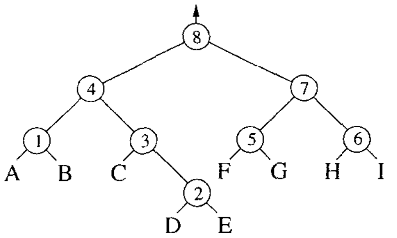

Figure 1: Initial Query Tree

model. This model imposes a schedule on the operators of a query and drives the flow of data between operators. The scheduling of operators is indicated in Figure 1 by the numbers associated to each operator. In the figure, the joins are numbered according to the order in which they would be completed by an iterator-based scheduler. The flow of data between the operators follows the model discussed in [SD90], i.e., the left input of a hash join is always materialized while the right input is consumed in a pipelined fashion.

The schedule implied by the tree in Figure 1 would thus begin by materializing the left subtree of the root node. Assuming that hash joins are used and that there is sufficient memory to hold the hash tables for relations A, C, and D (so no partitioning is necessary for these relations), this materialization would consist of the following steps:

1. Scan relation A and build hash-table H_A using selected tuples;
2. In a pipelined fashion, probe H_A with (selected) tuples of B and build a hash-table containing the result of A$\bowtie$B (H_{AB});
3. Scan C and build hash-table H_C;
4. Scan D and build hash-table H_D;
5. In a pipelined fashion, probe H_D, H_C and H_{AB} with tuples of E and build a hash-table containing the result of (A$\bowtie$B)$\bowtie$(C$\bowtie$D$\bowtie$E).
6. ...

The execution thus begins by requesting tuples from the remote site where relation A is stored. If there is a delay in accessing that site (say, because this site is temporarily down), then the scan of A (i.e., step 1) is blocked until the site recovers. Under a traditional iterator-based scheduling discipline, this delay of A would result in the entire execution of the query being blocked, pending the recovery of the remote site.

Given that unexpected delays are highly probable in a wide-area environment, such sensitivity to delays is likely to result in unacceptable performance. The scrambling algorithm addresses this problem by attempting to *hide* such delays by making progress on other parts of the query until the problem is resolved. The scrambling algorithm is invoked once a delayed relation is detected (via a timeout mechanism). The algorithm is iterative; during each iteration it selects part of the plan to execute and materializes the corresponding temporary results to be used later in the execution.

The scrambling algorithm executes in one of two phases. During *Phase 1*, each iteration modifies the schedule in order to execute operators that are not dependent on any data that is known to be delayed. For example, in the query of Figure 1, Phase 1 might result in materializing the join of relations C, D and E while waiting for the arrival of A. During *Phase 2*, each iteration synthesizes new operators (joins for example) in order to make further progress. In the example, a Phase 2 iteration might choose to join relation B with the result of (C$\bowtie$D$\bowtie$E) computed previously.

At the end of each iteration the algorithm checks to see if any delayed sources have begun to respond, and if so, it stops iterating and returns to normal scheduling of operators, possibly re-invoking scrambling if additional delayed relations are later detected. If, however, no delayed data has arrived during an iteration, then the algorithm iterates again. The algorithm moves from Phase 1 to Phase 2 when it fails to find an existing operator that is not dependent on a delayed relation. If, while in Phase 2, the algorithm is unable to create any new operators, then scrambling terminates and the query simply waits for the delayed data to arrive. In the following sections we describe, in detail, the two phases of scrambling and their interactions.

2.2 Phase 1: Materializing Subtrees

2.2.1 Blocked and Runnable Operators

The operators of a query tree have producer-consumer relationships. The immediate ancestor of a given operator consumes the tuples produced by that operator. Conversely, the immediate descendants of a given operator produce the tuples that operator consumes. The producer-consumer relationships create *execution dependencies* between operators, as one operator can not consume tuples before these tuples have been produced. For example, a select operator can not consume tuples of a base relation if that relation is not available. In such a case the select operator is blocked. If the select can not consume any tuples, it can not produce any tuples. Consequently, the consumer of the select is also blocked. By transitivity, all the ancestors of the unavailable relation are blocked.

When the system discovers that a relation is unavailable, query plan scrambling is invoked. Scrambling starts by splitting the operators of the query tree into two disjoint queues: a queue of *blocked operators* and a queue of *runnable operators*. These queues are defined as follows:

Definition 2.1 Queue of Blocked Operators: Given a query tree, the queue of blocked operators contains all the ancestors of each unavailable relation.

Definition 2.2 Queue of Runnable Operators: Given a query tree and a queue of blocked operators, the queue of runnable operators contains all the operators that are not in the queue of blocked operators.

Operators are inserted in the runnable and blocked queues according to the order in which their execution would be initiated by an iterator-based scheduler.

2.2.2 Maximal Runnable Subtree

Each iteration during Phase 1 of query plan scrambling analyzes the runnable queue in order to find a *maximal runnable subtree* to materialize. A maximal runnable subtree is defined as follows:

Definition 2.3 Maximal Runnable Subtree: Given the query tree and the queues of blocked and runnable operators, a runnable subtree is a subtree in which all the operators are runnable. A runnable subtree is maximal if its root is the first runnable descendant of a blocked operator.

None of the operators belonging to a maximal runnable subtree depend on data that is known to be delayed. Each iteration of Phase 1 initiates the materialization of the first maximal runnable subtree found. The notion of maximal used in the definition is important, as materializing the biggest subtrees during each iteration tends to minimize the number of materializations performed, hence reducing the amount of extra I/O caused by scrambling. The materialization of a runnable subtree completes only if no relations used by this subtree are discovered to be unavailable during the execution.[2] When the execution of a runnable subtree is finished and its result materialized, the algorithm removes all the operators belonging to that subtree from the runnable queue. It then checks if missing data have begun to arrive. If the missing data from others, blocked relations are still unavailable, another iteration is begun. The new iteration analyzes (again) the runnable queue to find the next maximal runnable subtree to materialize.

2.2.3 Subtrees and Data Unavailability

It is possible that during the execution of a runnable subtree, one (or more) of the participating base relations is discovered to be unavailable. This is because a maximal runnable subtree is defined with respect to the *current* contents of the blocked and runnable queues. The runnable queue is only a guess about the real availability of relations. When the algorithm inserts operators in the runnable queue, it does not know whether their associated relations are actually available or unavailable. This will be discovered only when the corresponding relations are requested.

In the case where a relation is discovered to be unavailable during the execution of a runnable subtree, the current iteration stops and the algorithm updates the runnable and blocked queues. All the ancestors of the unavailable relation are extracted from the runnable queue and inserted in the blocked queue. Once the queues are updated, the scrambling of the query plan initiates a new Phase 1 iteration in order to materialize another maximal runnable subtree.

[2]Note that in the remainder of this paper, we use "maximal runnable subtree" and "runnable subtree" interchangeably, except where explicitly noted.

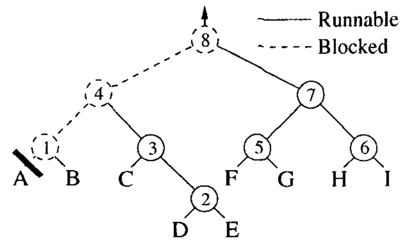

Figure 2: Blocked and Runnable Operators with Relation A Unavailable

2.2.4 Termination of Phase 1

At the end of each iteration, the algorithm checks for data arrival. If it is discovered that an unavailable relation has begun to arrive, the algorithm updates the blocked and runnable queues. The ancestors of the unblocked relation are extracted from the blocked queue and inserted in the runnable queue. Note that any ancestors of the unblocked relation that also depend on other blocked relations are not extracted from the queue. Phase 1 then terminates and the execution of the query returns to normal iterator-based scheduling of operators. If no further relations are blocked, the execution of the query will proceed until the final result is returned to the user. The scrambling algorithm will be re-invoked, however, if the query execution blocks again.

Phase 1 also terminates if the runnable queue is empty. In this case, Phase 1 can not perform any other iteration because all remaining operators are blocked. When this happens, query plan scrambling switches to Phase 2. The purpose of the second phase is to process the available relations when all the operators of the query tree are blocked. We present the second phase of query plan scrambling in Section 2.3. First, however, we present an example that illustrates all the facets of Phase 1 described above.

2.2.5 A Running Example

This example reuses the complex query tree presented at the beginning of Section 2. To discuss cases where data need or do need not to be partitioned before being joined, we assume that tuples of relations A, B, C, D and E do not need to be partitioned. In contrast, we assume that the tuples of relations F, G, H and I have to be partitioned. To illustrate the behavior of Phase 1, we follow the scenario given below:

1. When the execution of the query starts, relation A is discovered to be unavailable.
2. During the third iteration, relation G is discovered to be unavailable.
3. The tuples of A begin to arrive at the query execution site before the end of the fourth iteration.
4. At the time Phase 1 terminates, no tuples of G have been received.

The execution of the example query begins by requesting tuples from the remote site owning relation A. Following the above scenario, we assume relation A is

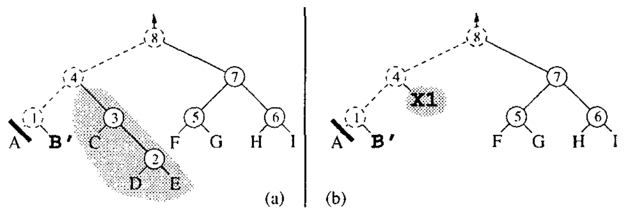

Figure 3: Query Tree During Iterations 1 and 2

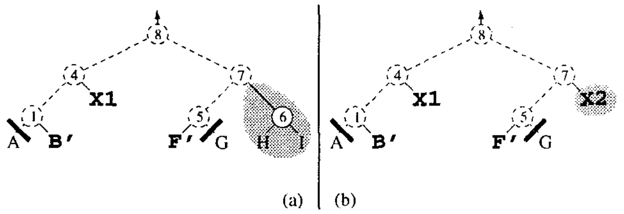

Figure 4: G Unavailable; X2 Materialized

unavailable (indicated by the thick solid line in Figure 2). The operators that are blocked by the delay of A are depicted using a dashed line.

The unavailability of A invokes Phase 1 which updates the blocked and runnable queues and initiates its first iteration. This iteration analyses the runnable queue and finds that the first maximal runnable subtree consists of a unary operator that selects tuples from relation B.[3] Once the operator is materialized (i.e., selected tuples of B are on the local disk stored in the relation **B'**), the algorithm checks for the arrival of the tuples of A. Following the above scenario, we assume that the tuples of A are still unavailable, so another iteration is initiated. This second iteration finds the next maximal runnable subtree to be the one rooted at operator 3. Note the subtree rooted at operator 2 is *not* maximal since its consumer (operator 3) is not blocked.

Figure 3 shows the materialization of the runnable subtrees found by the first two iterations of query scrambling. Part (a) of this figure shows the effect of materializing of the first runnable subtree: the local relation **B'** contains the materialized and selected tuples of the remote relation B. It also shows the second runnable subtree (indicated by the shaded grey area). Figure 3(b) shows the query tree after the materialization of this second runnable subtree. The materialized result is called **X1**.

Once **X1** is materialized, another iteration starts since, in this example, relation A is still unavailable. The third iteration finds the next runnable subtree rooted at operator 7 which joins F, G, H and I (as stated above, these relations need to be partitioned before being joined). The execution of this runnable subtree starts by building the left input of operator 5 (partitioning F into **F'**). It then requests relation G in order to partition it before probing the tuples of F. In this scenario, however, G is discovered to be unavailable, triggering the update of the blocked and runnable queues. Figure 4(a) shows that operators 5 and 7 are newly blocked operators (operator 8 was already blocked due to the unavailability of A). Once the queues of operators are updated, another iteration of scrambling is initiated to run the next runnable subtree, i.e., the one rooted at operator 6 (indicated by the shaded grey area in the figure). The result of this execution is called **X2**.

[3] As stated earlier, operators are inserted into the queues with respect to their execution order.

Figure 5 illustrates the next step in the scenario, i.e., it illustrates the case where after **X2** is materialized it is discovered that the tuples of relation A have begun to arrive. In this case, the algorithm updates the runnable and blocked queues. As shown in Figure 5(a), operators 1 and 4 that were previously blocked are now unblocked (operator 8 remains blocked however). Phase 1 then terminates and returns to the normal iterator-based scheduling of operators which materializes the left subtree of the root node (see Figure 5(b)). The resulting relation is called **X3**.

After **X3** is materialized, the query is blocked on G so Phase 1 is re-invoked. Phase 1 computes the new contents of the runnable and blocked queue and discovers that the runnable queue is empty since all remaining operators are ancestors of G. Phase 1 then terminates and the scrambling of the query plan enters Phase 2. We describe Phase 2 of the algorithm in the next section.

2.3 Phase 2: Creating New Joins

Scrambling moves into Phase 2 when the runnable queue is empty but the blocked queue is not. The goal of Phase 2 is to *create* new operators to be executed. Specifically, the second phase creates joins between relations that were not directly joined in the original query tree, but whose consumers are blocked (i.e., in the blocked queue) due to the unavailability of some other data.

In contrast to Phase 1 iterations, which simply adjust scheduling to allow *runnable* operators to execute, iterations during Phase 2 actually create new joins. Because the operations that are created during Phase 2 were not chosen by the optimizer when the original query plan was generated, it is possible that these operations may entail a significant amount of additional work. If the joins created and executed by Phase 2 are too expensive, query scrambling could result in a net degradation in performance. Phase 2,

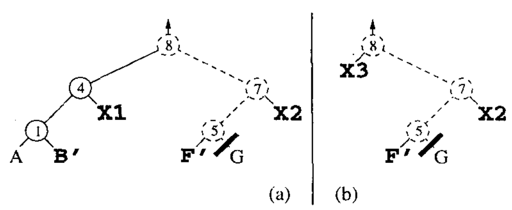

Figure 5: Relation A Available

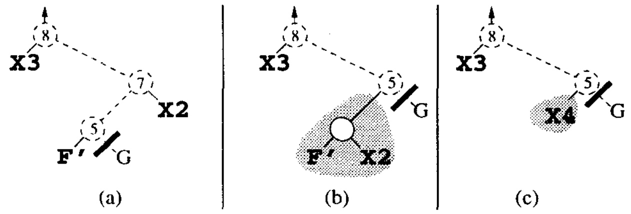

Figure 6: Performing a New Join in Phase 2

therefore, has the potential to negate or even reverse the benefits of scrambling if care is not taken. In this paper we use the simple heuristic of avoiding Cartesian products to prevent the creation of overly expensive joins during Phase 2. In Section 3, we analyze the performance impact of the cost of created joins relative to the cost of the joins in the original query plan. One way to ensure that Phase 2 does not generate overly expensive joins is to involve the query optimizer in the choice of new joins. Involving the optimizer in query scrambling is one aspect of our ongoing work.

2.3.1 Creating New Joins

At the start of Phase 2, the scrambling algorithm constructs a graph $\mathcal{G}$ of possible joins. Each node in $\mathcal{G}$ corresponds to a relation, and each edge in $\mathcal{G}$ indicates that the two connected nodes have common join attributes, and thus can be joined without causing a Cartesian product. Unavailable relations are not placed into $\mathcal{G}$.

Once $\mathcal{G}$ is constructed, Phase 2 starts to iteratively create and execute new join operators. Each iteration of Phase 2 performs the following steps:

1. In $\mathcal{G}$, find the two *leftmost* joinable (i.e., connected) relations i and j. The notion of *leftmost* is with respect to the order in the query plan. If there are no joinable relations in $\mathcal{G}$, then terminate scrambling.
2. Create a new join operator $i \bowtie j$.
3. Materialize $i \bowtie j$. Update $\mathcal{G}$ by replacing i and j with the materialized result of $i \bowtie j$. Update runnable and blocked queues. Update query tree.
4. Test to see if any unavailable data has arrived. If so, then terminate scrambling, else begin a new iteration.

Figure 6 demonstrates the behavior of Phase 2 by continuing the example of the previous section. The figure is divided into three parts. Part (a) shows the query tree at the end of Phase 1. In this case, $\mathcal{G}$ would contain **F'**, **X2**, and **X3**. Assume that, in $\mathcal{G}$, relations **F'** and **X2** are directly connected but relation **X3** is not connected to either (i.e., assume it shares join attributes only with the unavailable relation G). In this example, therefore, **F'** and **X2** are the two leftmost *joinable* relations; **X3** is the leftmost relation, but it is not joinable.

Figure 6(b) shows the creation of the new join of **F'** and **X2**. The creation of this join requires the removal of join number 7 from the blocked queue and its replacement in the ordering of execution by join number 5. Finally, Figure 6(c) shows the materialization of the created operator. The materialized join is called **X4**. At this point, $\mathcal{G}$ is modified by removing **F'** and **X2** and inserting **X4**, which is not joinable with **X3**, the only other relation in $\mathcal{G}$.

2.3.2 Termination of Phase 2

After each iteration of Phase 2, the number of relations in $\mathcal{G}$ is reduced. Phase 2 terminates if $\mathcal{G}$ is reduced to a single relation, or if there are multiple relations but none that are joinable. As shown in the preceding example, this latter situation can arise if the attribute(s) required to join the remaining relations are contained in an unavailable relation (in this case, relation G).

Phase 2 can also terminate due to the arrival of unavailable data. If such data arrive during a Phase 2 iteration, then, at the end of that iteration, the runnable and blocked queues are updated accordingly and the control is returned to the normal iterator-based scheduling of operators. As mentioned for Phase 1, query scrambling may be re-invoked later to cope with other delayed relations.

2.3.3 Physical Properties of Joins

The preceding discussion focused on restructuring *logical* nodes of a query plan. The restructuring of *physical* plans, however, raises additional considerations. First, adding a new join may require the introduction of additional unary operators to process the inputs of this new join so that it can be correctly executed. For example, a merge join operator requires that the tuples it consumes are sorted, and thus may require that sort operators be applied to its inputs. Second, deleting operators, as was done in the preceding example, may also require the addition of unary operators. For example, relations may need to be repartitioned in order to be placed as children of an existing hybrid hash node. Finally, changing the inputs of an existing join operator may also require modifications. If the new inputs are sufficiently different than the original inputs, the physical join operators may have to be modified. For example, an indexed nested loop join might have to be changed to a hash join if the inner relation is replaced by one that is not indexed on the join attribute.

2.4 Summary and Discussion

The query plan scrambling algorithm can be summarized as follows:

- When a query becomes blocked (because relations are unavailable), query plan scrambling is initiated. It first computes a queue of blocked operators and a queue of runnable operators.
- Phase 1 then analyses the queue of runnable operators, picks a maximal runnable subtree and materializes its result. This process is repeated, i.e., it iterates, until the queue of runnable operators is empty. At this point, the system switches to Phase 2.

- Phase 2 tries to create a new operator that joins two relations that are available and joinable. This process iterates until no more joinable relations can be found.
- After each iteration of the algorithm, it checks to see if any unavailable data have arrived, and if so, control is returned to normal iterator-based scheduling of operators, otherwise another iteration is performed.

There are two additional issues regarding the algorithm that deserve mention, here. The first issue concerns the knowledge of the actual availability of relations. Instead of discovering, as the algorithm does now, during the execution of the operations performed by each iteration that some sources are unavailable, it is possible to send some or all of the initial data requests to the data sources as soon as the first relation is discovered to be unavailable. Doing so would give the algorithm immediate knowledge of the availability status of all the sources. Fortunately, using the iterator model, opening multiple data sources at once does not force the query execution site to consume all the tuples simultaneously — the iterator model will suspend the flow of tuples until they are consumed by their consumer operators.

The second issue concerns the potential additional work of each phase. As described previously, Phase 1 materializes existing subtrees that have been optimized prior to runtime by the query optimizer. The relative overhead of each materialization may be more or less significant depending on the I/O pattern of the scrambled subtree compared to its unscrambled version. For example, if a subtree consists of a single select on a base relation, its materialization during Phase 1 is pure overhead since the original query plan was selecting tuples as they were received, without involving any I/O. On the other hand, the overhead of materializing an operator that partitions data is comparatively less important. In this case, both the original query plan and the scrambled plan have to perform disk I/Os to write the partitions on disk for later processing. The scrambled plan, however, writes to disk one extra partition that would be kept in memory by the original non-scrambled query plan.

Phase 2, however, can be more costly as it creates new joins from scratch using the simple heuristic of avoiding Cartesian products. The advantage of this approach is its simplicity. The disadvantage, however, is the potential overhead caused by the possibly sub-optimal joins. We study the performance impact of varying costs of the created joins in the following section.

The costs of materializations during Phase 1 and of new joins during Phase 2 may, in certain cases, negate the benefits of scrambling. Controlling these costs raises the possibility of integrating scrambling with an existing query optimizer. This would allow us to estimate the costs of iterations in order to skip, for example, costly materializations or expensive joins. Such an integration is one aspect of our ongoing work.

