

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

QUALCOMM INCORPORATED,
Petitioner*,

v.

NETWORK SYSTEM TECHNOLOGIES, LLC,
Patent Owner.

Case IPR2024-00338
Patent No. 7,373,449 B2

DECLARATION OF SYLVIA HALL-ELLIS, PH.D.

* Qualcomm Incorporated filed a petition and a motion for joinder in IPR2024-01081 and has been joined as a petitioner in this proceeding.

I. INTRODUCTION

1. My name is Sylvia D. Hall-Ellis. I have been retained as an expert on behalf of Qualcomm Incorporated (“Petitioner”).

2. I have written this report on behalf of Petitioner to provide my expert opinion regarding the authenticity and public availability of several publications. 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 to 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 \$350 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 Master of Library Science from the University of North Texas in 1972 and a Ph.D. in Library Science from the University of Pittsburgh in 1985. Over the last 50-plus 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, as a member of the REFORMA National Board of Directors, as a member of the Editorial Board for the ALCTS premier cataloging journal, *Library Resources and Technical Services*, as a Co-Chair of the Library Research Round Table (LRRT) for the American Library Association, and as a member of the LRRT Nominating Committee.

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. My full curriculum vitae is attached hereto as Exhibit 1.

III. PRELIMINARIES

10. *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 on or around the listed dates of their respective publications below.

11. 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.

12. 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.

13. 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.

14. *Persons of ordinary skill in the art.* I am told by counsel that the subject matter of this proceeding generally relates to system-on-chip (SoC), interconnect, or networking technologies.

15. 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.

16. I am told by counsel that persons of ordinary skill in this subject matter or art would be familiar with the relevant technology in the early 2000s.

17. It is my opinion that such a person would have been engaged in research, learning thorough study and practice in the field and possibly through formal instruction the bibliographic resources relevant to his or her research.

IV. LIBRARY CATALOGING PRACTICES

A. MARC RECORDS AND THE ONLINE LIBRARY CATALOG

18. 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 1960s 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

¹ The full text of the standard is available from the Library of Congress at <http://www.loc.gov/marc/bibliographic/> (last visited November 21, 2023).

primary communications protocol for the transfer and storage of bibliographic metadata in libraries.

19. MARC is a framework into which descriptive bibliographic data are transcribed to interact with the software in online library catalogs to provide access to books, journals, and other resources in the collection. Bibliographic data provide points of access and can be searched by a person of ordinary skill in the art at the time of the invention (“POSITA,” see paragraphs 14-17 above) to identify and obtain resources in the library collection. An information seeker (or POSITA) can search a local online library catalog or the holdings of a group of libraries in a state or region or in the global catalog WorldCat.

20. MARC records are not designed for public viewing. Although a significant number of libraries provide access to the MARC version of a bibliographic record, the public display is designed to show information in a succinct manner that is quickly understood and useful to the information seeker. Libraries determine the default search for the online catalog to make the entry of search terms efficient and result in a successful search. Information seekers can enter a keyword, title, author, or standard number for the item. Libraries may also provide a search capability called “Summon” that allows the information seeker to enter known information about the item to conduct a search.

21. 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.² 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

² 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 November 21, 2023).

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

Why Is a MARC Record Necessary? LIBRARY OF CONGRESS.³

22. Thus, almost every major library in the world is MARC-compatible. See, e.g., *MARC Frequently Asked Questions (FAQ)*, LIBRARY OF CONGRESS.⁴ (“MARC is the acronym for MACHine-Readable Cataloging. It defines a data format that emerged from a Library of Congress-led initiative that began 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.⁵

³ <http://www.loc.gov/marc/umb/um01to06.html#part2>

⁴ <https://www.loc.gov/marc/faq.html>

⁵ <http://www.loc.gov/marc/bibliographic/>

23. Examining the MARC records for a specific item reveals the comprehensive data transcribed about a particular item at the time that cataloging and classification occurred. In addition to the creator, title, subjects, and standard numbers, additional information may provide additional and relevant data depending on the type of resource. Understanding the full extent of bibliographic data for an item and the points of access associated with it provides essential information that can be used to determine the indexing and public availability for documents described in this declaration.

B. MARC RECORDS

24. 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.⁶

25. 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 Declaration. Where this “mouse over” option was not available, I consulted the Directory of OCLC Libraries to identify the institution that created the MARC record.⁷

26. 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.⁸ For example, 6XX fields are termed “Subject Access Fields.”⁹ Among these, for example, is the 650 field; this is the “Subject Added Entry – Topical Term” field.¹⁰ The 650 field is a “[s]ubject added

⁶ <http://www.loc.gov/marc/bibliographic/bd3xx.html>

⁷ <https://www.oclc.org/en/contacts/libraries.html>

⁸ <http://www.loc.gov/marc/bibliographic/>

⁹ <http://www.loc.gov/marc/bibliographic/bd6xx.html>

¹⁰ <http://www.loc.gov/marc/bibliographic/bd650.html>

entry in which the entry element is a topical term.” 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.”¹¹ A defined portion of the Library of Congress Call (LCC) Number is the classification number, and “source of the classification number is *Library of Congress Classification* and the *LC Classification-Additions and Changes*.” Thus, included in the 050 field is a subject matter classification. Further, the 082 field is the “Dewey Decimal Call Number.”¹² A defined portion of the Dewey Decimal Call (DDC) Number is the classification number, and “source of the classification number is the *Dewey Decimal Classification and Relative Index*.” Thus, included in the 082 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 Decimal classification scheme) and uses it consistently. When the Library

¹¹ <http://www.loc.gov/marc/bibliographic/bd050.html>

¹² <http://www.loc.gov/marc/bibliographic/bd082.html>

of Congress assigns the LCC classification number, it appears as part of the 050 field. When the Library of Congress assigns the DDC classification number, it appears as part of the 082 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.

C. OCLC

27. 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.”¹³ Among other services, OCLC and its members are responsible for maintaining the WorldCat database (<http://www.worldcat.org/>), used by independent and

¹³ 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>).

institutional libraries throughout the world.

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

29. 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."¹⁴ The date upon which the most recent update to Field 005 occurred also appears, next to the label "Replaced." Thus, when an item's MARC record has been printed to hardcopy—as is the case with the MARC attachments to this Declaration—the date reflected next to the label

¹⁴ Field 005 is visible when viewing a MARC record via an appropriate computerized interface. But when a MARC record is printed directly to hardcopy from the OCLC database, the "005" label is not shown. The date in the 005 field instead appears next to the label "Replaced."

“Entered” is necessarily on or after the date the library first cataloged and indexed the underlying item.

30. 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 create additional MARC records—the other libraries instead rely on the original MARC record. They may update or revise the MARC record to ensure 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).

31. 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 reflects when the underlying item is accessible to the public. Upwards of two-thirds to three-quarters of book sales and magazine subscriptions to libraries come from a jobber or wholesaler for online and print resources. These resellers make it their business to provide items 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 and journals directly from publishers themselves, which provide delivery on a similarly expedited schedule. In general, libraries make these purchases throughout the year and shelve the newly received items as soon thereafter as possible to make them available to their patrons.

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

33. *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 a 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.

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

35. 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.

36. Before the widespread development of online databases to index articles in journals, magazines, conference papers, and technical reports, libraries purchased printed volumes of indices. Graduate library school education mandated that students learn about the bibliographic control of disciplines, the prominent indexing volumes, and searching strategies required to use them effectively and efficiently. Half of the courses that I studied in library school were focused on the bibliography and resources in academic disciplines.

37. Librarians consulted with information seekers to verify citations, check availability in union catalogs, printed books catalogs, the OCLC database, and make formal requests for materials (e.g., books, conference proceedings, journal articles). Requests were transmitted using Telex machines, rudimentary email systems, and the United States Postal Service. During my career, I have performed and supervised staff who handled these resource sharing tasks.