Parameter	Value	Description
NumSites	8	number of sites
Mips	30	CPU speed (10^6 instr/sec)
NumDisks	1	number of disks per site
DskPageSize	4096	size of a disk page (bytes)
NetBw	1	network bandwidth (Mbit/sec)
NetPageSize	8192	size of a network page (bytes)
Compare	4	instr. to apply a predicate
HashInst	25	instr. to hash a tuple
Move	2	instr. to copy 4 bytes

Table 1: Simulation Parameters and Main Settings

3 Performance

In this section, we examine the main performance characteristics of the query scrambling algorithm. The first set of experiments shows the typical performance of any query that is scrambled. The second set of experiments studies the sensitivity of Phase 2 to the selectivity of the new joins it creates. We first describe the simulation environment used to study the algorithm.

3.1 Simulation Environment

To study the performance of the query scrambling algorithm, we extended an existing simulator [FJK96, DFJ+96] that models a heterogeneous, peer-to-peer database system such as SHORE [CDF+94]. The simulator we used provides a detailed model of query processing costs in such a system. Here, we briefly describe the simulator, focusing on the aspects that are pertinent to our experiments.

Table 1 shows the main parameters for configuring the simulator, and the settings used for this study. Every site has a CPU whose speed is specified by the *Mips* parameter, *NumDisks* disks, and a main-memory buffer pool. For the current study, the simulator was configured to model a client-server system consisting of a single client connected to seven servers. Each site, except the query execution site, stores one base relation.

In this study, the disk at the query execution site (i.e., client) is used to store temporary results. The disk model includes costs for random and sequential physical accesses and also charges for software operations implementing I/Os. The unit of disk I/O for the database and the client's disk cache are pages of size *DskPageSize*. The unit of transfer between sites are pages of size *NetPageSize*. The network is modeled simply as a FIFO queue with a specified bandwidth (*NetBw*); the details of a particular technology (Ethernet, ATM) are not modeled. The simulator also charges CPU instructions for networking protocol operations. The CPU is modeled as a FIFO queue and the simulator charges for all the functions performed by query operators like hashing, comparing, and moving tuples in memory.

In this paper, the simulator is used primarily to demonstrate the properties of the scrambling algorithm, rather than for a detailed analysis of the algo-

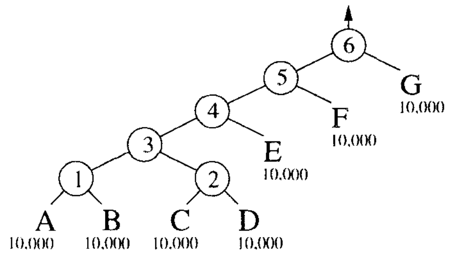

Figure 7: Query Tree Used for the Experiments

rithm. As such, the specific settings used in the simulator are less important than the way in which delay is either hidden or not hidden by the algorithm. In the experiments, the various delays were generated by simply requesting tuples from an "unavailable" source at the end of the various iterations of query plan scrambling. That is, rather than stochastically generating delays, we explicitly imposed a series of delays in order to study the behavior of the algorithm in a controlled manner. For example, to simulate the arrival of blocked tuples during, say, the third iteration of Phase 1, we scrambled the query 3 times, and then initiated the transfer of tuples from the "blocked" relation so that the final result of the query could eventually be computed.

3.2 A Query Tree for the Experiments

For all the experiments described in this section, we use the query tree represented in Figure 7. We use this query tree because it demonstrates all of the features of scrambling and allows us to highlight the impact on performance of the overheads caused by materializations and created joins.

Each base relation has 10,000 tuples of 100 bytes each. We assume that the join graph is fully connected, that is, any relation can be (equi-)joined with any other relation and that all joins use the same join attribute. In the first set of experiments, we study the performance of query plan scrambling in the case where all the joins in the query tree produce the same number of tuples, i.e., 1,000 tuples. In the second set of experiments, however, we study the case where the joins in the query tree have different selectivities and thus produce results of various sizes.

For all the experiments, we study the performance of our approach in the case where a single relation is unavailable. This relation is the *left-most* relation (i.e., relation A) which represents the case where query scrambling is the most beneficial. Examining the cases with others unavailable relations would not change the basic lessons of this study.

For each experiment described below, we evaluate the algorithm in the cases where it executes in the context of a small or a large memory. In the case of large memory, none of the relations used in the query tree (either a base relation or an intermediate result) need to be partitioned before being processed. In the case of small memory, every relation (including intermediate results) must be partitioned. Note, that since all joins in the test query use the same join attribute, no re-partitioning of relations is required when new joins are created in this case.

3.3 Experiment 1: The Step Phenomenon

Figure 8 shows the response time for the scrambled query plans that are generated as the delay for relation A (the leftmost relation in the plan) is varied. The delay for A is shown along the X-axis, and is also represented as the lower grey line in the figure. The higher grey line shows the performance of the unscrambled query, that is, if the execution of the query is simply delayed until the tuples of relation A begin to arrive. The distance between these two lines therefore is constant, and is equal to the response time for the original (unscrambled) query plan, which is 80.03 seconds in this case. In this experiment, the memory size of the query execution site is small. With this setting, the hash-tables for inner relations for joins can not entirely be built in memory so partitioning is required.

The middle line in Figure 8 shows the response time for the scrambled query plans that are executed for various delays of A. In this case, there are six possible scrambled plans that could be generated. As stated in Sections 2.2 and 2.3, the scrambling algorithm is iterative. At the *end of each iteration* it checks to see if delayed data has begun to arrive, and if so, it stops scrambling and normal query execution is resumed. If, however, at the end of the iteration, the delayed data has still not arrived, another iteration of the scrambling algorithm is initiated. The result of this execution model is the step shape that can be observed in Figure 8.

The *width* of each step is equal to the duration of the operations that are performed by the current iteration of the scrambling algorithm, and the *height* of the step is equal to the response time of the query if normal processing is resumed at the end of that iteration. For example, in this experiment, the first scrambling iteration results in the retrieval and partitioning of relation B. This operation requires 12.23 seconds. If at the end of the iteration, tuples of relation A have begun to arrive then no further scrambling is done and normal query execution resumes. The resulting execution in this case, has a response time of 80.10 seconds. Thus, the first step shown in Figure 8 has a width of 12.23 seconds and a height of 80.10 seconds. Note that in this case, scrambling is effective at hiding the delay of A; the response time of the scrambled query is nearly identical to that of original query with no delay of A.

If no tuples of A have arrived at the end of the first iteration, then another iteration is performed. In this case, the second iteration retrieves, partitions, and joins relations C and D. As shown in Figure 8, this iteration requires an additional 26.38 seconds, and if A begins to arrive during this iteration, then the resulting query plan has a total response time of 80.90 seconds. Thus, in this experiment, scrambling is able to hide delays of up to 38.61 seconds with a penalty of no more than 0.80 seconds (i.e., 1%) of the response

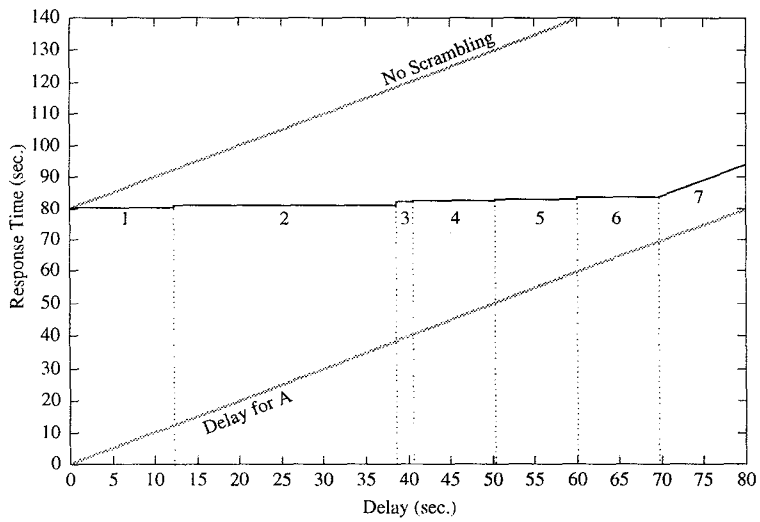

Figure 8: Response Times of Scrambled Query Plans (Small Memory, Varying the Delay of A.)

time of the original query with no delay. This corresponds to a response time improvement of up to 32% compared to not scrambling.

If, at the end of the second iteration, tuples of A have still failed to arrive, then the third iteration is initiated. In this case however, there are no more runnable subtrees, so scrambling switches to Phase 2, which results in the creation of new joins (see Section 2.3). In this third iteration, the result of C⋈D is partitioned and joined with relation B. This iteration has a width of only 2.01 seconds, because both inputs are already present, B is already partitioned, and the result of C⋈D is fairly small. The response time of the resulting plan is 82.22 seconds, which again represents a response time improvement of up to 32% compared to not scrambling.

The remaining query plans exhibit similar behavior. Table 2 shows the additional operations and the overall performance for each of the possible scrambled plans. In this experiment, the largest relative benefit (approximately 44%) over not scrambling is obtained when the delay of A is 69.79 seconds, which is the time required to complete all six iterations. After this point, there is no further work for query scrambling to do, so the scrambled plan must also wait for A to arrive. As can be seen in Figure 8, at the end of iteration six the response time of the scrambled plan increases linearly with the delay of A. The distance between the delay of A and the response time of the scrambled plan is the time that is required to complete the query once A arrives.

Although it is not apparent in Figure 8, the first scrambled query is slightly slower than the unscrambled query plan when A is delayed for a very short amount of time. For a delay below 0.07 seconds, the response time of the scrambled query is 80.10 seconds while it is 80.03 seconds for the non-scrambled query. When joining A and B, as the unscrambled query does, B is partitioned *during* the join, allowing one of the partitions of B to stay in memory. Partitioning B *before* joining it with A, as the first scrambled query plan does, forces this partition to be written back to disk and to be read later during the join with A. When A is delayed by less than the time needed to perform these additional I/Os, it is cheaper to stay idle waiting for A.

3.4 Experiment 2: Sensitivity of Phase 2

In the previous experiment all the joins produced the same number of tuples, and as a result, all of the operations performed in Phase 2 were beneficial. In this section, we examine the sensitivity of Phase 2 to changes in the selectivities of the joins it creates. Varying selectivities changes the number of tuples produced by these joins which affects the width and the height of each step. Our goal is to show cases where the benefits of scrambling vary greatly, from clear improvements to cases where scrambling performs worse than no-scrambling.

For the test query, the first join created in Phase 2 is the join of relation B with the result of C⋈D (which was materialized during Phase 1). In this set of experiments, we vary the selectivity of this new join to create a result of a variable size. The selectivity of this join is adjusted such that it produces from 1,000 tuples up to several thousand tuples. The other joins that Phase 2 may create behave like functional joins and they simply carry all the tuples created by (B⋈(C⋈D)) through the query tree. At the time these tuples are joined with A, the number of tuples carried along the query tree returns to normality and drops down to 1,000. Varying the selectivity of the first join produced by Phase 2 is sufficient to generate a variable number of tuples that are carried all along the tree by the other joins that Phase 2 may create.

The two next sections present the results of this sensitivity analysis for a small and a large memory case.

Scrambled Plan #	Performed by Iteration	Total Delay	Response Time	Savings
1	Partition B	0–12.23	80.10	up to 13.18%
2	X1←C⋈D	12.23–38.61	80.90	12.31–31.81%
3	X2←B⋈X1	38.61–40.62	82.22	30.69–31.85%
4	X3←X2⋈E	40.62–50.32	82.51	31.61–36.70%
5	X4←X3⋈F	50.32–60.07	83.05	36.28–40.72%
6	X5←X4⋈G	60.07–69.79	83.52	40.38–44.21%

Table 2: Delay Ranges and Response Times of Scrambled Query Plans

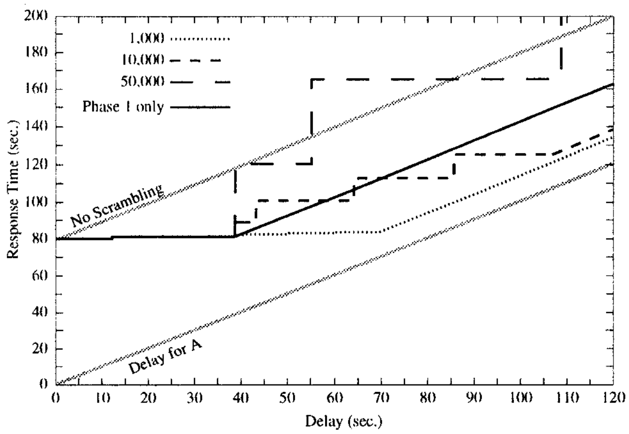

Figure 9: Response Times of Scrambled Query Plans (Small Memory, Varying Selectivity and Delay.)

As stated previously, when the memory is small, relations have to be partitioned before being joined (as in the previous experiment). This partitioning adds to the potential cost of scrambled plans because it results in additional I/O that would not have been present in the unscrambled plan. When the memory is large, however, hash-tables can be built entirely in memory so relations do not need to be partitioned. Thus, with large memory the potential overhead of scrambled plans is lessened.

3.4.1 Small Memory Case

In this experiment, we examine the effectiveness of query scrambling when the selectivity of the first join created by Phase 2 is varied. Figure 9 shows the response time results for 3 different selectivities. As in the previous experiment, the delay for A is shown along the X-axis and is also represented as the lower grey line in the figure. The higher grey line shows the response time of the unscrambled query, which as before, increases linearly with the delay of A. These two lines are exactly the same as the ones presented in the previous experiment.

The solid line in the middle of the figure shows the performance of a scrambled query plan that stops scrambling right at the end of Phase 1 (in this case, two iterations are performed during Phase 1) without initiating any Phase 2 iterations. Note that this line becomes diagonal after the end of Phase 1 since the system simply waits until the tuples of A arrive before computing the final result of the query.

Intuitively, it is not useful to perform a second phase for scrambled queries when the resulting response time would be located above this line. Costly joins that would be created by Phase 2 would consume a lot of resources for little improvement. On the other hand, Phase 2 would be beneficial for scrambled queries whose resulting response time would be below this line since the additional overhead would be small and the gain large.

The dashed and dotted lines in the figure illustrate the tradeoffs. These lines show the response time for the scrambled query plans that are executed for various delays of A and for various selectivities. Note all these scrambled query plans share the same response times for the iterations performed during Phase 1. These two first iterations correspond exactly to the scrambled plans 1 and 2 described in the previous experiment. At the end of the second iteration (38.61 seconds), however, if the tuples of A have still failed to arrive, a third iteration is initiated and the query scrambling enters Phase 2 which creates new joins.

The dotted line shows the performance when the selectivity for the new join is such that it produces a result of 1,000 tuples. This line is identical to the one showed in the previous experiment since all the joins were producing 1,000 tuples.

With the second selectivity, the first join created by the second phase produces 10,000 tuples. If at the end of this iteration, the tuples of A have still not arrived, another iteration is initiated and this iteration has to process and to produce 10,000 tuples. The corresponding line in the figure is the lowest dashed line. In this case, where 10 times more tuples have to be carried along the scrambled query plans, each step is higher (roughly 12 seconds) and wider since more tuples have to be manipulated than in the case where only 1,000 tuples are created. Even with the additional overhead of these 10,000 tuples, however, the response times of the scrambled query plans are far below the response times of the unscrambled query with equivalent delay.

When the new join produces 50,000 tuples (the higher dashed line in the figure), the response time of the scrambled plans are almost equal to or even worse than that of the original unscrambled query *including the delay for A*. In this case, it is more costly to carry the large number of tuples through the query tree than to simply wait for blocked data to arrive.

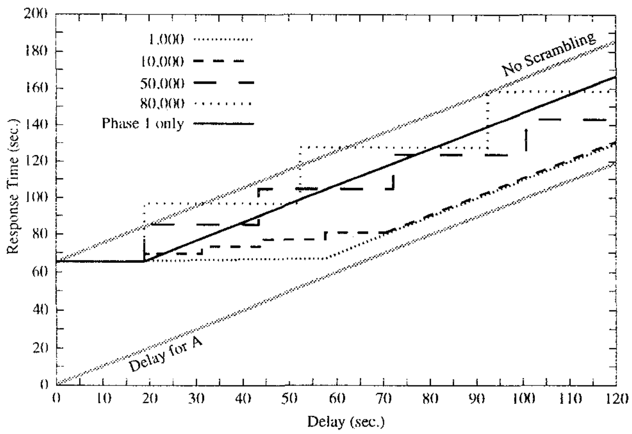

Figure 10: Response Times of Scrambled Query Plans (Large Memory, Varying Selectivity and Delay.)

3.4.2 Large Memory Case

Figure 10 shows the same experiment in the case where the memory is large enough to allow inner relations for joins to be built entirely in main memory. With large memory, no partitioning of relations needs to be done.

For the large memory case, the lines showing the increasing delay of A and the response time of the unscrambled query when this delay increases are separated by 65.03 seconds and Phase 2 starts when A is delayed by more than 18.95 seconds. Four different selectivities are represented in this figure.

In contrast to the previous experiment where 50 times more tuples negated the benefits of scrambling, in this case up to 80 times more tuples can be carried by the scrambled query plans before the benefits become close to zero. With a large memory, results computed by each iteration need only be materialized and can be consumed as is. In contrast, when the memory is small, materialized results have to be partitioned before being consumed. With respect to a small memory case, not partitioning the relation when the memory is large reduces the number of I/Os and allows the scrambled plans to manipulate more tuples for the same overhead.

3.5 Discussion

The experiments presented in this section have shown that query scrambling can be an effective technique that is able to improve the response time of queries when data are delayed. These improvements come from the fact that each iteration of a scrambled query plan can hide the delay of data. The improvement, however, depends on the overhead due to materializations and created joins.

The improvement that scrambling can bring also depends on the amount of work done in the original query. The bigger (i.e., the longer and the more costly) the original query is, the more improvement our technique can bring since it will be able to hide larger delays by computing costly operations. The improvement also depends on the shape of the query tree: bushy trees offer more options for scrambling than deep trees.

With respect to the Figures 9 and 10 presented above, when many iterations can be done during Phase 1, the point where Phase 2 starts shifts to the right. This increases the distance between the Phase 1-only diagonal line and the response time of the unscrambled query. In turn, the scrambling algorithm can handle a wider range of bad selectivities for the joins it creates during Phase 2.

4 Related Work

In this section we consider related work with respect to (a) the point in time that optimization decisions are made (i.e., compile time, query start-up time, or query run-time); (b) the variables used for dynamic decisions (i.e. if the response time of a remote source is considered); (c) the nature of the dynamic optimization (i.e. if the entire query can be rewritten); and (d) the basis of the optimization (i.e., cost-based or heuristic based).

The Volcano optimizer [CG94, Gra93] does dynamic optimization for distributed query processing. During optimization, if a cost comparison returns *incomparable*, the choice for that part of the search space is encoded in a *choose-plan* operator. At query start up time, all the incomparable cost comparisons are reevaluated. According to the result of the reevaluation, the choose-plan operator selects a particular query execution plan. All final decisions regarding query execution are thus made at query start-up time. Our work is complimentary to the Volcano optimizer since Volcano does not adapt to changes once the evaluation of the query has started.

Other work in dynamic query optimization either does not consider the distributed case [DMP93, OHMS92] or only optimizes access path selection and cannot reorder joins [HS93]. Thus, direct considerations of problems with response times from remote sources are not accounted for. These articles are, however, a rich source of optimizations which can be carried over into our work.

A novel approach to dynamic query optimization used in Rdb/VMS is described in [Ant93]. In this approach, multiple different executions of the same logical operator occur at the same time. They compete for producing the best execution – when one execution of an operator is determined to be (probably) better, the other execution is terminated.

In [DSD95] the response time of queries is improved by reordering left-deep join trees into bushy join trees. Several reordering algorithms are presented. This work assumes that reordering is done entirely at compile time. This work cannot easily be extended to handle run-time reordering, since the reorderings are restricted to occur at certain locations in the join tree.

[ACPS96] tracks the costs of previous calls to remote sources (in addition to caching the results) and can use this tracking to estimate the cost of new calls. As in Volcano, this system optimizes a query both at query compile and query start-up time, but does not change the query plan during query run-time.

The research prototype Mermaid [CBTY89] and its commercial successor InterViso [THMB95] are heterogeneous distributed databases that perform dynamic query optimization. Mermaid constructs its query plan entirely at run-time, thus each step in query optimization is based on dynamic information such as intermediate join result sizes and network performance. Mermaid neither takes advantage of a statically generated plan nor does it dynamically account for a source which does not respond at run-time.

The Sage system [Kno95] is an AI planning system for query optimization for heterogeneous distributed sources. This system interleaves execution and optimization and responds to unavailable data sources.

5 Conclusion and Future Work

Query plan scrambling is a novel technique that can dynamically adjust to changes in the run-time environment. We presented an algorithm which specifically deals with variability in performance of remote data sources and accounts for *initial* delays in their response times. The algorithm consists of two phases. Phase 1 changes the scheduling of existing operators produced as a result of query optimization. Phase 1 is iteratively applied until no more changes in the scheduling are possible. At this point, the algorithm enters Phase 2 which creates new operators to further process available data. New operators are iteratively created until there is no further work for query plan scrambling to do.

The performance experiments demonstrated how the technique hides delays in receiving the initial requested tuples from remote data sources. We then examined the sensitivity of the performance of scrambled plans to the selectivity of the joins created in Phase 2.

This work represents an initial exploration into the development of flexible systems that dynamically adapt to the changing properties of the environment. Among our ongoing and future research plans, we are developing algorithms that can scramble under different failure models to handle environments where data arrives at a bursty rate or at a steady rate that is significantly slower than expected. We are also studying the use of partial results which approximate the final results. We also plan to study the potential improvement of basing scrambling decisions on cost-based knowledge.

Finally, query plan scrambling is a promising approach to addressing many of the concerns addressed by dynamic query optimization. Adapting the query plan at run-time to account for the actual costs of operations could compensate for the often inaccurate and unreliable estimates used by the query optimizer. Moreover, it could account for remote sources that do not export any cost information, which is especially important when these remote sources run complex subqueries. Thus, we plan to investigate the use of scrambling as a complimentary approach to dynamic query optimization.

Acknowledgments We would like to thank Praveen Seshadri, Björn Jónsson and Jean-Robert Gruser for their helpful comments on this work. We would also like to thank Alon Levy for pointing out related work.

References

[ACPS96] S. Adali, K. Candan, Y. Papakonstantinou, and V. Subrahmanian. Query Caching and Optimization in Distributed Mediator Systems. *ACM SIGMOD Conf.*, Montréal, Canada, 1996.

[Ant93] G. Antoshenkov. Dynamic Query Optimization in Rdb/VMS. *ICDE Conf.*, Vienna, Austria, 1993.

[BE96] O. Bukhres and A. Elmagarmid. *Object-Oriented Multidatabase Systems.* Prentice Hall, 1996.

[CBTY89] A. Chen, D. Brill, M. Templeton, and C. Yu. Distributed Query Processing in a Multiple Database System. *IEEE Journal on Selected Areas in Communications*, 7(3), 1989.

[CDF+94] M. Carey, D. DeWitt, M. Franklin, N. Hall, M. McAuliffe, J. Naughton, D. Schuh, M. Solomon, C. Tan, O. Tsatalos, S. White, and M. Zwilling. Shoring Up Persistent Applications. *ACM SIGMOD Conf.*, Minneapolis, MN, 1994.

[CG94] R. Cole and G. Graefe. Optimization of Dynamic Query Execution Plans. *ACM SIGMOD Conf.*, Minneapolis, MN, 1994.

[DFJ+96] S. Dar, M. Franklin, B. Jónsson, D. Srivastava, and M. Tan. Semantic Data Caching and Replacement. *22nd VLDB Conf.*, Bombay, India, 1996.

[DMP93] M. Derr, S. Morishita, and G. Phipps. Design and Implementation of the Glue-Nail Database System. *ACM SIGMOD Conf.*, Washington, DC, 1993.

[DSD95] W. Du, M. Shan, and U. Dayal. Reducing Multidatabase Query Response Time by Tree Balancing. *ACM SIGMOD Conf.*, San Jose, CA, 1995.

[FJK96] M. Franklin, B. Jónsson, and D. Kossmann. Performance Tradeoffs for Client-Server Query Processing. *ACM SIGMOD Conf.*, Montréal, Canada, 1996.

[Gra93] G. Graefe. Query Evaluation Techniques for Large Databases. *ACM Computing Surveys*, 25(2), 1993.

[HS93] W. Hong and M. Stonebraker. Optimization of Parallel Query Execution Plans in XPRS. *Distributed and Parallel Databases*, 1(1), 1993.

[Kno95] C. Knoblock. Planning, Executing, Sensing, and Replanning for Information Gathering. In *Proc. of the 14th Int. Joint Conf. on Artificial Intelligence*, Montréal, Canada, 1995.

[OHMS92] J. Orenstein, S. Haradhvala, B. Margulies, and D. Sakahara. Query Processing in the ObjectStore Database System. *ACM SIGMOD Conf.*, San Diego, CA, 1992.

[SD90] D. Schneider and D. DeWitt. Tradeoffs in Processing Complex Join Queries via Hashing in Multiprocessor Database Machines. *16th VLDB Conf.*, Brisbane, Australia, 1990.

[SAD+95] M. Shan, R. Ahmen, J. Davis, W. Du, and W. Kent. *Modern Database Systems: The Object Model, Interoperability, and Beyond*, chapter Pegasus: A Heterogeneous Information Management System. ACM Press, 1995.

[THMB95] M. Templeton, H. Henley, E. Maros, and D. Van Buer. InterViso: Dealing with the Complexity of Federated Database Access. *VLDB Journal*, 4(2), 1995.

[TRV96] A. Tomasic, L. Raschid, and P. Valduriez. Scaling Heterogeneous Databases and the Design of DISCO. *ICDCS Conf.*, Hong Kong, 1996.

Attachment 2a

Library Catalog: Search for Books, Journals, Music and More


```
LEADER 00000cam  22000004a 4500
001    45316889
003    OCoLC
005    20010227162156.0
008    001106s2001    cc            001 0 eng
010       00065206
020    1565927443
040    DLC|cDLC|dC#P|dIXA
049    IXAA
050 00 QA76.73.S67|bK55 2001
100 1  Kline, Kevin E.,|d1966-
245 10 SQL in a nutshell :|ba desktop quick reference /|cKevin
       Kline with Daniel Kline.
250    1st ed.
260    Beijing ;|aCambridge [Mass.] :|bO'Reilly,|c2001.
300    viii, 214 p. ;|c23 cm.
500    Includes index.
630 00 SQL server.
650  0 SQL (Computer program language)
650  0 Client/server computing.
700 1  Kline, Daniel.
915    (TxU)11061992
915    45316889
```

Location	Call No.	Current Status
PCL Stacks	QA 76.73 S67 K55 2001	AVAILABLE

Attachment 2b

OCLC 45316889	**No holdings in XXX - 335 other holdings**										
Books							**Rec Stat** c	**Entered** 20001106		**Replaced** 20171117224538.9	
Type	a	**ELvl**	4	**Srce**		**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	mau
		Cont		**GPub**		**LitF**	0	**Indx**	1		
Desc	a	**Ills**		**Fest**	0	**DtSt**	s	**Dates**	2001 ,		

010 00065206
040 DLC $b eng $c DLC $d UKM $d NLGGC $d BTCTA $d YDXCP $d CBC $d IG# $d OCLCQ $d CNCGM $d OCLCF $d OCLCQ $d I8M $d OCLCQ $d OCLCO
015 GBA328945 $2 bnb
019 52056818 $a 677400556
020 1565927443
020 9781565927445
042 pcc
050 00 QA76.73.S67 $b K55 2001
082 00 005.75/85 $2 21
084 54.64 $2 bcl
090 $b
100 1 Kline, Kevin E., $d 1966-
245 10 SQL in a nutshell : $b a desktop quick reference / $c Kevin Kline with Daniel Kline.
260 Cambridge, Mass. : $b O'Reilly & Associates, $c ©2001.
300 viii, 214 pages ; $c 23 cm
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
500 Includes index.
505 0 Ch. 1. SQL, Vendor Implementations, and Some History -- Ch. 2. Foundational Concepts -- Ch. 3. SQL Statements Command Reference -- Ch. 4. SQL Functions -- Ch. 5. Unimplemented SQL99 Commands -- Appendix. SQL99 and Vendor-Specific Keywords.
>520 1 "SQL (Structured Query Language) is a standardized query language for maintaining and querying database information. Historically, SQL has been the choice for database management systems running on minicomputers and mainframes. Increasingly, however, SQL is being adapted to PC systems that support distributed databases and allow users on a local-area network to access the same data simultaneously. Although there are different dialects of SQL, it is the closest thing to a standard query language that currently exists." "SQL in a Nutshell is a practical and useful command reference to the latest release of the Structured Query Language standard (SQL99), which assists readers in learning how their favorite database product supports any standard SQL command. This book presents each of the SQL commands and describes its use in both commercial (Microsoft SQL Server 2000 and Oracle 8i) and open source (MySQL, PostgreSQL 7.0) implementations. Each command reference includes the command syntax (by vendor, if the syntax differs across implementations), a description, and informative examples that illustrate important concepts and uses." "SQL in a Nutshell is more than a convenient reference guide for experienced SQL programmers, analysts, and database administrators. It's also a great learning resource for novice and auxiliary SQL users, such as system administrators,

users of packaged client/server products, and consultants who need to be familiar with the various SQL dialects across many platforms."--Jacket.
630 00 SQL server.
650 0 SQL (Computer program language)
650 0 Client/server computing.
630 07 SQL server. $2 fast $0 (OCoLC)fst01376937
650 7 Client/server computing. $2 fast $0 (OCoLC)fst00864205
650 7 SQL (Computer program language) $2 fast $0 (OCoLC)fst01102869
650 17 SQL. $2 gtt
700 1 Kline, Daniel.
856 41 $3 Table of contents $u http://catdir.loc.gov/catdir/enhancements/fy1008/00065206-t.html
856 42 $3 Contributor biographical information $u
http://catdir.loc.gov/catdir/enhancements/fy1008/00065206-b.html
856 42 $3 Publisher description $u http://catdir.loc.gov/catdir/enhancements/fy0915/00065206-d.html

Attachment 2c

The Library of Congress

>> Go to Library of Congress Online Catalog

LIBRARY OF CONGRESS AUTHORITIES

Help | New Search | Search History | Headings List | Start Over

Previous Next

MARC Display | Labelled Display

LC control no.: sh 86006628
LCCN Permalink: https://lccn.loc.gov/sh86006628
HEADING: SQL (Computer program language)

000 00894cz a2200241n 450
001 4804241
005 20120329071628.0
008 861007|| anannbabn |a ana
010 __ **|a** sh 86006628
035 __ **|a** (DLC)sh 86006628
035 __ **|a** (DLC)7277456
035 __ **|a** (DLC)sp 86006628
035 __ **|a** (DLC)364310
040 __ **|a** DLC **|c** DLC **|d** DLC **|d** ABAU
150 __ **|a** SQL (Computer program language)
450 __ **|a** Structured Query Language (Computer program language)
550 __ **|w** g **|a** Declarative programming languages
550 __ **|w** g **|a** Query languages (Computer science)
670 __ **|a** Work cat.: Date, C.J. A guide to the SQL standard, 1987 **|b** (a relational database language)
670 __ **|a** Wikipedia, July 19, 2007 **|b** (SQL ... Technically, SQL is a declarative computer language for use with SQL databases)
906 __ **|t** 0738 **|u** te04 **|v** 0
952 __ **|a** LC pattern: AL (Computer program language)
953 __ **|a** fg09 **|b** ym58

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

Help - Search - Search History - Headings List - Start Over

Library of Congress Library of Congress Authorities

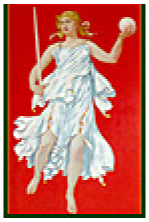

URL: https://www.loc.gov/

Mailing Address:
101 Independence Ave, S.E.
Washington, DC 20540

URL: http://authorities.loc.gov/
Library of Congress Online Catalog
URL: https://catalog.loc.gov/

Questions, comments, error reports: Contact Us

Attachment 2d

LC control no.: sh 93000502
LCCN Permalink: https://lccn.loc.gov/sh93000502
HEADING: Client/server computing
000 01122cz a2200277n 450
001 4847855
005 20120731055937.0
008 930203i| anannbabn |a ana
010 __ **|a** sh 93000502
035 __ **|a** (DLC)sh 93000502
035 __ **|a** (DLC)194540
040 __ **|a** DLC **|c** DLC **|d** DLC **|d** ABAU
053 _0 **|a** QA76.9.C55
150 __ **|a** Client/server computing
450 __ **|a** Computing, Client/server
550 __ **|w** g **|a** Electronic data processing
550 __ **|w** g **|a** Electronic data processing **|x** Distributed processing
670 __ **|a** Work cat.: 93009429: Dewire, D.T. Client/server computing **|b** (Distributed processing in which an application is split into tasks with certain tasks performed on the user's (client) machine and others on the host (server) computer)
670 __ **|a** LC database, Jan. 28, 1993 **|b** (client/server computing; client-server architecture)
670 __ **|a** Byte, Jul., 1989: **|b** p. 215.
675 __ **|a** Computer dict.; **|a** Stand. gloss. of computer term.
906 __ **|t** 9548 **|u** te03 **|v** 1
910 __ **|a** Proposal saved by tc09 on 07/26/2012 at 07:33:13
952 __ **|a** LC pattern: End-user computing
953 __ **|a** jf11 **|b** ec11
906 __ **|t** 1207 **|v** 0

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

Help - Search - Search History - Headings List - Start Over

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

Attachment 2e

The Library's catalogs are currently affected by performance issues. Staff are investigating the problem. We apologize for any inconvenience to our users.

LC control no.: n 90684343
LCCN Permalink: https://lccn.loc.gov/n90684343
HEADING: SQL server
000 00596cz a2200181n 450
001 3987471
005 20080321051519.0
008 900821n| acannaabn |a ana
010 __ **|a** n 90684343
035 __ **|a** (OCoLC)oca02802619
035 __ **|a** (Uk)002611964
040 __ **|a** DLC **|c** DLC **|d** IAhCCS **|d** Uk
130 _0 **|a** SQL server
430 _0 **|a** Microsoft SQL server
667 __ **|a** Do not establish individual versions, e.g. SQL server 2002
670 __ **|a** Nath, A. The guide to SQL server, c1990: **|b** CIP t.p. (SQL server) galley (relational data base server)
670 __ **|a** Microsoft SQL server 2002 unleashed, c2002.
953 __ **|a** vf03

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

Attachment 2f

```
Leader 00000cam a22002774a 4500
001    990095687500203941
005    20080903165901.0
007    cr |n|||||||||
008    001106s2001    mau           001 0 eng
010    ##$a^^^00065206^
020    ##$a1565927443$0(uri) http://www.isbnsearch.org/isbn/1565927443
035    ##$a(MH)009568750HVD01-Aleph
035    ##$a(OCoLC)45316889$0(uri) http://www.worldcat.org/oclc/45316889
040    ##$aDLC$cDLC$dOCLCQ
042    ##$apcc
050    00$aQA76.73.S67$bK55 2001
082    00$a005.75/85$221
100    1#$aKline, Kevin E.,$d1966-$0(uri) http://id.loc.gov/authorities/names/no97002232$0(uri) http://viaf.org/viaf/sourceID/LC|no9700
245    10$aSQL in a nutshell :$ba desktop quick reference /$cKevin Kline with Daniel Kline.
260    ##$aCambridge, Mass. :$bO'Reilly & Associates,$cc2001.
300    ##$aviii, 214 p. ;$c23 cm.
500    ##$aIncludes index.
505    0#$aCh. 1. SQL, Vendor Implementations, and Some History -- Ch. 2. Foundational Concepts -- Ch. 3. SQL Statements Command Refere
520    1#$a"SQL (Structured Query Language) is a standardized query language for maintaining and querying database information. Histori
630    00$aSQL server.$0(uri) http://id.loc.gov/authorities/names/n90684343$0(uri) http://viaf.org/viaf/sourceID/LC|n90684343$2LCNAMES$
630    00$aMicrosoft SQL server$2LCNAMES$041-LIBRARY_OF_CONGRESS-n  90684343$9N
650    #0$aSQL (Computer program language)$0(uri) http://id.loc.gov/authorities/subjects/sh86006628$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86
650    #0$aClient/server computing.$0(uri) http://id.loc.gov/authorities/subjects/sh93000502$2LCSH$041-LIBRARY_OF_CONGRESS-sh 93000502$
650    #0$aComputing, Client/server$2LCSH$041-LIBRARY_OF_CONGRESS-sh 93000502$9N
650    #0$wg$aElectronic data processing$2LCSH$041-LIBRARY_OF_CONGRESS-sh 93000502$9R
650    #0$wg$aElectronic data processing$xDistributed processing$2LCSH$041-LIBRARY_OF_CONGRESS-sh 93000502$9R
650    #0$aStructured Query Language (Computer program language)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9N
650    #0$wg$aDeclarative programming languages$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9R
650    #0$wg$aQuery languages (Computer science)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9R
700    1#$aKline, Daniel.$0(uri) http://id.loc.gov/authorities/names/n00008705$0(uri) http://viaf.org/viaf/sourceID/LC|n00008705$2LCNAM
856    41$uhttp://nrs.harvard.edu/urn-3:hul.ebookbatch.SAFAR_batch:1565927443$8222082674510003941
900    ##$a009568750
998    ##$aNO-OVERLAY
AVA    ##$0990095687500203941$8222082674570003941$a01HVD_INST$bCAB$cGEN$dQA76.73.S67 K55 2001$eavailable$f1$g0$jGEN$k0$p1$qCabot Scienc
AVA    ##$0990095687500203941$8222082674510003941$a01HVD_INST$bNET$cGEN$echeck_holdings$jCOLIN$k#$p2$qNetworked Resource (NET)
INST   ##$a01HVD_INST
INT    ##$aP
```

Attachment 2g

Search

My Account Log in

SQL In a Nutshell: A Desktop Quick Reference

Advanced

Everything Catalog Articles Databases Online Journals Library Website

Catalog / Record

Book

SQL in a nutshell : a desktop quick reference / Kevin Kline with Daniel Kline.

Record Info

Main Author
Kline, Kevin E., 1966-

Published
2001 - Beijing ; Cambridge [Mass.] : O'Reilly,

Contributors
Kline, Daniel.

Edition
1st ed.

Show complete record

Art Architecture & Engineering Book Collection About location
2nd Floor

Action	Description	Status	Call Number
Get this	N/A	Missing	QA 76.73 .S67 K551 2001

Offsite Shelving

Action	Description	Status	Call Number
Get this	N/A	On shelf	QA 76.73 .S67 K551 2001

Actions Select what to do with this record.

Email Text Zotero Endnote Export RIS

Ask Ask a Librarian

MARC Data

LDR			00000nam a22003134a 4500
001			004145590
003			MiU
005			20070613111602.0
008			010305s2001 cc 001 0 eng
010			|a 00065206
020			|a 1565927443
035			|a (OCoLC)ocm45316889
035			|a (III)iiio40182113
040			|a DLC |c DLC |d C#P |d EYM
042			|a pcc
050	0	0	|a QA76.73.S67 |b K55 2001
082	0	0	|a 005.75/85 |2 21
100	1		|a Kline, Kevin E., |d 1966-
245	1	0	|a SQL in a nutshell : |b a desktop quick reference / |c Kevin Kline with Daniel Kline.
250			|a 1st ed.
260			|a Beijing ; |a Cambridge [Mass.] : |b O'Reilly, |c 2001.
300			|a viii, 214 p. ; |c 23 cm.
500			|a Includes index.
630	0	0	|a SQL server.
650		0	|a SQL (Computer program language)
650		0	|a Client/server computing.
700	1		|a Kline, Daniel.
852	0		|a MiU |b AAEL |c BKS |h QA 76.73 .S67 K551 2001 |x iiio40

Ask Ask a Librarian

852	0	|a MiU |b OFFS |c STATE |h QA 76.73 .S67 K551 2001 |x Transfer from AAEL 2017828
970		|a BK |b Book
971		|a MiU
972		|c 20040625
972		|c 20070613
973		|a AC |b avail_circ
974		|a 39015049126926 |b AAEL |c BKS |f 01 |s MI
974		|a 39015069321951 |b OFFS |c STATE |f 01
998		|c alexismz-monorec AddCopySame20070613
998		|c mcp jb

Attachment 2h

Roger

The UC San Diego Library Catalog


```
LEADER 00000cam a22004814a 4500
001    45316889
005    20150122180258.0
008    001106t20012001mau          001 0 eng
010    00065206
015    GBA3-28945
019    52056818
020    1565927443
020    9781565927445
035    (OCoLC)45316889|z(OCoLC)52056818
040    DLC|cDLC|dUKM|dNLGGC|dBTCTA|dYDXCP|dUtOrBLW|dOrLoB-B
       |dUtOrBLW
042    pcc
050 00 QA76.73.S67|bK55 2001
082 00 005.75/85|221
100 1  Kline, Kevin E.,|d1966-
245 10 SQL in a nutshell :|ba desktop quick reference /|cKevin
       Kline with Daniel Kline
260    Cambridge, Mass. :|bO'Reilly & Associates,|c[2001], ©2001
300    viii, 214 pages ;|c23 cm
336    text|btxt|2rdacontent
337    unmediated|bn|2rdamedia
338    volume|bnc|2rdacarrier
500    Includes index
530    Also issued online
630 00 SQL server
650  0 SQL (Computer program language)
650  0 Client/server computing
700 1  Kline, Daniel
910    rec/jkm
910    MARS
910    CUS holdings added post-reclamation. rjf091211
910    RDA ENRICHED
938    Baker and Taylor|bBTCP|n00065206
938    YBP Library Services|bYANK|n1719470
970 01 |tPreface
970 11 |lCh. 1|tSQL, Vendor Implementations, and Some History|p1
970 11 |lCh. 2|tFoundational Concepts|p9
970 11 |lCh. 3|tSQL Statements Command Reference|p27
970 11 |lCh. 4|tSQL Functions|p163
970 11 |lCh. 5|tUnimplemented SQL99 Commands|p194
970 11 |lAppendix|tSQL99 and Vendor-Specific Keywords|p197
970 01 |tIndex|p205
```

Location	Call Number	Status
Books, Geisel Floor1 East	QA76.73.S67 K55 2001	AVAILABLE

RSS Feeds
New Books, Videos, etc.