38. A major firm known for the breadth of subjects and comprehensive

treatment in the preparation of index volumes, the H. W. Wilson Company offered these reference resources since the firm was founded in 1898. The *Reader's Guide to Periodical Literature* is one of the best-known titles available from H. W. Wilson. Each volume includes a comprehensive index for 300 of the most popular and important periodicals published in the United States and Canada. Information seekers have subject access expressed in plain language terminology, author access, and cross references to find the desired results from their searches. The family of index titles included *Science & Technology Index*, *Business Periodicals*, *Applied Science & Technology Index*, *Humanities Index*, *Biological & Agricultural Index*, and *Industrial Arts Index*. These printed indices have been superseded by digital database offerings available to information seekers through Ebsco.

39. Online indexing services such as *Google Scholar*¹⁵ or *ScienceDirect*¹⁶ 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.

¹⁵ <https://scholar.google.com/>

¹⁶ <https://www.sciencedirect.com/>

40. A citation of a document by another is evidence that the document was publicly available and in use no later than the publication date of the citing document.

V. “NETWORKS ON CHIP” Textbook (Jantsch and Tenhunen, eds.)

41. Attachment 1G to this declaration is a copy of chapters 1 through 4 of a book titled *Networks on Chip* edited by Axel Jantsch and Hannu Tenhunen (hereafter “Jantsch” book) and issued by Kluwer Academic. Attachment 1G contains a true and correct copy of the title page, copyright page, and table of contents, along with chapters 1-4, as held by the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison. The text in Attachment 1G is complete for Chapters 1 through 4; 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.

42. Attached hereto as Attachment 1A is a true and correct copy of the MARC record for this monograph from the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison online catalog. Library ownership is indicated by the presence of the library’s code (“GZM”) in the 049 field. The library continues to update this MARC record to meet current cataloging rules. The most recent enhancement to the MARC record occurred on May 12, 2014, as shown in the 005 field (“20140512”). I personally identified and retrieved the

library catalog record which is Attachment 1A.

43. The International Standard Book Number (ISBN) on Exhibit 1014 (1- 4020-7392-5) matches the ISBN in field 020 of Attachment 1A. Therefore, Attachment 1G is a true and correct excerpt of the same book as the one that the cataloger at the Library of Congress used to create a MARC record (Attachment 1B) as part of the Cataloging in Publication (CIP) Program.¹⁷ A cataloger prepares basic information about the forthcoming book which is included on the verso of the title page (also called the copyright page). The information is provided to save time needed for the cataloging process at the library after the book is purchased. The CIP data includes the author (field 100), title and subtitle (field 245), ISBN (field 022), suggested subject headings (field 650), classification numbers for the Library of Congress classification number (field 050), and the Dewey Decimal classification number (field 082). These data elements are included in the MARC record that the Library of Congress contributes to the OCLC bibliographic database. 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 51477844 on January 8, 2003, as shown in the 008 field (“030108”). The library continues to update and enhance this MARC record to meet current cataloging

¹⁷ <https://www.loc.gov/publish/cip/>

rules. The most recent enhancement to Attachment 1B occurred on August 3, 2020, as shown in the 005 field (“200803”). I personally identified and retrieved the MARC record that is Attachment 1B.

44. Field 955 indicates that the Jantsch book arrived in the Library of Congress on April 7, 2003 (“2003-04-07”) and exited the cataloging process on May 12, 2003. Therefore, the book titled *Networks on Chip* (including the chapters 1 through 4 excerpted as Attachment 1G) was present and cataloged in the Library of Congress on May 12, 2003. It would have been able to be located by a person of ordinary skill in the art using the library catalog system and was thus publicly accessible as of this date.

45. Attached hereto as Attachment 1C is a true and correct copy of the MARC record for the Jantsch book titled *Networks on Chip* obtained from the OCLC bibliographic database. I personally identified and retrieved the MARC record that is Attachment 1C. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Attachment 1C, that library code is “DLC,” which means that the MARC record for this book was created 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 51477844 on January 8, 2003 (see ¶ 43 above) as shown in the “Entered” field (“20030108”). The library continues to update and enhance

this MARC record to meet current cataloging rules. The most recent enhancement to Attachment 1C occurred on June 8, 2022, as shown in the “Replaced” field (“20220608”). I personally identified and retrieved the MARC record that is Attachment 1C.