Recommend Purchases
Suggest a title for purchase

Help
View Roger help topics

Attachment 2i

LIBRARY CATALOG

Search Navigation

1 of 5

BOOK

SQL in a nutshell : a desktop quick reference

Full Record MARC Tags

000		01297cam a22003134a 4500
001		12223581
005		20100706150633.0
008		001106s2001 mau 001 0 eng
906	__	|**a** 7 |**b** cbc |**c** orignew |**d** 1 |**e** ocip |**f** 20 |**g** y-gencatlg
925	0_	|**a** acquire |**b** 2 shelf copies |**x** policy default
955	__	|**a** to ASCD pc16 11-06-00; jf02 11-06-00; jf12 to Dewey 11-06-00; aa05 11-07-00; CIP ver. jf03 06-14-01; jf12 to BCCD 06-18-01
010	__	|**a** 00065206
020	__	|**a** 1565927443
040	__	|**a** DLC |**c** DLC |**d** DLC
042	__	|**a** pcc
050	00	|**a** QA76.73.S67 |**b** K55 2001
082	00	|**a** 005.75/85 |**2** 21
100	1_	|**a** Kline, Kevin E., |**d** 1966-
245	00	|**a** SQL in a nutshell : |**b** a desktop quick reference / |**c** Kevin Kline with Daniel Kline.
260	__	|**a** Cambridge, Mass. : |**b** O'Reilly & Associates, |**c** c2001.
300	__	|**a** viii, 214 p. ; |**c** 23 cm.
500	__	|**a** Includes index.
650	_0	|**a** SQL (Computer program language)
630	00	|**a** SQL server.
650	_0	|**a** Client/server computing.
700	1_	|**a** Kline, Daniel.
856	42	|**3** Publisher description |**u** http://www.loc.gov/catdir/enhancements/fy0915/00065206-d.html

856	42	|**3** Contributor biographical information |**u** http://www.loc.gov/catdir/enhancements/fy1008/00065206-b.html
856	41	|**3** Table of contents only |**u** http://www.loc.gov/catdir/enhancements/fy1008/00065206-t.html

Request this Item

LC Find It

Item Availability

CALL NUMBER	QA76.73.S67 K55 2001 Copy 2
Request in	Jefferson or Adams Building Reading Rooms
Status	Not Charged

CALL NUMBER	QA76.73.S67 K55 2001 LANDOVR Copy 1
Request in	Jefferson or Adams Building Reading Rooms - STORED OFFSITE
Status	Not Charged

Attachment 2j

ASP in a **nutshell**: a **desktop quick reference**
K Weissinger - 2000 - books.google.com
… 3. Microsoft Internet information server. 4. **SQL** server. 5. Web servers … So you're not alone. You hold in your hands the **desktop reference** for this exciting technology … All rights reserved. Page 16. Preface How This Book Is Structured ASP in a **Nutshell** is divided into three parts …
Cited by 81 Related articles All 6 versions

Java Enterprise in a **nutshell**: a **desktop quick reference**
J Farley, W Crawford, D Flanagan - 2002 - books.google.com
… Preface This book is a **desktop quick reference** for Java™ programmers who are writing enterprise applications … help you work with technologies key to the Enterprise APIs, such as **SQL** and IDL … Java in a **Nutshell**, Java Foundation Classes in a **Nutshell**, and J2ME in a **Nutshell** …
Cited by 22 Related articles All 2 versions

Oracle in a **Nutshell**: A **Desktop Quick Reference**
R Greenwald, DC Kreines - 2002 - books.google.com
… We hope that this book will become an invalu- able **desktop reference** for Oracle users of all … the greatest challenges that faced us was the most basic—how to create a **Nutshell** book for … bold Used in examples that show interaction between an Oracle utility (eg, **SQL***Plus, RMAN …
Cited by 7 Related articles All 4 versions

SQL in a **nutshell**: a **desktop quick reference**
KE Kleine, D Kline - 2001 - O'Reilly & Associates
Cited by 3 Related articles

Dreamweaver in a **Nutshell**: A **Desktop Quick Reference**
H Williamson, B Epstein - 2002 - books.google.com
… **Nutshell** Handbook, the **Nutshell** Handbook logo, and the O'Reilly logo are registered trademarks … view uses visual tools, which would be available in a **desktop** publishing program … recordset status life objects • Dynamic data formatting • Server behavior library • **SQL** Query Editor …
Cited by 6 Related articles All 3 versions

Design of a web-based medical database for computer-assisted orthopedic surgery
T Wu, G Müller, E Schkommodau… - International Congress …, 2001 - Elsevier
… AMIA Symp. (1999), pp. 804-808. [7] J. Byrne. Building Microsoft **SQL** Server 7 Web Sites, Prentice-Hall, Upper Saddle River, NJ (1999). [8] A. Weissinger. ASP in a **Nutshell**: A **Desktop Quick Reference**, O'Reilly and Associates, Sebastopol, CA (2000) …
Cited by 2 Related articles All 2 versions

A remote cooperative design system using interactive 3D graphics
Y Luo, R Galli, D Sanchez, A Bennasar… - … Journal of Image and …, 2001 - World Scientific
… de junio de 1999. 9. AK Weissinger, ASP in a **Nutshell**, A **Desktop Quick Reference** (O'Reilly, 1999). 10. M. Otey and P. Conte, **SQL** Server 7, Developer's **Guide** (Obsborne McGraw-Hill, 1999). 11. S. Spainhour and V …
Cited by 6 Related articles All 3 versions

Learning oracle pl/**sql**
B Pribyl, S Feuerstein - 2002 - books.google.com
… Although neither of O'Reilly's two books on **SQL** qualify as tutorials, you may still find them helpful to have on your bookshelf: Oracle **SQL**: The Essential **Reference**, and **SQL** in a **Nutshell**: A **Desktop Quick Reference**, the latter of which addresses multiple vendors' versions …
Cited by 15 Related articles All 16 versions

Using transact-**SQL** and simulation techniques to create virtual M&M'S
RM Snyder - Journal of Computing Sciences in Colleges, 2002 - dl.acm.org
… Myrtle Beach, SC. [11] Stephens, R., et al. (1997). Sams teach yourself **SQL** in 21 days. Indianapolis, IN: Sams. [12] Weissinger, A. (2000). ASP in a **nutshell**: A **desktop quick reference**. Sebastopol, CA: O'Reilly & Associates.
Related articles All 2 versions

CORBA and Web technologies applied to long distance learning

AJPD Fonte, JC Metrolho - … '01. 27th Annual Conference of the …, 2001 - ieeexplore.ieee.org

… The GUI uses Active Server Page (ASP) technology, which includes ActiveX Data Objects (ADO) to open a connection and query requests on a **SQL** database to validate remote users [8]. Also, it was … [3] D. Flanagan, Java in a **nutshell** - A **desktop Quick Reference**, O'Reilly …

Attachment 3a

Library Catalog: Search for Books, Journals, Music and More


```
LEADER 00000cam  2200000 a 4500
001    41291586
003    OCoLC
005    20020626161702.0
008    990503s1999    caua          001 0 eng d
010       99222909
015    GB99-74559
020    0072118458
040    WU|cDLC|dUKM|dIXA
042    lccopycat
049    IXAA
050 00 QA76.73.S67|bG758 1999
082 00 005.75/65|221
100 1  Groff, James R.
245 10 SQL, the complete reference /|cJames R. Groff, Paul N.
       Weinberg.
246 14 Complete reference SQL
260    Berkeley, Calif. :|bOsborne/McGraw-Hill,|cc1999.
300    xxviii, 994 p. :|bill. ;|c24 cm +|e1 CD-ROM (4 3/4 in.)
500    Includes index.
538    System requirements: Windows 95/NT.
650  0 SQL (Computer program language)
650  0 Relational databases.
650  4 Relational databases.
650  4 SQL (Computer program language)
700 1  Weinberg, Paul N.
915    (TxU)11833473
915    41291586
```

Location	Call No.	Current Status
PCL Check Out Desk On Site Storage - ASK AT PCL CHECK OUT DESK or click Pick It Up	QA 76.73 S67 G758 1999 CD-ROM	AVAILABLE
PCL Stacks	QA 76.73 S67 G758 1999 TEXT	AVAILABLE

Attachment 3b

OCLC 41291586 **No holdings in XXX - 193 other holdings**

Books						**Rec Stat**		**Entered**	19990503	**Replaced**	20181219233938.8
Type	a	**ELvl**		**Srce**	d	**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	cau
		Cont		**GPub**		**LitF**	0	**Indx**	1		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	1999 ,		

007 c $b o $e g
010 99222909
040 WU $b eng $c DLC $d UKM $d BTCTA $d YDXCP $d TAMCT $d DEBBG $d N5L $d OCLCQ $d OCLCF $d OCLCQ $d OCLCO $d OCLCQ $d I8M $d OCLCO $d OCLCQ $d APUMS $d EUM $d OCLCQ $d LVT
015 GB9974559 $2 bnb
015 b98W5638
019 1027923215 $a 1038160464 $a 1038550086 $a 1049750111 $a 1056823219 $a 1057079021 $a 1076724096
020 0072118458
020 9780072118452
042 lccopycat
050 00 QA76.73.S67 $b G758 1999
082 00 005.75/65 $2 21
084 ST 271 S70 $2 rvk
090 $b
100 1 Groff, James R.
245 10 SQL, the complete reference / $c James R. Groff, Paul N. Weinberg.
246 14 Complete reference SQL
260 Berkeley, Calif. : $b Osborne/McGraw-Hill, $c ©1999.
300 xxviii, 994 pages : $b illustrations ; $c 24 cm + $e 1 computer optical disc (4 3/4 in.)
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
500 Includes index.
505 0 An Overview of SQL -- Retrieving Data -- Updating Data -- Database Structure -- Programming with SQL -- SQL Today and Tomorrow.
650 0 SQL (Computer program language)
650 0 Relational databases.
650 7 Relational databases. $2 fast $0 (OCoLC)fst01093575
650 7 SQL (Computer program language) $2 fast $0 (OCoLC)fst01102869
650 7 SQL $2 gnd $0 (DE-588)4134010-3
650 07 SQL. $2 swd
653 COMPUTER LANGUAGES ‡# UkMaJRU
653 SQL ‡# UkMaJRU
700 1 Weinberg, Paul N.
856 41 $3 Table of contents $u http://catdir.loc.gov/catdir/toc/mh023/99222909.html
856 42 $3 Publisher description $u http://catdir.loc.gov/catdir/description/mh023/99222909.html

Attachment 3c

LC control no.: sh 86007768
LCCN Permalink: https://lccn.loc.gov/sh86007768
HEADING: Relational databases

000 00975cz a2200217n 450
001 4805240
005 20120329071820.0
008 861205i| anannbabn |a ana
010 __ **|a** sh 86007768
035 __ **|a** (DLC)sh 86007768
035 __ **|a** (DLC)151925
040 __ **|a** DLC **|c** DLC **|d** DLC
150 __ **|a** Relational databases
450 __ **|w** nne **|a** Relational data bases
550 __ **|w** g **|a** Databases
670 __ **|a** Work cat.: Yang, C. Relational databases, 1986.
670 __ **|a** Standard glossary of computer terminology **|b** (Relational data base management system)
670 __ **|a** LC data base **|b** (relational database(s), relational data base(s))
675 __ **|a** National Computing Centre. Thesaurus of computing terms; **|a** Computer dictionary; **|a** Dictionary of data processing; **|a** Encyclopedia of computer science; **|a** Engr. index; **|a** Jordain. Condensed computer encyclopedia; **|a** McGraw-Hill Encyc. of Sci & Tech.; **|a** Web. 3; **|a** Webster's New World compact dict. of computer terms
906 __ **|t** 9349 **|u** te04 **|v** 0
953 __ **|a** fg01 **|b** ta28

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

Attachment 3d

Copyright United States Copyright Office

Help | Search | History | Titles | Start Over

Public Catalog

Copyright Catalog (1978 to present)
Search Request: Left Anchored Title = sql the complete reference
Search Results: Displaying 2 of 3 entries

previous | next

Labeled View

SQL : the complete reference / James R. Groff, Paul N.Weinberg.

Type of Work: Text
Registration Number / Date: TX0004891473 / 1999-05-05
Title: SQL : the complete reference / James R. Groff, Paul N.Weinberg.
Imprint: Berkeley : Osborne/McGraw-Hill, c1999.
Description: 994 p. + CD-ROM.
Copyright Claimant: McGraw-Hill Companies, Inc.
Date of Creation: 1999
Date of Publication: 1999-03-09
Previous Registration: Appl. states no claim is made in material on CD-ROM.
Variant title: SQL : the complete reference
Names: Groff, James R.
Weinberg, Paul N
McGraw-Hill Companies, Inc.

previous | next

Save, Print and Email (Help Page)	
Select Download Format	Full Record ▼ Format for Print/Save
Enter your email address:	Email

Help Search History Titles Start Over

Attachment 3e

MARC View

LEADER 02209nam a2200337 a 4500			
001			2362352
005			20030429125527.0
008			990503s1999 caua 001 0 eng d
010			|a 99222909
015			|a GB99-74559
020			|a 0072118458
035			|a (OCoLC)ocm41291586
035			|9 APP3277CU
040			|a WU |c DLC |d UKM |d OrLoB-B
042			|a lccopycat
050	0	0	|a QA76.73.S67 |b G758 1999
082	0	0	|a 005.75/65 |2 21
100	1		|a Groff, James R.
245	1	0	|a SQL, the complete reference / |c James R. Groff, Paul N. Weinberg.
246	1	4	|a Complete reference SQL
260			|a Berkeley, Calif. : |b Osborne/McGraw-Hill, |c c1999.
300			|a xxviii, 994 p. : |b ill. ; |c 24 cm + |e 1 computer laser optical disc (4 3/4 in.)
500			|a Includes index.
505	0	0	|g 1. |t Introduction -- |g 2. |t A Quick Tour of SQL -- |g 3. |t SQL In Perspective -- |g 4. |t Relational Databases -- |g 5. |t SQL Basics -- |g 6. |t Simple Queries -- |g 7. |t Multi-Table Queries (Joins) -- |g 8. |t Summary Queries -- |g 9. |t Subqueries and Query Expressions -- |g 10. |t Database Updates -- |g 11. |t Data Integrity -- |g 12. |t Transaction Processing -- |g 13. |t Creating a Database -- |g 14. |t Views -- |g 15. |t SQL Security -- |g 16. |t The System Catalog -- |g 17. |t Embedded SQL -- |g 18. |t Dynamic SQL -- |g 19. |t Sql Apis -- |g 20. |t Database Processing and Stored Procedures -- |g 21. |t SQL and Data Warehousing -- |g 22. |t SQL Networking and Distributed Databases -- |g 23. |t SQL and Objects -- |g 24. |t The Future of SQL -- |g A. |t The Sample Database -- |g B. |t Database Vendor Profiles -- |g C. |t Company and Product List -- |g D. |t SQL Syntax Reference -- |g E. |t SQL Call Level Interface -- |g F. |t SQL Information Schema Standard -- |g G. |t CD-ROM Installation Guide.
520	1		|a "SQL: The Complete Reference provides an in-depth discussion of SQL fundamentals, modern SQL products, and SQLs role in trends such as data warehousing, "thinclient" architectures, and Internet-based e-commerce. This book is your one-stop resource for all you need to know about SQL."--BOOK JACKET.
650		0	|a SQL (Computer program language)
650		0	|a Relational databases.
650		4	|a Relational databases.
650		4	|a SQL (Computer program language)
700	1		|a Weinberg, Paul N.
900			|a AUTH |b TOC
852	0		|0 2845140 |a Offsite - Place Request for delivery within 2 business days |b off,eng |h QA76.73.S67 G758 1999 |x [ang] --mix
876			|0 2845140 |3 book & CD |a 2834603 |p CU72114142 CGD: Shared
990			|a QA
991			|a QA76.73.S67 G758 1999 |b loc+QA-+0076+73000000+S6700000+G7580000+001999
992			|a Offsite |b Offsite - Place Request for delivery within 2 business days >> QA76.73.S67 G758 1999|DELIM|2845140
995			|a catalog

MARC Record Load
2019-02-03 2:39am

Authority Variants

Author Variants
Groff, Jim

Subject Variants
Structured Query Language (Computer program language)
Relational data bases

Geographic Variants
-- none --
Last Variant Update
2019-02-03 2:39am

Attachment 3f

SQL, the complete reference /

James R. Groff, Paul N. Weinberg.

Main Author:	Groff, James R.
Other Names:	Weinberg, Paul N.
Published:	Berkeley, Calif. : Osborne/McGraw-Hill, c1999.
Topics:	SQL (Computer program language) | Relational data bases.
Tags:	005.7565, G893s

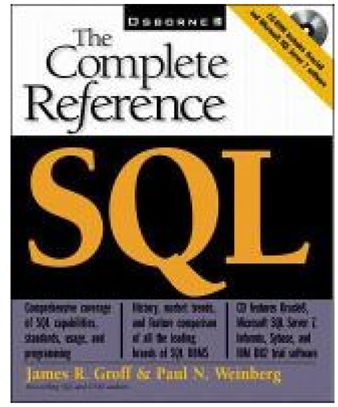

- More Details
- Location & Availability
- Table of Contents
- User Reviews
- Published Reviews
- Request Item

000 00930cam a2200313 a 4500
001 4181220
003 UIUdb
005 20020415165455.0
008 990503s1999 caua 001 0 eng d
010 **|a** 99222909
020 **|a**0072118458
035 **|a**(OCoLC)ocm41291586
035 **|9**DVY-0060
040 **|a**WU**|c**DLC**|d**UIU
042 **|a**lccopycat
049 **|a**UIUJ
050 00**|a**QA76.73.S67**|b**G758 1999
082 00**|a**005.75/65**|2**21
092 0 **|a**005.7565**|b**G893s**|2**21
100 1 **|a**Groff, James R.
245 10**|a**SQL, the complete reference /**|c**James R. Groff, Paul N. Weinberg.
246 14**|a**Complete reference SQL

260 **|a**Berkeley, Calif. :**|b**Osborne/McGraw-Hill,**|c**c1999.
300 **|a**xxviii, 994 p. :**|b**ill. ;**|c**24 cm +**|e**1 computer laser optical disc (4 3/4 in.)
500 **|a**Includes index.
650 0**|a**SQL (Computer program language)
650 0**|a**Relational data bases.
700 1 **|a**Weinberg, Paul N.
994 **|a**02**|b**UIU

Attachment 3g

```
LEADER 00000nam  2200385 a 4500
001    41291586
005    19991123  TAPE AAPP1123: 0050
006    m        b
007    couugu
008    991124s1999    caua          001 0 eng d
009    Reclvl: f Addate: 991124 Addid: PCT Moddate: 080606 Modid:
       XRF
010       99222909
015    GB99-74559
020    0072118458
029 1  UKM|bb9974559
035    GLADN184811874
040    WU|cDLC|dUKM
042    lccopycat
050 00 QA76.73.S67|bG758 1999
082 00 005.75/65|221
090    QA76.73.S67|bG758 1999
100 1  Groff, James R.
245 10 SQL, the complete reference /|cJames R. Groff, Paul N.
       Weinberg.
246 14 Complete reference SQL
260    Berkeley, Calif. :|bOsborne/McGraw-Hill,|cc1999.
300    xxviii, 994 p. :|bill. ;|c24 cm +|e1 computer laser
       optical disc (4 3/4 in.)
500    Includes index.
538    System requirements: Windows 95/NT.
650  0 SQL (Computer program language)
650  0 Relational databases.
700 1  Weinberg, Paul N.
935    GLADN184811874
957    OCLC xref loaded 20110805; IR: 739629551
994    92|bCUY
```

Location	Call No.	Status
Engineering	QA76.73.S67 G758 1999	AVAILABLE
Engineering CD Cabinet	QA76.73.S67 G758 1999 compu/d	AVAILABLE

Attachment 3h

LIBRARY | CATALOG

Search Navigation

4 of 19

BOOK

SQL, the complete reference

Full Record | MARC Tags

000		01385cam a2200373 a 4500
001		2160717
005		20021126124653.0
008		990503s1999 caua 001 0 eng d
035	__	**|a** (DLC) 99222909
906	__	**|a** 7 **|b** cbc **|c** copycat **|d** 2 **|e** ncip **|f** 19 **|g** y-gencatlg
955	__	**|a** pb17 07-01-99 to cat.; jf02 7-7-99; jf11 to sl 07-09-99; jf25 07-15-99 to ddc
010	__	**|a** 99222909
020	__	**|a** 0072118458
035	__	**|a** (OCoLC)41291586
040	__	**|a** WU **|c** WU **|d** DLC
042	__	**|a** lccopycat
050	00	**|a** QA76.73.S67 **|b** G758 1999
082	00	**|a** 005.75/65 **|2** 21
100	1_	**|a** Groff, James R.
245	10	**|a** SQL, the complete reference / **|c** James R. Groff, Paul N. Weinberg.
246	14	**|a** Complete reference SQL
260	__	**|a** Berkeley, Calif. : **|b** Osborne/McGraw-Hill, **|c** c1999.
300	__	**|a** xxviii, 994 p. : **|b** ill. ; **|c** 24 cm + **|e** 1 computer laser optical disc (4 3/4 in.)
500	__	**|a** Includes index.
650	_0	**|a** SQL (Computer program language)
650	_0	**|a** Relational databases.
700	1_	**|a** Weinberg, Paul N.

856	42	**|3** Publisher description **|u** http://www.loc.gov/catdir/description/mh023/99222909.html
856	41	**|3** Table of contents **|u** http://www.loc.gov/catdir/toc/mh023/99222909.html
920	__	**|a** ** LC HAS REQ'D # OF SHELF COPIES **
922	__	**|a** co
985	__	**|e** GAP
991	__	**|b** c-MRC **|h** QA76.73.S67 **|i** G758 1999 **|o** am **|t** Copy 1 **|w** BOOKS
991	__	**|b** c-MRC **|h** QA76.73.S67 **|i** G758 1999 **|t** Copy 1 **|w** BOOKS

Request this Item

LC Find It

Item Availability

CALL NUMBER	QA76.73.S67 G758 1999 FT MEADE Copy 2
Request in	Machine Readable Collections - STORED OFFSITE
Status	Not Charged

CALL NUMBER	QA76.73.S67 G758 1999 Copy 1
Request in	See Reference Staff. By Appt in Jefferson Main RR (MRC)
Status	Not Charged

Attachment 3i

SQL: the complete reference

JR Groff, PN Weinberg, AJ Oppel - 2002 - 160592857366.free.fr

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer market. This book offers a conceptual framework for understanding and using SQL, describes the history of SQL and SQL standards, and explains the role of SQL in the computer industry today. It will show you, step-by-step, how to use SQL features, with many …

SQL.: The Complete Reference

P Weinberg - 2010 - lib.pyu.edu.vn

SQL: The Complete Reference, Third Edition provides a comprehensive, in-depth treatment of the SQL language for both technical and nontechnical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in today's computer industry. This book offers a conceptual framework for understanding and using SQL, describes the history of SQL and SQL standards, and explains the role of SQL in various computer industry segments, from enterprise data processing to data warehousing …

Showing the best results for this search. See all results

SQL: the complete reference

JR Groff, PN Weinberg - 160592857366.free.fr

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer …

SQL: The Complete Reference

JR Groff, PN Weinberg - vedantcollege.org

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer …

SQL The Complete Reference

J Groff, P Weinberg - 2009 - dl.acm.org

Get comprehensive coverage of every aspect of SQL from three leading industry experts. Revised with coverage of the latest RDBMS software versions, this one-stop guide explains how to build, populate, and administer high-performance databases and develop robust …

SQL: The Complete Reference

AR Groff, PN Weinberg - 1994 - thuvien.ued.udn.vn

Skip navigation …

SQL: The Complete Reference

JR Groff, PN Weinberg - Citeseer

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer …

SQL: The Complete Reference

JR Groff, PN Weinberg - files.learnmsbieasily.webnode.in

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer …

SQL: The Complete Reference

JR Groff, PN Weinberg - 1999 - lib.pyu.edu.vn

SQL: The Complete Reference provides a comprehensive, in-depth treatment of the SQL language for both technical and non-technical users, programmers, data processing professionals, and managers who want to understand the impact of SQL in the computer …

SQL: The Complete Reference

J Groff, P Weinberg - 2002 - dl.acm.org

Learn all you need to know about SQL with this one-stop resource. Updated to include information on Web services and XML, this complete guide also comes with a Windows 2000-and XP-compatible CD containing 4 databases (SQL Server 2000, IBM DB2 Version …

Attachment 4a

Map & Floorplans | Hours | InterLibrary Services | Article Databases

Library Catalog: Search for Books, Journals, Music and More

Help | Ask a Librarian | Renew & My Account

```
LEADER 00000nam  2200000Ia 4500
001    43333718
003    OCoLC
005    20000302094036.0
008    000124s2000    caua          001 0 eng d
020    0130161659
040    VYF|cVYF|dIXA
049    IXAA
090    QA76.73.S67|bI54 2000
245 00 Informix guide to SQL :|btutorial.
260    Menlo Park, CA :|bInformix Press ;|aUpper Saddle River,
       N.J. :|bPrentice Hall PTR,|cc2000.
300    1 v. (various pagings) :|bill. ;|c24 cm.
500    "Informix Enterprise Decision Server, version 8.3;
       Informix Dynamic Server.2000, version 9.2."
500    Includes index.
500    2nd ed.
650  0 SQL (Computer program language)
710 2  Informix Software, Inc.
915    (TxU)10448245
915    43333718
```

Location	Call No.	Current Status
Library Storage Facility -- click Pick It Up	QA 76.73 S67 I54 2000	DUE 07-15-19

Attachment 4b

OCLC	43333718	**No holdings in XXX - 59 other holdings**

Books

Rec Stat | **Entered** 20000802 | **Replaced** 20151212032617.8

Type	a	**ELvl**		**Srce**		**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	nju
		Cont	b	**GPub**		**LitF**	0	**Indx**	1		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	2000 ,		

010 00699312
040 DLC $b eng $c DLC $d BAKER $d BTCTA $d YDXCP $d AU@ $d OCLCF $d OCLCO $d OCLCQ
015 GB99V5259 $2 bnb
019 222581770 $a 835937013
020 0130161659
020 9780130161659
050 00 QA76.73.S67 $b I54 2000
082 00 005.75/65 $2 21
090 $b
245 00 Informix guide to SQL, tutorial.
250 2nd ed.
260 Menlo Park, CA : $b Informix Press ; $a Upper Saddle River, NJ : $b Prentice Hall PTR, $c ©2000.
300 1 volume (various pagings) : $b illustrations ; $c 24 cm
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
500 "Informix-SQL for UNIX, version 5.0."
504 Includes bibliographical references and index.
650 0 SQL (Computer program language)
650 7 SQL (Computer program language) $2 fast $0 (OCoLC)fst01102869
710 2 Informix Software, Inc.

Attachment 4c

Copyright
United States Copyright Office

Help | Search | History | Titles | Start Over

Public Catalog

Copyright Catalog (1978 to present)
Search Request: Left Anchored Title = Informix guide to SQL, tutorial
Search Results: Displaying 1 of 4 entries

previous next

Labeled View

Informix guide to SQL : tutorial.

Type of Work:	Text
Registration Number / Date:	TX0005125481 / 2000-01-19
Title:	Informix guide to SQL : tutorial.
Edition:	2nd ed.
Imprint:	Upper Saddle River, NJ : Prentice Hall PTR, c2000.
Description:	1 v.
Series:	Informix Press
Copyright Claimant:	Informix Corporation
Date of Creation:	1999
Date of Publication:	1999-11-18
Authorship on Application:	updated & new material: Informix Software, employer for hire.
Previous Registration:	Prev. reg. 1994, TXu 662-870.
Basis of Claim:	New Matter: updated & new material added.
Other Title:	Informix Press
Names:	Informix Corporation Informix Software

previous next

Save, Print and Email (Help Page)	
Select Download Format	Full Record ▼ Format for Print/Save
Enter your email address:	Email

Help Search History Titles Start Over

Contact Us | Request Copies | Get a Search Estimate | Frequently Asked Questions (FAQs) about Copyright | Copyright Office Home Page | Library of Congress Home Page

Attachment 4d

```
Leader 00915cam a2200265Ia 4500
001    990012934760302611
005    20181018110354.0
008    000124t20002000cau           001 0 eng d
020    ##$a0130161659
035    ##$a(OCoLC)43333718
040    ##$aVYF$beng$cVYF$dILU
049    ##$aILUU
090    ##$aQA76.73.S67$bI54 2000
245    00$aInformix guide to SQL :$btutorial.
264    #1$aMenlo Park, CA :$bInformix Press ;$aUpper Saddle River, N.J. :$bPrentice Hall PTR,$c[2000]
264    #4$c©2000
300    ##$a1 volume (various pagings) :$billustrations ;$c24 cm
336    ##$atext$btxt$2rdacontent
337    ##$aunmediated$bn$2rdamedia
338    ##$avolume$bnc$2rdacarrier
500    ##$a"Informix Enterprise Decision Server, version 8.3; Informix Dynamic Server.2000, version 9.2."
500    ##$aIncludes index.
650    #0$aSQL (Computer program language)
710    2#$aInformix Software, Inc.
994    ##$aE0$bILU
AVA    ##$0990012934760302611$822154294320002611$a01TTU_INST$bULIB$cUNIV LIB--Stacks$dQA76.73.S67 I54 2000$eavailable$f1$g0$jSTACK$k#$p
INST   ##$a01TTU_INST
INT    ##$aP
```

Attachment 4e

000 00829cam a2200265 a 4500
001 5481261
005 20160102113456.0
008 000802s2000 njua b 001 0 eng
010 __ |**a** 00699312
020 __ |**a** 0130161659
035 __ |**a** (OCoLC)43333718
035 __ |**a** ucoclc43333718
040 __ |**a** DLC |**c** DLC |**d** BAKER |**d** CLU |**d** UtOrBLW
049 __ |**a** CLUE |**l** L0089806020
050 00 |**a** QA76.73.S67 |**b** I54 2000
082 00 |**a** 005.75/65 |**2** 21
245 00 |**a** Informix guide to SQL, tutorial.
250 __ |**a** 2nd ed.
260 __ |**a** Menlo Park, CA : |**b** Informix Press ; |**a** Upper Saddle River, NJ : |**b** Prentice Hall PTR, |**c** c2000.
300 __ |**a** 1 v. (various pagings) : |**b** ill. ; |**c** 24 cm.
504 __ |**a** Includes bibliographical references and index.
650 _0 |**a** SQL (Computer program language)
710 2_ |**a** Informix Software, Inc.
910 __ |**a** anb 061010
910 __ |**a** MARS

This item
· Detailed Record
Staff MARC View
· **Permalink to this record**

Actions
· Request an Item
(e.g. Recall, Request from SRLF, Purchase)
· Print
· Email
· Add to My Bookbag
· View Melvyl Record

Google Books:

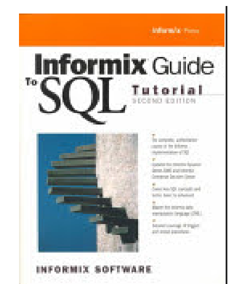

"About This Book"

Attachment 4f

```
Leader 01804nam a2200385Ia 4500
001    991028183449703611
005    20180627160954.0
008    000124t20002000cau           001 0 eng d
020    ##$a0130161659$0(uri) http://www.isbnsearch.org/isbn/0130161659
035    ##$ab19744882-01sful_inst
035    ##$a(OCoLC)43333718$0(uri) http://www.worldcat.org/oclc/43333718
040    ##$aVYF$cVYF$dUtOrBLW
050    #4$aQA76.73.S67$bI54 2000
099    ##$aQA 76.73 S67 I54 2000
245    00$aInformix guide to SQL :$btutorial.
250    ##$a2nd ed.
264    #1$aMenlo Park, CA :$bInformix Press ;$aUpper Saddle River, N.J. :$bPrentice Hall PTR,$c[2000]
264    #4$c©2000
300    ##$a1 volume (various pagings) :$billustrations ;$c24 cm
336    ##$atext$btxt$2rdacontent
337    ##$aunmediated$bn$2rdamedia
338    ##$avolume$bnc$2rdacarrier
500    ##$a"Informix Enterprise Decision Server, version 8.3; Informix Dynamic Server.2000, version 9.2."
500    ##$aIncludes index.
500    ##$aPrevious edition: Menlo Park, Calif.: Informix Press; London: Prentice Hall International, 1995.
505    0#$aCh. 1. Database Concepts -- Ch. 2. Composing SELECT Statements -- Ch. 3. Selecting Data From Complex Types -- Ch. 4. Using F
650    #0$aSQL (Computer program language)$0(uri) http://id.loc.gov/authorities/subjects/sh86006628$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86
650    #0$aStructured Query Language (Computer program language)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9N
710    2#$aInformix Software, Inc.$0(uri) http://id.loc.gov/authorities/names/no90006953$0(uri) http://viaf.org/viaf/sourceID/LC|no9000
901    ##$a2017-05-15$bMILL
902    ##$aQA 76.73 S67 I54 2000
907    ##$a.b19744882
979    ##$aOrder moved to BRS - Jun.6/01.
990    ##$a.b19744882$bben  $cm$da$e-$f000414$g991220$h160525$i11$j1$k0
991    ##$aben
996    ##$zRDA ENRICHED
AVA    ##$0991028183449703611$8221624449400003611$a01SFUL_INST$bbennett$cStacks$dQA 76.73 S67 I54 2000$eavailable$f1$g0$jstacks$k0$p1$qW
INST   ##$a01SFUL_INST
INT    ##$aP
```

Attachment 4g

LIBRARY | CATALOG

Search Navigation

2 of 4

BOOK

Informix guide to SQL, tutorial

Full Record | MARC Tags

000		00986cam a2200253 a 4500
001		12127876
005		20011116141322.0
008		000802s2000 njua b 001 0 eng
906	__	**|a** 7 **|b** cbc **|c** origcop **|d** 2 **|e** ncip **|f** 20 **|g** y-gencatlg
925	0_	**|a** acquire **|b** 2 shelf copies **|x** policy default
955	__	**|a** to ASCD pb11 08-02-00; jf00 08-04-00; jf01 (subj) 3-16-01; jf03 (desc.) to SL 03-27-01; jf25 2 copies to Dewey 05-03-01; aa20 05-07-01
010	__	**|a** 00699312
020	__	**|a** 0130161659
040	__	**|a** DLC **|c** DLC **|d** DLC
050	00	**|a** QA76.73.S67 **|b** I54 2000
082	00	**|a** 005.75/65 **|2** 21
245	00	**|a** Informix guide to SQL, tutorial.
250	__	**|a** 2nd ed.
260	__	**|a** Menlo Park, CA : **|b** Informix Press ; **|a** Upper Saddle River, NJ : **|b** Prentice Hall PTR **|c** c2000.
300	__	**|a** 1 v. (various pagings) : **|b** ill. ; **|c** 24 cm.
500	__	**|a** "Informix-SQL for UNIX, version 5.0."
504	__	**|a** Includes bibliographical references and index.
650	_0	**|a** SQL (Computer program language)
710	2_	**|a** Informix Software, Inc.

Request this Item

LC Find It

Item Availability

CALL NUMBER	QA76.73.S67 I54 2000 LANDOVR Copy 1
Request in	Jefferson or Adams Building Reading Rooms - STORED OFFSITE
Status	Not Charged

CALL NUMBER	QA76.73.S67 I54 2000 FT MEADE Copy 2
Request in	Jefferson or Adams Building Reading Rooms - STORED OFFSITE
Status	Not Charged

Attachment 5a

```
LEADER 00000cas a2200769 a 4500
001    25073822
005    19920717171056.0
008    920109c19919999nyumr p       0    0eng d
010    92659029
022    |y0362-1340
040    TJC|cTJC|dDLC|dNST|dWAU|dCOD
042    lc
049    DVDI
050 00 QA76.7|b.S54
090    QA76.7|b.A8
090    QA76.7|b.A8
222  0 ACM SIGPLAN notices
245 00 ACM SIGPLAN notices
246 2  Association for Computing Machinery Special Interest Group
       on Programming Languages notices
246 2  Special Interest Group on Programming Languages notices
246 10 Notices
246 13 SIGPLAN notices
260    New York, NY :|bACM Press,|c[1991-]
300    volumes :|billustrations ;|c28 cm
310    Monthly
336    text|btxt|2rdacontent
337    unmediated|bn|2rdamedia
338    volume|bnc|2rdacarrier
362 0  Vol. 26, no. 10 (Oct. 1991)-
500    Title from cover
500    "A monthly publication of the Special Interest Group on
       Programming Languages."
500    Some issues have former title: SIGPLAN notices
580    One issue of volume (i.e. v. 31, no. 9) for <1996>-
       contains the proceedings of the International Conference
       on Architectural Support for Programming Languages and
       Operating Systems (ASPLOS)
580    One issue of volume (i.e. v. 31, no. 10) for <1996>-
       contains the proceedings of OOPSLA (the Conference on
       Object-oriented Programming Systems, Languages, and
       Applications)
580    One issue of volume (i.e. v. 32, no. 7) for <1997>-
       contains the proceedings of the ACM SIGPLAN Symposium on
       Principles & Practice of Parallel Procramming (PPoPP)
580    Vol. 33, no. 7 (July 1998) contains ACM SIGPLAN/SIGSOFT
       Workshop on Program Analysis for Software and Tool
       Engineering (PASTE'98)
580    Vol. 34, no. 3 (March 1999) contains ACM SIGPLAN
       International Symposium on Memory Management (ISMM '98)
580    Vol. 34, no. 7 (July 1999) contains ACM SIGPLAN Workshop
       LCTES
580    One issue of volume (i.e. v.35, no.9 for 2000- contains
       ACM SIGPLAN International Conference on Functional
       Programming
580    Vol. 36, no. 6 contains Proceedings of the 28th ACM
       SIGPLAN-SIGACT Symposium on Principles of Programming
       (POPL 2001)
583 1  |326-36 (1991-2001) : |acommitted to retain|c20180101
       |d20430101|5CoU|fAlliance Shared Trust|uhttps://
       www.coalliance.org/shared-print-archiving-policies
650  0 Programming languages (Electronic computers)|vPeriodicals
710 2  ACM Special Interest Group on Programming Languages
711 2  International Conference on Architectural Support for
       Programming Languages and Operating Systems.|tProceedings
711 2  OOPSLA (Conference).|tOOPSLA ... conference proceedings
711 2  ACM SIGPLAN Symposium on Principles & Practice of Parallel
       Programming.|tProceedings
711 2  ACM SIGPLAN International Symposium on Memory Management
711 2  ACM SIGPLAN Workshop LCTES.|tProceedings of the ACM
       SIGPLAN ... Workshop on Languages, Compilers, and Tools
       for Embedded Systems (LCTES)
711 2  ACM SIGPLAN International Conference on Functional
       Programming.|tProceedings of the ... ACM SIGPLAN
```

```
       International Conference on Functional Programming
740 02 ASPLOS proceedings
740 02 Proceedings - International Conference on Architectural
       Support for Programming Languages and Operating Systems
740 02 Architectural support for programming languages and
       operating systems
740 02 Proceedings - OOPSLA
740 02 Conference proceedings - OOPSLA, Object-oriented
       Programming Systems, Languages, and Applications
740 02 Proceedings of the ACM SIGPLAN Symposium on Principles and
       Practice of Parallel Programming (PPoPP)
740 02 PPoPP proceedings
740 02 ACM SIGPLAN Symposium on Principles & Practice of Parallel
       Programming
740 02 PASTE'98 - ACM SIGPLAN/SIGSOFT Workshop on Program
       Analysis for Software and Tool Engineering
740 02 Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on
       Principles of Programming (POPL 2001)
780 00 |tSIGPLAN notices|x0362-1340|w(DLC)   83642027
       |w(OCoLC)1460457
907    .b18830286|b06-10-10|c11-28-94
907    .b18830286|b03-24-08|c11-28-94
944    RDA ENRICHED
944    MARS
947    03/04be cg|bmr|nurl change
947    rev cg 9/02 af www
947    rev. 4/97 lkh (580/711/740s); rev jch 12/99 af www
947    03/04tr hmc|bmr|ncomputer science move
```