46. Attachment 1C further includes an entry in field 050 (“TK7895.E42 \$b N48 2003”)—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) and an entry in field 082 (“004.6”), a subject matter consistent with the Dewey Decimal classification system.

Attachment 1C further includes two descriptor terms reading “Systems on a chip \$x Design and construction” (see Attachment 1D, Library of Congress subject heading sh2002000568 and Attachment 1E, Library of Congress subject heading sh87005579) and “Computer networks” (see Attachment 1F, Library of Congress subject heading sh85029513) in the 650 fields. Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 1C 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 January 8, 2003, the MARC record attached hereto as Attachment 1C was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that added

this book to its collection, which means that the corresponding publication was publicly available on or before that same date through any library with access to the OCLC bibliographic database or through an individual library.

47. Attachment 1C indicates that the Jantsch book cataloged at the Library of Congress is currently available from 127 libraries. In view of above, the Jantsch book titled *Networks on Chip* was publicly available as of May 12, 2003, because by that date it had been cataloged and indexed at the Library of Congress, made part of the OCLC bibliographic database, and received in the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison. For these reasons, it is my opinion that the Jantsch book was published and accessible to the public no later than May 12, 2003.

VI. DOCUMENT 1: EXHIBIT 1007 (“GOOSSENS2003”)

48. Exhibit 1007 is a true and correct copy of chapter 4 of the Jantsch book I discussed above. I obtained Exhibit 1007 from counsel. The text in Exhibit 1007 is complete for those portions excerpted; 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 is a true and correct copy of Chapter 4 of the Jantsch book I discussed above in a condition that creates no suspicion about its authenticity.

49. For the reasons I explained in Section V above, it is my opinion that

Exhibit 1007, as a true and correct copy of chapter 4 of the Jantsch book, was published and accessible to the public no later than May 12, 2003. Exhibit 1007 was published and accessible to the public no later than May 12, 2003 because it is contained within the Jantsch book that, as I explained above in Section V, was published and accessible to the public no later than May 12, 2003.

VII. DOCUMENT 2: EXHIBIT 1014

50. Exhibit 1014 is a true and correct copy of the cover, title page and copyright page, table of contents, preface, and chapters 1 through 3 of the Jantsch book I discussed above. I obtained Exhibit 1014 from counsel. The text in Exhibit 1014 is complete for those portions excerpted; 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 is a true and correct copy of the excerpts of the Jantsch book in a condition that creates no suspicion about its authenticity.

51. For the reasons I explained in Section V above, it is my opinion that Exhibit 1014, as a true and correct copy of the cover, title page and copyright page, table of contents, preface, and chapters 1 through 3 of the Jantsch book, was published and accessible to the public no later than May 12, 2003. Exhibit 1014 was published and accessible to the public no later than May 12, 2003 because it is contained within the Jantsch book that, as I explained above in Section V, was

published and accessible to the public no later than May 12, 2003.

VIII. SUMMARY OF OPINIONS

52. In view of the foregoing, it is my opinion that the Jantsch textbook discussed above, including the excerpts provided as Exhibits 1007 and Exhibit 1014, was publicly available no later than May 12, 2003.

53. 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.

54. 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 1034 of Title 18 of the United States Code.

DATED: October 24, 2024


By: 
Sylvia D. Hall-Ellis, Ph.D.