View additional copies or search for a specific volume/copy

Location	**PASCAL Offsite QA76.7 .A8**
Holdings	**v.26 (1991) - v.36 (2001)**
	Request a PDF (UCB only)

LOCATION	CALL #	STATUS
PASCAL Offsite	QA76.7 .A8 v.26 no.10-12 1991	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.27 no.1-5 1992	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.27 no.6-8 1992	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.27 no.9-12 1992	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.28 no.1-6 1993	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.29 no.1-6 1994	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.28 no.7-9 1993	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.28 no.10-12 1993	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.29 no.7-9 1994	AVAILABLE
PASCAL Offsite	QA76.7 .A8 v.29 no.10-12 1994	AVAILABLE

Attachment 5b

OCLC 25073822	**No holdings in XXX - 348 other holdings**

Continuing Resources

						Rec Stat	c	**Entered**	19920109	**Replaced**	20190308220546.6
Type	a	**ELvl**		**Srce**	c	**GPub**		**Ctrl**		**Lang**	eng
BLvl	s	**Form**		**Conf**	0	**Freq**	m	**MRec**		**Ctry**	nyu
S/L	0	**Orig**		**EntW**		**Regl**	r	**Alph**	a		
Desc	a	**SrTp**	p	**Cont**		**DtSt**	c	**Dates**	1991 , 9999		

010 92659029
040 TJC $b eng $c TJC $d DLC $d NST $d WAU $d NST $d DLC $d IUL $d GUA $d NSD $d OCL $d LYU $d DLC $d CUS $d NLM $d WAU $d MYG $d DLC $d OCLCQ $d LVB $d OCLCQ $d DLC $d OCLCQ $d HUL $d IUL $d DLC $d OCLCF $d OCLCO $d GILDS $d OCLCO $d AU@
012 2 $i 9201 $l y
016 7 100941183 $2 DNLM
019 926214371 $a 990751744 $a 1012071622
022 0 1523-2867 $y 0362-1340 $2 1
037 $b Association for Computing Machinery, 11 W 42nd St., New York, NY 10036
042 pcc $a nsdp
050 00 QA76.7 $b .S54
082 00 001.64/24/05 $2 19
090 $b
210 0 ACM SIGPLAN not.
222 0 ACM SIGPLAN notices
245 00 ACM SIGPLAN notices : $b a monthly publication of the Special Interest Group on Programming Languages.
246 3 Association for Computing Machinery Special Interest Group on Programming Languages notices
246 3 Special Interest Group on Programming Languages notices
246 30 Notices
246 13 SIGPLAN notices
260 New York, NY : $b ACM Press, $c [1991]-
300 v. : $b ill. ; $c 28 cm.
310 Monthly
362 0 Vol. 26, no. 10 (Oct. 1991)-
500 Title from cover.
525 Some issues accompanied by supplements; some v. include special issues.
588 Latest issue consulted: Vol. 38, no. 2 suppl. (Feb. 2003).
650 0 Programming languages (Electronic computers) $v Periodicals.
650 6 Langages de programmation $v Périodiques.
650 7 Programming languages (Electronic computers) $2 fast $0 (OCoLC)fst01078704
655 7 Periodicals. $2 fast $0 (OCoLC)fst01411641
655 7 Periodicals. $2 lcgft
655 4 Electronic journals.
710 2 ACM Special Interest Group on Programming Languages.
776 08 $i Online version: $t ACM SIGPLAN notices (Online) $w (OCoLC)50318304
776 08 $i Online version: $t ACM SIGPLAN notices $w (OCoLC)603778396

776 08 $i Online version: $t ACM SIGPLAN notices $w (OCoLC)603778493
780 00 $t SIGPLAN notices $x 0362-1340 $w (DLC)   83642027 $w (OCoLC)1460457
850   AAP $a CoU $a DLC $a GU $a InU $a MdBJ $a MiEM $a NSyU $a OCU $a WaU
856 41 $u http://portal.acm.org/sigplan/newsletter
891 20 $9 853 $8 1 $a v. $b no. $u 12 $v r $i (year) $j (month) $w m
891 41 $9 863 $8 1.1 $a <27> $b <1> $i <1992> $j <01>.

Attachment 5c

LIBRARY OF CONGRESS AUTHORITIES

LC control no.: sh2010108608
LCCN Permalink: https://lccn.loc.gov/sh2010108608
HEADING: Programming languages (Electronic computers) Periodicals
000 00521cz a2200169n 450
001 8299720
005 20110729164023.0
008 100602|| anannbabn |n ana
010 __ **|a** sh2010108608
035 __ **|a** (DLC)470963
035 __ **|a** (DLC)sh2010108608
040 __ **|a** DLC **|b** eng **|c** DLC
150 __ **|a** Programming languages (Electronic computers) **|v** Periodicals
667 __ **|a** Record generated for validation purposes.
670 __ **|a** Work cat.: ACM letters on programming languages and systems, c1992-c1994
906 __ **|t** 8888 **|u** tc00 **|v** 0
953 __ **|a** tc00

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

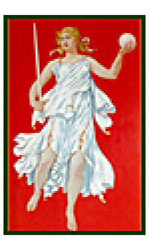

Library of Congress
URL: https://www.loc.gov/

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

Attachment 5d

Implementation of a database factory

AR Selvaraj, D Ghosh - ACM SIGPLAN Notices, 1997 - dl.acm.org

Object oriented software systems that utilize relational databases for data-store have to deal with the problem of interfacing to the relational data. This aspect of the software is more relevant to the solution domain rather than the problem domain ie, the business requirements of an application do not dictate that there be a mechanism that allows the application to store and retrieve relational data. Hence, de-coupling the application from the database and its interface is relevant from the perspective of portability that of the application …

Attachment 5e

Document details

SIGPLAN Notices (ACM Special Interest Group on Programming Languages)
Volume 32, Issue 6, June 1997, Pages 14-16

Implementation of a Database Factory (Article)

Selvaraj, A.R.[a,b,c], Ghosh, D.[d]

[a]College of Engineering, Guindy, India
[b]Anna University, Madras, India
[c]University of Texas, Austin, TX, United States

View additional affiliations

Abstract

View references (3)

Object oriented software systems that utilize relational databases for data-store have to deal with the problem of interfacing to the relational data. This aspect of the software is more relevant to the solution domain rather than the problem domain i.e., the business requirements of an application do not dictate that there be a mechanism that allows the application to store and retrieve relational data. Hence, de-coupling the application from the database and its interface is relevant from the perspective of portability of the application to other kinds of databases. For example, it is conceivable the application may be required to work with relational databases from different vendors. This article shows an adaptation of the Factory Method Pattern and the Abstract Factory Pattern [1] as a generic solution to the problem of de-coupling application code from the underlying database and its associated interface mechanisms.

SciVal Topic Prominence

Topic: Students | Teaching | cloud computing

Prominence percentile: 38.577

ISSN: 03621340
CODEN: SINOD
Source Type: Journal
Original language: English

Document Type: Article

References (3)

1 Gamma, E., Helm, R., Johnson, R., Vlissides, J.
(1995) *Design Patterns : Elements of Reusable Object-Oriented Software*. Cited 13227 times.
Addison Wesley, Reading, MA

2 Martin, R.C.
The Dependency Inversion Principle
(1996) *C++ Report*. Cited 42 times.
June

3 Meyer, B.
(1988) *Object Oriented Software Construction*. Cited 3353 times.
Prentice Hall

Selvaraj, A.R.; College of Engineering, India; email:asokanrs@worldnet.att.net

< Back to results | < Previous 118 of 171 Next >

^ Top of page

Deriving patterns from design principles
Pescio, C.
(1998) JOOP - Journal of Object-Oriented Programming

A class of Kalman filters for real-time image processing
Laplante, P.A. , Neill, C.J.
(2003) Proceedings of SPIE - The International Society for Optical Engineering

Improving object-oriented micro architectural design through knowledge systematization
Garzás, J. , Piattini, M.
(2005) Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)

View all related documents based on references

Find more related documents in Scopus based on:

Authors >

Attachment 6a

Library Catalog: Search for Books, Journals, Music and More


```
LEADER 00000caM  2200000Ia 4500
001    41208236
003    OCoLC
005    19991202132409.0
008    990419s1999    caua          001 0 eng d
020    1565924010
040    OKJ|cOKJ|dOCL|dIXA
049    IXAA
090    QA76.73.S65|bK55 1999
100 1  Kline, Kevin E.,|d1966-
245 10 Transact-SQL programming /|cKevin Kline, Lee Gould, Andrew
       Zanevsky.
250    1st ed.
260    Sebastopol, CA :|bO'Reilly & Associates,|c1999.
300    xxvi, 807 p. : ill. ;|c24 cm. +|e1 computer optical disc
       (4 3/4 in.)
500    "Microsoft SQL Server 6.5/7.0 & Sybase Adaptive Server
       11.5"--Cover.
500    Includes index (p.789-807).
505 2  Forward, Preface, pt. I. The basics: programming in
       Transact-SQL--pt. II. The building blocks: Transact-SQL
       language elements--pt. III. Functions and extensions--
       pt.IV. Programming Transact-SQL objects--pt. V.
       Performance tuning and optimization--pt. VI. Appendixes.
520    Provides detailed information about Transact-SQL
       programming and shows specific differences between the
       Microsoft and Sybase versions of the language.
630 00 SQL server.
650  0 SQL (Computer program language)
650  0 Query languages (Computer science)
650  0 Database management.
700 1  Gould, Lee.
700 1  Zanevsky, Andrew.
915    (TxU)10332312
915    41208236
```

Location	Call No.	Current Status
Library Storage Facility -- click Pick It Up	QA 76.73 S6774 K555 1999 TEXT	DUE 08-26-19
PCL Check Out Desk On Site Storage - ASK AT PCL CHECK OUT DESK or click Pick It Up	QA 76.73 S6774 K555 1999 CDROM	DUE 08-26-19

Attachment 6b

OCLC 41388930	**No holdings in XXX - 135 other holdings**										
Books						**Rec Stat**	c	**Entered**	20000526	**Replaced**	20190525034427.3
Type	a	**ELvl**		**Srce**		**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	cau
		Cont		**GPub**		**LitF**	0	**Indx**	1		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	1999 ,		

010 00267561
040 UKM $b eng $c DLC $d UKM $d OKJ $d BAKER $d BTCTA $d YDXCP $d ZWZ $d OCLCF $d OCLCQ $d OCLCO $d I8M $d OCLCQ $d VYT $d OCLCQ $d CPO $d PAU $d OCLCQ $d CWI $d OCLCQ $d ZQP
015 GB9935592 $2 bnb
019 41208236 $a 1049057789 $a 1080765076 $a 1102169587
020 1565924010
020 9781565924017
020 8173660816
020 9788173660818
042 lccopycat
050 00 QA76.73.S67 $b K555 1999
082 00 005.75/6 $2 21
090 $b
100 1 Kline, Kevin E., $d 1966-
245 10 Transact-SQL programming / $c Kevin Kline, Lee Gould, and Andrew Zanevsky.
260 Sebastopol : $b O'Reilly, $c 1999.
300 xxvi, 807 pages : $b illustrations ; $c 24 cm + $e 1 computer optical disc (4 3/4 in.)
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
500 Includes index.
505 2 Forward, Preface, pt. I. The basics: programming in Transact-SQL--pt. II. The building blocks: Transact-SQL language elements--pt. III. Functions and extensions--pt. IV. Programming Transact-SQL objects--pt. V. Performance tuning and optimization--pt. VI. Appendixes.
520 Provides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybase versions of the language.
650 0 SQL (Computer program language)
650 7 SQL (Computer program language) $2 fast $0 (OCoLC)fst01102869
650 7 SQL (langage de programmation) $2 ram
650 7 Bases de données $x Gestion. $2 ram
700 1 Gould, Lee.
700 1 Zanevsky, Andrew.
856 42 $3 Contributor biographical information $u http://catdir.loc.gov/catdir/enhancements/fy1004/00267561-b.html
856 42 $3 Publisher description $u http://catdir.loc.gov/catdir/enhancements/fy0715/00267561-d.html

Attachment 6c

Help | Search | History | Titles | Start Over

Public Catalog

Copyright Catalog (1978 to present)
Search Request: Left Anchored Title = Transact-SQL Programming
Search Results: Displaying 4 of 4 entries

previous next

Labeled View

Transact-SQL programming / Kevin Kline, Lee Gould, and Andrew Zanevsky.

Type of Work: Text
Registration Number / Date: TX0005106977 / 1999-12-09
Title: Transact-SQL programming / Kevin Kline, Lee Gould, and Andrew Zanevsky.
Description: 807 p.
Copyright Claimant: O'Reilly & Associates
Date of Creation: 1999
Date of Publication: 1999-03-31
Names: Kline, Kevin
Gould, Lee
Zanevsky, Andrew
O'Reilly & Associates

previous next

Save, Print and Email (Help Page)	
Select Download Format Full Record	Format for Print/Save
Enter your email address:	Email

Help Search History Titles Start Over

Contact Us | Request Copies | Get a Search Estimate | Frequently Asked Questions (FAQs) about Copyright | Copyright Office Home Page | Library of Congress Home Page

Attachment 6d

```
Leader 01695nam a2200385 a 4500
001    991002238889702766
005    20170314105507.0
008    990514s1999    cc a          001 0 eng d
015    ##$aGB99-35592
019    ##$a41208236
020    ##$a1565924010$0(uri) http://www.isbnsearch.org/isbn/1565924010
035    ##$a(CoDU)b25440482-01uode_inst
035    ##$a(OCoLC)ocm41388930 $9ExL
040    ##$aUKM$cUKM$dOKJ
049    ##$aDVPP
090    ##$aQA76.73.S65$bK55 1999
100    1#$aKline, Kevin E.,$d1966-$0(uri) http://id.loc.gov/authorities/names/no97002232$0(uri) http://viaf.org/viaf/sourceID/LC|no9700
245    10$aTransact-SQL programming /$cKevin Kline, Lee Gould, and Andrew Zanevsky.
260    ##$aBeijing ;$aCambridge :$bO'Reilly,$c1999.
300    ##$axxvi, 807 p. :$bill. ;$c24 cm. +$e1 computer optical disk (4 3/4 in.)
500    ##$aIncludes index.
500    ##$a"Please verify all system requirements for disks before using them.  The library does not offer technical support."
520    ##$aProvides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybase
590    ##$aAccompanying CD-ROM is requestable from the Hampden Center.
630    00$aSQL server.$0(uri) http://id.loc.gov/authorities/names/n90684343$0(uri) http://viaf.org/viaf/sourceID/LC|n90684343$2LCNAMES$
630    00$aMicrosoft SQL server$2LCNAMES$041-LIBRARY_OF_CONGRESS-n  90684343$9N
650    #0$aSQL (Computer program language)$0(uri) http://id.loc.gov/authorities/subjects/sh86006628$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86
650    #0$aQuery languages (Computer science)$0(uri) http://id.loc.gov/authorities/subjects/sh96000160$2LCSH$041-LIBRARY_OF_CONGRESS-sh
650    #0$aDatabase management.$0(uri) http://id.loc.gov/authorities/subjects/sh85035848$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9Y
650    #0$wnne$aData base management$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aData services (Database management)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aDatabase management services$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aDBMS (Computer science)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aGeneralized data management systems$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aServices, Database management$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aSystems, Database management$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$aSystems, Generalized database management$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9N
650    #0$wg$aElectronic data processing$2LCSH$041-LIBRARY_OF_CONGRESS-sh 85035848$9R
650    #0$aQLSs (Computer science)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 96000160$9N
650    #0$aQuery language systems (Computer science)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 96000160$9N
650    #0$wg$aProgramming languages (Electronic computers)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 96000160$9R
650    #0$aStructured Query Language (Computer program language)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9N
650    #0$wg$aDeclarative programming languages$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9R
650    #0$wg$aQuery languages (Computer science)$2LCSH$041-LIBRARY_OF_CONGRESS-sh 86006628$9R
700    1#$aGould, Lee.$0(uri) http://id.loc.gov/authorities/names/nb99043490$0(uri) http://viaf.org/viaf/sourceID/LC|nb99043490$2LCNAME
700    1#$aZanevsky, Andrew.$0(uri) http://id.loc.gov/authorities/names/nb99043489$0(uri) http://viaf.org/viaf/sourceID/LC|nb99043489$2
907    ##$a.b25440482$b01-25-16$c02-16-00
919    ##$aMARCIVE$dSpring 2011
925    ##$aQA76.73.S65$bK55 1999$g1$h10$iU186009905508$j0$k0$lhcstk$o-$p$0.00$ru$s-$t0$u7$v15$w0$x0$y.i25979127$z03-31-00
925    ##$aQA76.73.S65$bK55 1999$cCD-ROM$g1$h0$iU186011451523$j0$k0$lhccdr$o-$p$0.00$ru$s-$t0$u2$v0$w0$x0$y.i25979140$z03-31-00
998    ##$ahc$b03-31-00$cm$da$e-$feng$gcc $h0$i0
AVA    ##$0991002238889702766$822639692740002766$a01UODE_INST$bh$cCDs/CD-ROMs$dQA76.73.S65 K55 1999$eavailable$f1$g0$iMAINCAMPUS$jhccdr
AVA    ##$0991002238889702766$a01UODE_INST$bp$cProspector$dQA76.73.S65 K55 1999$eunavailable$f1$g1$iMAINCAMPUS$jpoill$k0$p2$qMain
INST   ##$a01UODE_INST
INT    ##$aP
PLK    ##$aReplaces.$b991007048269702766
PLK    ##$aAdditional Form, Replaces.$b991041475088202766
```

```
Leader 01930cam a2200457Ma 4500
001    991041475088202766
005    20170220063037.4
006    m     o  d
007    cr cn|||||||||
008    001201s1999    caua    o     001 0 eng d
020    ##$a1565929799$q(electronic bk.)
020    ##$a9781565929791$q(electronic bk.)
020    ##$z1565924010
035    ##$a(OCoLC)45729649$z(OCoLC)532733300$z(OCoLC)758975574$z(OCoLC)970797262
035    ##$a(OCoLC)ocm45729649
040    ##$aN$T$beng$epn$cN$T$dOCL$dOCLCQ$dOCLCG$dOCLCQ$dTUU$dOCLCQ$dTNF$dOCLCQ$dZCU$dOCLCO$dOCLCF$dOCLCQ$dN$T$dOCLCQ$dMWM$dOCLCQ
049    ##$aDVPP
050    #4$aQA76.73.S67$bK555 1999eb
100    1#$aKline, Kevin E.,$d1966-
245    10$aTransact-SQL programming /$cKevin Kline, Lee Gould, and Andrew Zanevsky.
260    ##$aSebastopol :$bO'Reilly,$c1999.
300    ##$a1 online resource (xxvi, 807 pages) :$billustrations
336    ##$atext$btxt$2rdacontent
337    ##$acomputer$bc$2rdamedia
338    ##$aonline resource$bcr$2rdacarrier
500    ##$aIncludes index.
```

```
520    ##$aProvides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybase
588    0#$aPrint version record.
650    #0$aSQL (Computer program language)
650    #7$aCOMPUTERS$xProgramming Languages$xSQL.$2bisacsh
650    #7$aSQL (Computer program language)$2fast$0(OCoLC)fst01102869
700    1#$aGould, Lee.
700    1#$aZanevsky, Andrew.
776    08$iPrint version:$aKline, Kevin E., 1966-$tTransact-SQL programming.$dSebastopol : O'Reilly, 1999$z1565924010$w(DLC)   002(
856    40$uhttp://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&AN=24587
INST   ##$a01UODE_INST
INT    ##$aE
PLK    ##$aAdditional Form, Replaced By.$b991002238889702766

-------------------------------------------------------------------------------------------------------------

Leader 02155nam a2200433 a 4500
001    991007048269702766
005    20010619180020.0
006    m        u
007    cr cn|||||||||
008    010619s1999    caua    s     001 0 eng d
010    ##$z00267561
020    ##$a1565929799 (electronic bk.)
020    ##$z1565924010
035    ##$a(CoDU)b26945654-01uode_inst
035    ##$anL24587$9ExL
035    ##$a.b33663221$bweb$cg
035    ##$a.b33663221$bweb$cg
035    ##$a.b33663221$bweb$cg
040    ##$aN$T$cN$T
049    ##$aDVPP
050    14$aQA76.73.S67$bK555 1999eb
100    1#$aKline, Kevin E.,$d1966-
245    10$aTransact-SQL programming$h[electronic resource] /$cKevin Kline, Lee Gould, and Andrew Zanevsky.
260    ##$aSebastopol, Calif. :$bO'Reilly,$c1999.
300    ##$axxvi, 807 p. :$bill. ;$c24 cm. +$e1 computer optical disc (4 3/4 in.)
504    ##$aIncludes index.
520    ##$aProvides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybase
533    ##$aElectronic reproduction.$bBoulder, Colo. :$cNetLibrary,$d2000.$nAvailable via World Wide Web.$nAccess may be limited to NetL
650    #0$aSQL (Computer program language)
655    #7$aElectronic books.$2local.
700    1#$aGould, Lee.
700    1#$aZanevsky, Andrew.
710    2#$aNetLibrary, Inc.
776    0#$cOriginal$w(DLC) 00267561.
856    4#$uhttp://du.idm.oclc.org/login?url=http://www.netlibrary.com/urlapi.asp?action=summary&v=1&bookid=24587$zView electronic book.
856    4#$uhttp://du.idm.oclc.org/login?url=http://www.netlibrary.com/urlapi.asp?action=summary&v=1&bookid=24587$zView electronic book.
907    ##$a.b26945654$b04-11-16$c06-06-02
919    ##$aMARCIVE$dSpring 2011
919    ##$aBackstage$dpre-2012
945    ##$aQA76.73.S67$bK555 1999eb$g1$h90$j0$k0$linter$o-$p$0.00$q $r $sj$t0$u0$v0$w0$x0$y.i27957317$z06-06-02
945    ##$aQA76.73.S67$bK555 1999eb$g1$h0$j0$k0$lwnint$o-$p$0.00$q $r $sj$t130$u0$v0$w0$x0$y.i61128417$z01-11-12
998    ##$ain$awnint$b06-06-02$cm$da$e-$feng$gcau$h0$i2
998    ##$awnint$b01-01-02$ct$d@$ez$feng$gcau$h0
INST   ##$a01UODE_INST
INT    ##$aE
PLK    ##$aReplaced By.$b991002238889702766
```