Exhibit 1

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 as a faculty member teaching 400 graduate students in technical services (cataloging, bibliographic control, classification), automation, and networking. Participated in distance education program and coordinated annual conference. Conducted university and Texas Library Association-funded field research focused on library collection development and academic achievement.
- 1992-1993** **Head Librarian, Rocky Mountain College of Art & Design, Denver, Colorado.** Responsible for the daily operation, selection and acquisition of materials, formulation of policies for library operations, media center, and photography/slides archives. Designed and implemented library automation and delivery of electronic resources to college community.
- 1981-1985** **Development Officer, PRLC, Inc., Pittsburgh, Pennsylvania.** Served as senior member of corporate management team in strategic planning, development of proposals and contracts, implementation, and evaluation of new services, products, and programs for 100 institutional member organizations. Coordinated the development of proposals and contracts totaling \$4,000,000 annually. Provided creative leadership to corporate committees to solicit ideas, identify goals and objectives, plan, develop, present, and evaluate professional development opportunities.
- 1981** **Director of Library Development, Pennsylvania Department of Education, Harrisburg, Pennsylvania.** Responsible for statewide development, technical assistance, professional development, resource sharing, children's services, institutional library services, networking, and state aid program for all libraries throughout the Commonwealth. Functioned as liaison to Governor's Advisory Council, LSCA Advisory Council, District Administrators, private colleges, universities, consortia managers, and network directors. Supervised \$14,000,000 formula-based state aid program and \$3,000,000 grant awards to individual libraries, consortia, and networks.
- 1978-1981** **Assistant Director, Southern Tier Library System, Corning, New York.** Coordinated operation of system-wide programs (technical assistance, professional development, resource sharing, technical services, outreach) to 40 public libraries in 5 counties serving 500,000 residents. Solicited ideas, identified goals, sponsored, and evaluated professional development opportunities and technical assistance sessions.
- 1976-1978** **Division Librarian for Technical Services, Corpus Christi Public Libraries, Corpus Christi, Texas.** Provided leadership in acquisitions, cataloging, serials control, and processing for main library and 4 branches serving 250,000 residents. Participated as senior member of library management team. Compiled and prepared technical evaluations, reports, and statistical analyses of Division operations to measure the achievement and cost of annual goals, objectives, and staff performance.
- 1975-1976** **System Coordinator, San Antonio Major Resource Center, San Antonio, Texas.** Served as senior member of the management team for District X Office, charged to provide technical assistance, resource sharing, media services, and professional development to librarians and staff representing 30 public library jurisdictions in 21 counties serving 1,500,000 residents. Functioned as liaison to System Director, staff, and members of governing bodies with the System Board of Directors and the Texas State Library and Historical Commission. Prepared LSCA grant applications and monitored awards totaling \$1,100,000 annually.

- 1973-1975** **Bilingual Branch Librarian, San Antonio Public Library, San Antonio, Texas.** Worked as librarian providing reference, information, and readers' advisory services in branch serving 50,000 Spanish-speaking residents in southwest San Antonio. Participated in collection development and resource acquisition activities, specializing in children's work, Spanish language resources, and multicultural studies.
- 1972-1973** **Librarian, Holding Institute, Laredo, Texas.** Worked as high school librarian serving 500 boarding students in Spanish-speaking environment of private school. Provided reference, research assistance, and library instruction to students and 35 faculty members.
- 1966-1971** **Rockford Public Library, Rockford, Illinois.** Worked in branches as part-time as a Library Assistant, Clerk, and Page in city library serving 150,000 residents.

PUBLICATIONS

Editor-reviewed Monographs (Completed and in Progress)

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

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

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

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

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

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

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

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

Editor-reviewed Chapters (Completed and In Progress)

Hall-Ellis, Sylvia D. "Prepared to Lead: Talent, Skills, and Competencies." In *Telling the Technical Services Story*. Chicago, IL: American Library Association, 2020. Program Presentation for Re-Accreditation by the American Library Association, School of Information, San Jose State University, 2022.

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. Lanham, 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. Lanham, 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>
- 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. *Preparing Catalogers and Metadata Specialists Through Competency-Based Education: The San José State University iSchool Experience*. Report for the 2021 Re-Accreditation Program Presentation to the American Library Association. San José, Calif.: San José State University iSchool, 2020.

Hall-Ellis, Sylvia D. *Colorado Community College System Grant Handbook*. Denver, Colo.: 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, Colo.: 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, Schuylers, Steuben, and Yates Counties*. Ithaca, N.Y.: South Central Reference and Research Council, 1981.
- Hall, Sylvia D. *OCLC/ILL Subsystem at Southern Tier Library System*. Albany, N.Y.: SUNY/OCLC Network, 1980.
- Hall, Sylvia D. *Retrospective Conversion for Major Public Libraries in South Central New York State: A Technical Report*. 2 vols. Corning, N.Y.: Southern Tier Library System, 1979-1980.
- Hall, Sylvia D. *The Librarian and the Mexican-American Patron in Texas*. Austin, Tex.: Texas State Library, 1978.