Attachment 6e

```
leader  01722cam a2200397 a 4500
001     991004792899704436
008     000526s1999    caua          001 0 eng d
010     ##$a   00267561
015     ##$aGB99-35592
020     ##$a1565924010
035     ##$a534296-o1cmu_inst
035     ##$a(Sirsi) o41388930
040     ##$aUKM $cDLC $dUKM $dOKJ
042     ##$alccopycat
050     00$aQA76.73.S67 $bK555 1999
082     00$a005.75/6 $221
100     1#$aKline, Kevin E., $d1966-
245     10$aTransact-SQL programming / $cKevin Kline, Lee Gould, and Andrew Zanevsky.
260     ##$aSebastopol CA : $bO'Reilly, $c1999.
300     ##$axxvi, 807 p. : $bill. ; $c24 cm. + $e1 computer optical disc (4 3/4 in.)
504     ##$aIncludes index.
505     2#$aForward, Preface, pt. I. The basics: programming in Transact-SQL--pt. II. The building blocks: Transact-SQL language elemen
520     ##$aProvides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybas
596     ##$a9
650     #0$aSQL (Computer program language)
700     1#$aGould, Lee.
700     1#$aZanevsky, Andrew.
902     ##$a534296
903     ##$ao41388930
916     ##$a20010515
917     ##$a19991104
918     ##$aACQSTUDENT
919     ##$a20150629
920     ##$aCATASSIST
948     ##$hHELD BY PMC - 109 OTHER HOLDINGS
999     ##$aQA76.73 .S67 K555 1999 DISC $wLC $c1 $iF090191 $d8/24/2007 $e8/24/2007 $lBY-REQUEST $mOFFSITE $n5 $rY $sY $tCD-ROM $u5/15/2
999     ##$aQA76.73 .S67 K555 1999 $wLC $c1 $iF090190 $d11/14/2008 $e8/18/2008 $lBY-REQUEST $mOFFSITE $n11 $p$37.31 $q2 $rY $sY $tBOOK
```

Attachment 6f

Roger

The UC San Diego Library Catalog


```
LEADER 00000cam a2200493 a 4500
001    41388930
005    20180501111136.0
008    000526s1999    caua          001 0 eng
010    00267561
015    GB9935592|2bnb
019    41208236
020    1565924010
020    9781565924017
035    (OCoLC)41388930|z(OCoLC)41208236
040    UKM|beng|cDLC|dUKM|dOKJ|dBAKER|dBTCTA|dYDXCP|dZWZ|dOCLCF
       |dOCLCQ|dOCLCO|dI8M|dOCLCQ|dVYT|dOCLCQ|dCUS|dUtOrBLW
042    lccopycat
050 00 QA76.73.S67|bK555 1999
082 00 005.75/6|221
100 1  Kline, Kevin E.,|d1966-
245 10 Transact-SQL programming /|cKevin Kline, Lee Gould, and
       Andrew Zanevsky
260    Sebastopol :|bO'Reilly,|c1999
300    xxvi, 807 pages :|billustrations ;|c24 cm +|e1 computer
       optical disc (4 3/4 in.)
336    text|btxt|2rdacontent
337    unmediated|bn|2rdamedia
338    volume|bnc|2rdacarrier
500    Includes index
505 2  Forward, Preface, pt. I. The basics: programming in
       Transact-SQL--pt. II. The building blocks: Transact-SQL
       language elements--pt. III. Functions and extensions--pt.
       IV. Programming Transact-SQL objects--pt. V. Performance
       tuning and optimization--pt. VI. Appendixes
520    Provides detailed information about Transact-SQL
       programming and shows specific differences between the
       Microsoft and Sybase versions of the language
650  0 SQL (Computer program language)
650  7 SQL (Computer program language)|2fast|0(OCoLC)fst01102869
700 1  Gould, Lee
700 1  Zanevsky, Andrew
910    RDA ENRICHED
910    MARS
938    Baker & Taylor|bBKTY|c49.95|d37.46|i1565924010|n0003234940
       |sactive
938    Baker and Taylor|bBTCP|n00267561
938    YBP Library Services|bYANK|n1579056
994    C0|bCUS
```

Location	Call Number	Status
Books, Geisel Floor1 East	QA76.73.S67 K555 1999	AVAILABLE

RSS Feeds
New Books, Videos, etc.

Recommend Purchases
Suggest a title for purchase

Help
View Roger help topics

Attachment 6g

LIBRARY CATALOG

Search Navigation 2 of 6

BOOK

Transact-SQL programming

Full Record MARC Tags

000		01726cam a2200337 a 4500
001		12022936
005		20100212084643.0
008		000526s1999 caua 001 0 eng d
906	__	|**a** 7 |**b** cbc |**c** copycat |**d** 2 |**e** ncip |**f** 20 |**g** y-gencatlg
925	0_	|**a** acquire |**b** 2 shelf copies |**x** policy default
955	__	|**a** to ASCD pb08 05-26-00; jf00 06-01-00; jf03 09-14-00 ; jf11 to sl 10-2-00; jf25 to Dewey 10-05-00; aa19 10-16-2000; copy 2 added jf16 to BCCD 12-04-00
010	__	|**a** 00267561
015	__	|**a** GB99-35592
020	__	|**a** 1565924010
035	__	|**a** (OCoLC)ocm41388930
040	__	|**a** UKM |**c** UKM |**d** OKJ |**d** DLC
042	__	|**a** lccopycat
050	00	|**a** QA76.73.S67 |**b** K555 1999
082	00	|**a** 005.75/6 |**2** 21
100	1_	|**a** Kline, Kevin E., |**d** 1966-
245	10	|**a** Transact-SQL programming / |**c** Kevin Kline, Lee Gould, and Andrew Zanevsky.
260	__	|**a** Sebastopol : |**b** O'Reilly, |**c** 1999.
300	__	|**a** xxvi, 807 p. : |**b** ill. ; |**c** 24 cm. + |**e** 1 computer optical disc (4 3/4 in.)
504	__	|**a** Includes index.
505	2_	|**a** Forward, Preface, pt. I. The basics: programming in Transact-SQL--pt. II. The building blocks: Transact-SQL language elements--pt. III. Functions and extensions--pt.IV. Programming Transact-SQL objects--pt. V. Performance tuning and optimization--pt. VI. Appendixes.

520	__	|**a** Provides detailed information about Transact-SQL programming and shows specific differences between the Microsoft and Sybase versions of the language.
650	_0	|**a** SQL (Computer program language)
700	1_	|**a** Gould, Lee.
700	1_	|**a** Zanevsky, Andrew.
856	42	|**3** Publisher description |**u** http://www.loc.gov/catdir/enhancements/fy0715/00267561-d.html
856	42	|**3** Contributor biographical information |**u** http://www.loc.gov/catdir/enhancements/fy1004/00267561-b.html

Request this Item

LC Find It

Item Availability

CALL NUMBER	QA76.73.S67 K555 1999 FT MEADE Copy 2
Request in	Machine Readable Collections - STORED OFFSITE
Status	Not Charged

CALL NUMBER	QA76.73.S67 K555 1999 Copy 1
Request in	See Reference Staff. By Appt in Jefferson Main RR (MRC)
Status	Not Charged

Attachment 6h

Transact-SQL Programming: Covers Microsoft SQL Server 6.5/7.0 and Sybase Adaptive Server 11.5
K Kline, L Gould, A Zanevsky - 1999 - books.google.com
Transact-SQL is a procedural language used on both Microsoft SQL Server and Sybase SQL Server systems. It is a full-featured programming language that dramatically extends the power of SQL (Structured Query Language). The language provides programmers with a …
Cited by 18 Related articles All 4 versions

Microsoft SQL Server 2005 Stored Procedure **Programming** in T-SQL &. NET
D Sunderic - 2006 - dl.acm.org
… Table of contents Chapter 1: The SQL Server 2005 Environment and Tools Chapter 2: Stored Procedure Design Concepts Chapter 3: Basic **Transact-SQL Programming** Constructs Chapter 4: Functions Chapter 5: Composite **Transact-SQL** Constructs: Batches, Scripts, and …
Cited by 13 Related articles All 3 versions

Advanced **Transact-SQL** for SQL Server 2000
I Ben-Gan, T Moreau - 2008 - books.google.com
… 197 Chapter 7 Writing Code in **Transact-SQL** ..... 233 Chapter 8 Views … 87 Using the SELECT INTO Statement ..... 90 The Bulk Copy **Program**-bcp …
Cited by 19 Related articles All 6 versions

SQL Server 2000: the complete reference
J Shapiro - 2002 - dl.acm.org
… Learn to **program** SQL Server with **Transact-SQL**,handle multi-dimensional data processing with the Analysis Services,use Data Transformation Services,implement data mining and data warehousing techniques,back up and restore databases,and manage security,performance …
Cited by 22 Related articles All 2 versions

Microsoft (R) SQL Server 2005 Unleashed
R Rankins, P Bertucci, C Gallelli, AT Silverstein - 2006 - dl.acm.org
… Chapters on the CD42 Managing Linked and Remote Servers43 Configuring, Tuning, and Optimizing SQL Server Options44 Administering Very Large SQL Server Databases45 SQL Server Disaster Recovery Planning46 **Transact-SQL Programming** Guidelines, Tips, and …
Cited by 18 Related articles All 2 versions

Sybase SQL Server 11: Unleashed
R Rankins, JR Garbus, BW Mcewan, D Solomon - 1996 - fntic.univ-ouargla.dz
… your command, you need a comprehensive book to detail the best ways to **program**, tune, and administer your databases. SYBASE SQL Server 11 Unleashed is that book. Between its covers you'll find information on installation, **Transact-SQL programming**, performance tuning …
Cited by 13 Related articles All 3 versions

Programming Microsoft SQL Server 2000 with XML
G Malcolm, M Graeme - 2001 - dl.acm.org
Programming Microsoft SQL Server 2000 with XML … Publication: Cover Image. · Book, **Programming** Microsoft SQL Server 2000 with XML. Microsoft Press Redmond, WA, USA ©2001 ISBN:0735613699. 2001 Book. Bibliometrics …
Cited by 15 Related articles All 3 versions

SQL Server the Complete Reference
G Coffman - 1998 - dl.acm.org
Cited by 9 Related articles

SQL Server 2000 Stored Procedure & XML **Programming**
D Sunderic - 2003 - McGraw-Hill/Osborne
Cited by 8 Related articles All 2 versions

The Guru's Guide to SQL Server Stored Procedures, Xml, and HTML with Cdrom
K Henderson, R Foreword By-Soukup - 2001 - dl.acm.org
… In their book, The Pragmatic **Programmer** (Addison-Wesley, 1999), Andrew Hunt and Dave … to be better **programmers** should learn at least one new **programming** language per … If you want to master **Transact**-**SQL** stored procedure **programming**, you should master **programming** …

Cited by 9 Related articles All 6 versions

Related searches

stored procedure programming **microsoft** sql **server**

Attachment 7a

Start Over | Regular Display | Return to Browse | Limit/Sort Search

(Search History)

```
LEADER 00000pam  2200301 a 4500
001    17650642
005    19890216113659.0
008    880303s1988    nju           001 0 eng
010    88005934
020    0131103709
020    0131103628 (pbk.)
040    DLC|cDLC|dCOR
049    CORN
082 0  005.13/3|219
090    QA76.73 .C15 K47 1988
100 1  Kernighan, Brian W.
245 14 The C programming language /|cBrian W. Kernighan, Dennis
       M. Ritchie.
250    2nd ed.
260    Englewood Cliffs, N.J. :|bPrentice Hall,|cc1988.
300    xii, 272 p. ;|c25 cm
500    Includes index.
590    c.1
650  0 C (Computer program language)
700 1  Ritchie, Dennis M.,|d1941-2011.
907    .b11062940|b05-02-12|c10-21-03
910 99 CARL0001878518
998    dm|b02-16-89|cm|da|e-|feng|gnju|h4|i2
```

Location	Call No.	Status
DML 3rd Floor Books	QA76.73 .C15 K47 1988	Checked Out

Attachment 7b

OCLC 17650642 **No holdings in XXX - 1375 other holdings**											
Books						**Rec Stat**	c	**Entered**	19880303	**Replaced**	20190326164422.5
Type	a	**ELvl**		**Srce**		**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	nju
		Cont	b	**GPub**		**LitF**	0	**Indx**	1		
Desc	i	**Ills**	a	**Fest**	0	**DtSt**	t	**Dates**	1988 ,	1988	

007 t $b a
010 88005934
040 DLC $b eng $e rda $c DLC $d UKM $d WQM $d UBA $d NOR $d MUQ $d BAKER $d NLGGC $d BTCTA $d YDXCP $d AU@ $d BNO $d UKV3G $d HEBIS $d U5D $d OKN $d BDX $d GBVCP $d OCLCF $d OCLCQ $d CN8BJ $d SE4M6 $d MEAUC $d OCLCQ $d LTU $d XFF $d LVT $d VYT $d OCLCO $d ZXX $d OCLCO $d LIP $d NJK $d THHCU $d SFR $d THHCU $d CSA $d OCLCQ $d CSJ $d TFW $d OCLCQ $d OCLCO $d WY@ $d OCLCO $d DHA $d UKSAE $d UKBTC $d JVQ $d NJR $d OCLCO $d BUF $d BGAUB $d OKX $d UAB $d CUY $d IOG $d UX0 $d GILDS $d APUMS $d DAC $d NUS $d OSU $d OCLCO $d IOK $d NIU $d T9K $d OCLCQ $d OCLCO $d UNITY $d CNO $d OCLCO $d OKR $d XOJ $d CCH $d GBS $d MNB $d NLHHG $d CAI $d BRL $d AGM $d VNS $d UKWOH $d WOM $d UKOBU $d LHU $d CEF $d TKN $d MUO $d NLMVD $d CRU $d CPS $d IAD $d PAU $d NAM $d XIM $d QE2 $d SNN $d NTG $d INT $d URS $d OCLCQ $d OCLCO $d BGU $d OCLCO $d E@W $d JG0 $d OCLCO $d DXU $d XC$$ $d UWO $d UKUOY $d YOU $d NZ1 $d HV6 $d BID $d CWI $d HUELT $d ERD $d IOB $d CAUOI $d LGG $d UKNRU $d LENOT $d U3G $d UWW $d FIE $d VP@
015 GB8863031 $2 bnb
015 GB9119224 $2 bnb
016 7 013-11037 $2 Uk
019 18815644 $a 59107357 $a 775886256 $a 788343545 $a 1018297095
020 0131103709
020 9780131103702
020 0131103628 $q (paperback)
020 9780131103627 $q (paperback)
020 0131158171
020 9780131158177
050 00 QA76.73.C15 $b K47 1988
060 4 QA 76.73.C15 K39c 1988
080 681.3.06:800.92
080 681.3.061
080 681.3.04C
080 004.43
080 681.3.06
080 681.3.04
080 518.579.1.K31
080 519.682.8
080 800.92
082 00 005.13/3 $2 19
083 0 Digital computer systems. Programming languages: C language

084 54.53 $2 bcl
084 D.3 $2 acmccs
084 D.3.2 $2 acmccs
084 PN 46(C10) $2 blsrissc
090 $b
100 1 Kernighan, Brian W., $e author.
245 14 The C programming language / $c Brian W. Kernighan, Dennis M. Ritchie.
250 Second edition.
264 1 Englewood Cliffs, N.J. : $b Prentice Hall, $c [1988]
264 4 $c ©1988
300 xii, 272 pages : $b illustrations ; $c 25 cm.
336 text $b txt $2 rdacontent
337 unmediated $b n $2 rdamedia
338 volume $b nc $2 rdacarrier
490 1 Prentice-Hall software series
500 Includes index.
505 00 $g ch. 1. $t A tutorial introduction -- $g ch. 2. $t Types, operators, and expressions -- $g ch. 3. $t Control flow -- $g ch. 4. $t Functions and program structure -- $g ch. 5. $t Pointers and arrays -- $g ch. 6. $t Structures -- $g ch. 7. $t Input and output -- $g ch. 8. $t The UNIX system interface.
520 Introduces the features of the C programming language, discusses data types, variables, operators, control flow, functions, pointers, arrays, and structures, and looks at the UNIX system interface.
650 0 C (Computer program language)
650 6 C (Langage de programmation)
650 7 C (Computer program language) $2 fast $0 (OCoLC)fst00843279
650 17 C (programmeertaal) $2 gtt
650 7 C $g Programmiersprache $2 gnd $0 (DE-588)4113195-2
650 7 Linguagem De Programacao (Outras) $2 larpcal
650 7 Computers. $2 nasat
650 7 C++ (programming language) $2 nasat
650 7 C. $2 tekord
650 4 C (Computer program language)
650 2 Computers.
650 2 Software.
653 Programming languages
655 7 Lærebøker. $2 humord
655 7 Programmering. $2 humord
700 1 Ritchie, Dennis M., $d 1941-2011, $e author.
776 08 $i Online version: $a Kernighan, Brian W. $t C programming language. $b 2nd ed. $d Englewood Cliffs, N.J. : Prentice Hall, ©1988 $w (OCoLC)573196545
776 08 $i Online version: $a Kernighan, Brian W. $t C programming language. $b 2nd ed. $d Englewood Cliffs, N.J. : Prentice Hall, ©1988 $w (OCoLC)607895187
830 0 Prentice-Hall software series.
856 41 $3 Table of contents $u http://www.gbv.de/dms/hebis-darmstadt/toc/10657762.pdf
856 42 $z Additional Information at Google Books $u http://books.google.com/books?isbn=0131103709

Attachment 7c

LIBRARY OF CONGRESS AUTHORITIES

LC control no.: sh 85018532
LCCN Permalink: https://lccn.loc.gov/sh85018532
HEADING: C (Computer program language)
000 00499cz a2200181n 450
001 4671169
005 20120322080115.0
008 860211|| anannbabn |a ana
010 __ **|a** sh 85018532
035 __ **|a** (DLC)sh 85018532
035 __ **|a** (DLC)413338
040 __ **|a** DLC **|c** DLC **|d** DLC
053 _0 **|a** QA76.73.C15
150 __ **|a** C (Computer program language)
550 __ **|w** g **|a** Programming languages (Electronic computers)
550 __ **|a** Objective-C (Computer program language)
906 __ **|t** 0851 **|u** tc09 **|v** 1
953 __ **|a** xx00 **|b** tc09

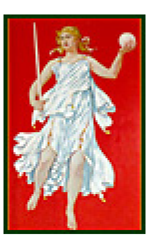

Library of Congress
URL: https://www.loc.gov/

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

Attachment 7d

Help | Search | History | Titles | Start Over

Public Catalog

Copyright Catalog (1978 to present)
Search Request: Left Anchored Title = C Programming Language
Search Results: Displaying 4 of 14 entries

previous | next

Labeled View

The C programming language / Brian W. Kernighan, Dennis M. Ritchie.

Type of Work:	Text
Registration Number / Date:	TX0002341624 / 1988-06-17
Title:	The C programming language / Brian W. Kernighan, Dennis M. Ritchie.
Edition:	2nd ed.
Imprint:	Englewood Cliffs, N. J. : Prentice-Hall, c1988.
Description:	272 p.
Copyright Claimant:	Bell Telephone Laboratories, Inc.
Date of Creation:	1987
Date of Publication:	1988-03-10
Previous Registration:	Prev. reg. 1978, TX 30-443.
Basis of Claim:	New Matter: "revised and new material added."
Names:	Kernighan, Brian W. Ritchie, Dennis M. Bell Telephone Laboratories, Inc.

previous | next

Save, Print and Email (Help Page)	
Select Download Format	Full Record ▼ Format for Print/Save
Enter your email address:	Email

Help Search History Titles Start Over

Attachment 7e

Library Catalog: Search for Books, Journals, Music and More


```
LEADER 00000pam  2200000 a 4500
001    17650642
005    19881108210939.0
008    880303s1988    nju           001 0 eng
010       88005934
020    0131103709
020    0131103628 (pbk.)
040    DLC|cDLC|dIXA
049    IXAA
050 0  QA76.73.C15|bK47 1988
082 0  005.13/3|219
100 1  Kernighan, Brian W.
245 14 The C programming language /|cBrian W. Kernighan, Dennis
       M. Ritchie.
250    2nd ed.
260    Englewood Cliffs, N.J. :|bPrentice Hall,|cc1988.
300    xii, 272 p. ;|c25 cm.
500    Includes index.
650  0 C (Computer program language)
700 1  Ritchie, Dennis M.,|d1941-2011.
915    (TxU)2831283
915    17650642
```

Location	Call No.	Current Status
PCL Stacks	QA 76.73 C15 K47 1988 c.3	DUE 09-08-19
PCL Stacks	QA 76.73 C15 K47 1988	DUE 09-08-19
PCL Stacks	QA 76.73 C15 K47 1988 c.4	DUE 09-08-19
PCL Stacks	QA 76.73 C15 K47 1988 c.5	AVAILABLE
PCL Stacks	QA 76.73 C15 K47 1988 c.8	BEING SEARCHED
PCL Stacks	QA 76.73 C15 K47 1988 c.9	DUE 10-18-19
PCL Stacks	QA 76.73 C15 K47 1988	AVAILABLE
PCL Stacks	QA 76.73 C15 K47 1988 c.11	DUE 09-08-19

Attachment 7f

```
LEADER 00000pam a2200349 a 4500
001    17650642
008    880303s1988    nju           001 0 eng
010    88005934
020    0131103709
020    0131103628 (pbk.)
040    DLC|cDLC|dCOD
049    CODI
050 0  QA76.73.C15|bK47 1988
050 0  QA76.73.C15|bK47 1988
100 1  Kernighan, Brian W
245 14 The C programming language /|cBrian W. Kernighan, Dennis
       M. Ritchie
250    2nd ed
260    Englewood Cliffs, N.J. :|bPrentice Hall,|c©1988
300    xii, 272 pages ;|c25 cm
336    text|btxt|2rdacontent
337    unmediated|bn|2rdamedia
338    volume|bnc|2rdacarrier
500    Includes index
650  0 C++ (Computer program language)
700 1  Ritchie, Dennis M.,|d1941-2011
907    .b17591120|b09-24-13|c11-25-94
907    .b17591120|b05-10-07|c11-25-94
944    RDA ENRICHED
944    MARS
```

LOCATION	CALL #	STATUS
Engineering, Math, Physics Library--Stacks	QA76.73.C15 K47 1988 c.2	AVAILABLE
Engineering, Math, Physics Library--Stacks	QA76.73.C15 K47 1988 c.3	AVAILABLE
Engineering, Math, Physics Library--Stacks	QA76.73.C15 K47 1988 c.4	AVAILABLE

Attachment 7g

```
Leader 04434cam a2200901 i 4500
001    991009232479702901
005    20190327061034.9
007    ta
008    880303t19881988njua     b    001 0 eng
010    ##$a   88005934
015    ##$aGB8863031$2bnb
015    ##$aGB9119224$2bnb
016    7#$a013-11037$2Uk
019    ##$a18815644$a59107357$a775886256$a788343545$a1018297095
020    ##$a0131103709
020    ##$a9780131103702
020    ##$a0131103628$q(paperback)
020    ##$a9780131103627$q(paperback)
020    ##$a0131158171
020    ##$a9780131158177
035    ##$aocm17650642-01cals_network
035    ##$a(OCoLC)17650642$z(OCoLC)18815644$z(OCoLC)59107357$z(OCoLC)775886256$z(OCoLC)788343545$z(OCoLC)1018297095
035    ##$a(OCoLC)ocm17650642
040    ##$aDLC$beng$erda$cDLC$dUKM$dWQM$dUBA$dNOR$dMUQ$dBAKER$dNLGGC$dBTCTA$dYDXCP$dAU@$dBNO$dUKV3G$dHEBIS$dU5D$dOKN$dBDX$dGBVCP$dOCLCF
050    00$aQA76.73.C15$bK47 1988
060    #4$aQA 76.73.C15 K39c 1988
066    ##$cThai
080    ##$a681.3.06:800.92
080    ##$a681.3.061
080    ##$a681.3.04C
080    ##$a004.43
080    ##$a681.3.06
080    ##$a681.3.04
080    ##$a518.579.1.K31
080    ##$a519.682.8
080    ##$a800.92
082    00$a005.13/3$219
083    0#$aDigital computer systems. Programming languages: C language
084    ##$a54.53$2bcl
084    ##$aD.3$2acmccs
084    ##$aD.3.2$2acmccs
084    ##$aPN 46(C10)$2blsrissc
100    1#$aKernighan, Brian W.,$eauthor.
245    14$aThe C programming language /$cBrian W. Kernighan, Dennis M. Ritchie.
250    ##$aSecond edition.
264    #1$aEnglewood Cliffs, N.J. :$bPrentice Hall,$c[1988]
264    #4$c©1988
300    ##$axii, 272 pages :$billustrations ;$c25 cm.
336    ##$atext$btxt$2rdacontent
337    ##$aunmediated$bn$2rdamedia
338    ##$avolume$bnc$2rdacarrier
490    1#$aPrentice-Hall software series
500    ##$aIncludes index.
505    00$gch. 1.$tA tutorial introduction --$gch. 2.$tTypes, operators, and expressions --$gch. 3.$tControl flow --$gch. 4.$tFunctions
520    ##$aIntroduces the features of the C programming language, discusses data types, variables, operators, control flow, functions,
650    #0$aC (Computer program language)
650    #6$aC (Langage de programmation)
650    #7$aC (Computer program language)$2fast$0(OCoLC)fst00843279
650    17$aC (programmeertaal)$2gtt
650    #7$aC$gProgrammiersprache$2gnd
650    #7$aLinguagem De Programacao (Outras)$2larpcal
650    #7$aComputers.$2nasat
650    #7$aC++ (programming language)$2nasat
650    #7$aC.$2tekord
650    #4$aC (Computer program language)
650    #2$aComputers.
650    #2$aSoftware.
653    ##$aProgramming languages
655    #7$aLærebøker.$2humord
655    #7$aProgrammering.$2humord
700    1#$aRitchie, Dennis M.,$d1941-2011,$eauthor.
776    08$iOnline version:$aKernighan, Brian W.$tC programming language.$b2nd ed.$dEnglewood Cliffs, N.J. : Prentice Hall, ©1988$w(OCoL
776    08$iOnline version:$aKernighan, Brian W.$tC programming language.$b2nd ed.$dEnglewood Cliffs, N.J. : Prentice Hall, ©1988$w(OCoL
830    #0$aPrentice-Hall software series.
856    41$3Table of contents$uhttp://www.gbv.de/dms/hebis-darmstadt/toc/10657762.pdf
856    42$uhttp://books.google.com/books?isbn=0131103709$zAdditional Information at Google Books
880    ##$6599-00/Thai$aCS -- 4 เล่ม.
969    ##$a050803/0t 1v/cls WD copy per JBB, lost$9LOCAL$A01CALS_PSU
969    ##$amo1213add-0824-1v-cam add copy$9LOCAL$A01CALS_PSU
974    ##$a9-18-16  cac  added 003 field$9LOCAL$A01CALS_PSU
974    ##$akd/kd 910503 17650642$9LOCAL$A01CALS_PSU
974    ##$akd/kd 910503 17650642$9LOCAL$A01CALS_PSU
994    ##$a92$bCAC
995    ##$aCSU ongoing load collection$b2019-03-26$cMaster record variable field(s) change: 599, 880$dCSU interim load$eUPDATED
```

```
AVA    ##$0991001848949702904$82265388240002904$a01CALS_CHI$bCSUChico$c3rd Floor - Main Collection$dQA76.73.C15 K47 1988$eavailable$f1$
AVA    ##$0991004651709702906$822124798160002906$a01CALS_UHL$bCSUEB-CIRC$cBook Stacks$dQA76.73.C15 K47 1988$eavailable$f1$g0$jmain$k0$p
AVA    ##$0991008808189702907$822120086740002907$a01CALS_UFR$bCSUFRESNO1$cLower Level$dQA76.73.C15 K47 1988$eavailable$f1$g0$jst$k0$p1$
AVA    ##$0991863663402909$82269854760002909$a01CALS_HUL$bHSULIB$cThird Floor, south$dQA76.73.C15 K47 1988$eunavailable$f1$g1$jMain3$k0
AVA    ##$0991009940859702910$82272807460002910$a01CALS_ULB$bCSULB$c4th Floor Stacks$dQA76.73 .C15 K47 1988$eunavailable$f1$g1$j4stacks
AVA    ##$0991006075009702914$82273942790002914$a01CALS_UNO$bOviatt$cFloor4$dQA76.73.C15 K47 1988$eavailable$f2$g1$jf4$k0$p1$qOviatt Ma
AVA    ##$099114449880101671$822132678510001671$a01CALS_USL$bsacstate$c4 NORTH$dQA76.73.C15 K47 1988$eavailable$f2$g1$j4n$k0$p1$qUniver
AVA    ##$0991003956269702918$822121773850002918$a01CALS_SFR$bMAIN$cLibrary Retrieval System$dQA76.73.C15 K47 1988$eavailable$f1$g0$jlr
AVA    ##$0991008658579702919$82265115730002919$a01CALS_SJO$bSJSU$c8th Floor$dQA76.73.C15 K47 1988$eavailable$f2$g0$jab8$k0$p1$qSJSU Li
AVA    ##$0991003733599702920$82232911770002920$a01CALS_PSU$bCPSLO_MAIN$cMain Collection$dQA76.73.C15 K47 1988$eavailable$f2$g1$js
AVA    ##$0991002159429702922$a01CALS_UST$bCSUSTAN$cGreen Collection$dQA76.73 .C15 K47 1988$eavailable$f1$g0$jtgc$k0$p1$qCSU Stani
INST   ##$a01CALS_NETWORK$bP$c71411943090002901
INST   ##$a01CALS_CHI$bP$c2165388250002904
INST   ##$a01CALS_UHL$bP$c21124798170002906
INST   ##$a01CALS_UFR$bP$c21120086750002907
INST   ##$a01CALS_HUL$bP$c2169854770002909
INST   ##$a01CALS_ULB$bP$c2172807470002910
INST   ##$a01CALS_MAL$bP$c71107010310002912
INST   ##$a01CALS_UNO$bP$c2173942800002914
INST   ##$a01CALS_USL$bP$c21132678520001671
INST   ##$a01CALS_SFR$bP$c21121773860002918
INST   ##$a01CALS_SJO$bP$c2165115740002919
INST   ##$a01CALS_PSU$bP$c2132911780002920
INST   ##$a01CALS_UST$bP$c2155884490002922
MMS    ##$b991009232479702901$a01CALS_NETWORK
MMS    ##$b991001848949702904$a01CALS_CHI
MMS    ##$b991004651709702906$a01CALS_UHL
MMS    ##$b991008808189702907$a01CALS_UFR
MMS    ##$b991863663402909$a01CALS_HUL
MMS    ##$b991009940859702910$a01CALS_ULB
MMS    ##$b991003065825502912$a01CALS_MAL
MMS    ##$b991006075009702914$a01CALS_UNO
MMS    ##$b99114449880101671$a01CALS_USL
MMS    ##$b991003956269702918$a01CALS_SFR
MMS    ##$b991008658579702919$a01CALS_SJO
MMS    ##$b991003733599702920$a01CALS_PSU
MMS    ##$b991002159429702922$a01CALS_UST
PLK    ##$aAdditional Form, Replaces.$b991069838028702901


------------------------------------------------------------------------------------------------------------


Leader 03437cam a2200685Ka 4500
001    991069838028702901
005    20181220055728.4
006    m     o  d
007    cr bn||||||abp
007    cr bn||||||ada
008    100325s1988    nju     o     001 0 eng d
019    ##$a835960408$a1044293912$a1056373107$a1074309905
020    ##$a9780133086218$q(electronic bk.)
020    ##$a0133086216$q(electronic bk.)
020    ##$z9780133086249
020    ##$z0133086240
020    ##$z0131103709
020    ##$z9780131103702
020    ##$z0131103628$q(pbk.)
020    ##$z9780131103627$q(pbk.)
035    ##$a(OCoLC)573196545$z(OCoLC)835960408$z(OCoLC)1044293912$z(OCoLC)1056373107$z(OCoLC)1074309905
035    ##$a(OCoLC)ocn573196545
037    ##$aCL0500000205$bSafari Books Online
040    ##$aOCLCE$beng$epn$cOCLCE$dOCLCQ$dOCLCA$dOCLCQ$dUMI$dDEBSZ$dOCLCQ$dOCLCF$dOCLCQ$dOCLCO$dOCLCA$dOCLCQ$dYDX$dN$T$dOCLCO$dZCU$dOCLC
042    ##$adlr
050    #4$aQA76.73.C15$bK47 1988
060    #4$aQA 76.73.C15 K39c 1988
072    #7$aCOM$x051060$2bisacsh
082    04$a005.13/3$219
084    ##$a54.53$2bcl
100    1#$aKernighan, Brian W.
245    14$aThe C programming language /$cBrian W. Kernighan, Dennis M. Ritchie.
250    ##$a2nd ed.
260    ##$aEnglewood Cliffs, N.J. :$bPrentice Hall,$c©1988.
300    ##$a1 online resource (xii, 272 pages)
336    ##$atext$btxt$2rdacontent
337    ##$acomputer$bc$2rdamedia
338    ##$aonline resource$bcr$2rdacarrier
505    00$gch. 1.$tTutorial introduction --$gch. 2.$tTypes, operators, and expressions --$gch. 3.$tControl flow --$gch. 4.$tFunctions a
506    ##$3Use copy$fRestrictions unspecified$2star$5MiAaHDL
520    ##$aIntroduces the features of the C programming language, discusses data types, variables, operators, control flow, functions,
533    ##$aElectronic reproduction.$b[S.l.] :$cHathiTrust Digital Library,$d2010.$5MiAaHDL
```

```
538    ##$aMaster and use copy. Digital master created according to Benchmark for Faithful Digital Reproductions of Monographs and Seri
583    1#$adigitized$c2010$hHathiTrust Digital Library$lcommitted to preserve$2pda$5MiAaHDL
588    0#$aPrint version record.
650    #0$aC (Computer program language)
650    #6$aC (Langage de programmation)
650    #7$aCOMPUTERS$xProgramming Languages$xC.$2bisacsh
650    #7$aC (Computer program language)$2fast$0(OCoLC)fst00843279
650    17$aC (programmeertaal)$2gtt
650    #7$aLinguagem De Programacao (Outras)$2larpcal
650    #2$aComputers.
650    #2$aSoftware.
653    ##$aProgramming languages
655    #4$aElectronic books.
700    1#$aRitchie, Dennis M.,$d1941-2011.
776    08$iPrint version:$aKernighan, Brian W.$tC programming language.$b2nd ed.$dEnglewood Cliffs, N.J. : Prentice Hall, ©1988$w(D
856    40$uhttps://www.safaribooksonline.com/library/view/-/9780133086249/?ar
994    ##$a92$bCAC
995    ##$aO'Reilly Safari Learning Platform: Academic edition$b2018-12-19$cMaster record variable field(s) change: 019$dCSU inter
AVE    ##$i01CALS_UMB
AVE    ##$i01CALS_ULA
AVE    ##$i01CALS_ULB
AVE    ##$i01CALS_USL
AVE    ##$i01CALS_MLM
AVE    ##$i01CALS_USM
AVE    ##$i01CALS_NETWORK
AVE    ##$i01CALS_PSU
AVE    ##$i01CALS_UBA
AVE    ##$i01CALS_UNO
AVE    ##$i01CALS_UST
AVE    ##$i01CALS_SJO
AVE    ##$i01CALS_MAL
AVE    ##$i01CALS_UHL
AVE    ##$i01CALS_UDH
AVE    ##$i01CALS_SDL
AVE    ##$i01CALS_UCI
AVE    ##$i01CALS_FUL
AVE    ##$i01CALS_SFR
AVE    ##$i01CALS_HUL
AVE    ##$i01CALS_UFR
AVE    ##$i01CALS_USB
AVE    ##$i01CALS_CHI
INST   ##$a01CALS_NETWORK$bE$c51504062970002901
MMS    ##$b991069838028702901$a01CALS_NETWORK
OWN    ##$i01CALS_NETWORK$lCO
PLK    ##$aAdditional Form, Replaced By.$b991009232479702901
PLK    ##$aReplaced By.$b991070871646502901


-----------------------------------------------------------------------------------------------------------


Leader 01728pam a2200373 a 4500
001    991070871646502901
005    20170822133628.0
006    m     o  d
007    cr unu
008    880303s1988    nju     o     00110 eng
010    ##$a   88005934
020    ##$a0-13-308623-2
020    ##$a0-13-308621-6
020    ##$a0-13-110362-8
035    ##$a(CKB)3710000000220165
035    ##$a(MH)001715469-3
040    ##$aDLC$cDLC$dCLS
050    0#$aQA76.73.C15$bK47 1988
060    #4$aQA 76.73.C15 K39c 1988
082    0#$a005.13/3$219
090    ##$aQA76.73.C15$bK47 1988
100    1#$aKernighan, Brian W.
245    14$aThe C programming language /$cBrian W. Kernighan, Dennis M. Ritchie.$h[electronic resource]
250    ##$a2nd ed.
260    0#$aEnglewood Cliffs, N.J. :$bPrentice Hall,$cc1988.
300    ##$a1 online resource (xii, 272 p. )
500    ##$aIncludes index.
504    ##$aIncludes index.
505    0#$aTutorial introduction -- Types, operators, and expressions -- Control flow -- Functions and program structure -- Pointers an
520    ##$aIntroduces the features of the C programming language, discusses data types, variables, operators, control flow, functions,
650    #0$aC (Computer program language)
650    #2$aComputers.
650    #2$aSoftware.
655    #4$aElectronic books.
```

700	1#$aRitchie, Dennis M.,$d1941-2011.
776	##$z0-13-308624-0
906	##$aBOOK
AVE	##$i01CALS_UMB
AVE	##$i01CALS_ULA
AVE	##$i01CALS_PUP
AVE	##$i01CALS_USM
AVE	##$i01CALS_NETWORK
AVE	##$i01CALS_PSU
AVE	##$i01CALS_UBA
AVE	##$i01CALS_UST
AVE	##$i01CALS_SJO
AVE	##$i01CALS_MAL
AVE	##$i01CALS_UDH
AVE	##$i01CALS_SDL
AVE	##$i01CALS_UCI
AVE	##$i01CALS_FUL
AVE	##$i01CALS_UFR
INST	##$a01CALS_NETWORK$bE$c51517699270002901
MMS	##$b991070871646502901$a01CALS_NETWORK
PLK	##$aReplaces.$b991069838028702901

Attachment 8

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. "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/silo/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 Ėxito: 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-
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
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)

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, *Críticas*, 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
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

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)
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, 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 September 2018