Other Publications

- Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association*. St. Paul, MN: St. Catherine University, 2017. 255 p.
- Hall-Ellis, Sylvia D., and Robert O. Ellett, Jr., guest editors. *RDA Testing: Lessons Learned and Challenges Revealed: A Special Issue of Cataloging & Classification Quarterly* 49, no. 7-8 (November 2011).
- Hall-Ellis, Sylvia D. and Robert O. Ellett, Jr. *Fundamentals of Cataloging I: An Online ALCTS Course*. Chicago, Ill.: American Library Association, 2009.
- Hall-Ellis, Sylvia D., Angela J. Kinney and Martha Ann Bace, editors. *Best Practices in Cataloging Education Identified at the 2007 ALCTS Pre-conference: A Special Issue of the Journal for Library and Information Science Education* 49, no. 2 (Spring 2008).

- Hall-Ellis, Sylvia D. *Cataloging Education: A New Emphasis for the Library and Information Science Curriculum*. Presentation delivered to the ALCTS Pre-conference, 132nd Annual Conference, American Library Association, Washington, D.C., June 22, 2007. <http://www.loc.gov/aba/professional/>
- Hall-Ellis, Sylvia D. “*Cataloger Competencies: Do the Employers Require What the Professors Teach?*” Presentation delivered to the ALCTS CCS Heads of Cataloging Discussion Group, 131st Annual Conference, American Library Association, New Orleans, Louisiana, June 26, 2006. <http://dlist.sir.arizona.edu/>
- Hall-Ellis, Sylvia D., guest editor, with Deborah S. Grealy, Clara L. Sitter, and Tassos Petrou. *Library Students’ Voices in Colorado: An Issue of Colorado Libraries* 31:1 (Spring 2005).
- Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association*. Denver, Colo.: University of Denver, College of Education, 2003. 125 p. with CD-ROM.
- Hall-Ellis, Sylvia D. and Deborah S. Grealy, guest editors. *Library Education in Colorado: An Issue of Colorado Libraries* 29, no. 2 (Summer 2003).
- Hall-Ellis, Sylvia D., ed. *Library and Information Science Program Self-Study Document for Accreditation by the American Library Association*. Denver, Colo.: University of Denver, College of Education, 2003.
- Hall-Ellis, Sylvia D. “Adaptive Furniture for the ADA-Qualified: An Investigative Study.” 2001. Unpublished manuscript.
- Hall-Ellis, Sylvia D. “Cataloging Electronic Serials: A Study in User Access and Article Delivery.” 2001. Unpublished manuscript.
- Hall-Ellis, Sylvia D. “E-books: Acquisition and Circulation Issues in Public Libraries.” 2001. Unpublished manuscript.
- Hall-Ellis, Sylvia D. “Authority Control and Its Importance in School Library Media Online Catalogs.” 1999. Unpublished manuscript.
- Hall, Sylvia D. “Cooperation Improving Service in San Antonio System.” *Texas Libraries* XXXVII (Spring 1976): 27-29.
- Hall, Sylvia D. “Summer is Sea Monstrously Fun!” *Texas Library Journal* L (December 1974): 224-226.
- Hall, Sylvia D. “A Brief Annotated Bibliography of Non-book Materials on the Mexican American.” *Texas Library Journal* XLVIII (October 1972): 213-232.

Book Reviews

- ALA Book of Library Grant Money*. Ann Kepler, ed. 8th ed. Chicago, IL: American Library Association, 2011. 978-0-8389-1058-0, \$175.00 [*Colorado Libraries*, June 2012]
- Developing a Compensation Plan for Your Library*. Paula M. Singer and Laura L. Francisco, comps. 2nd ed. Chicago, IL: American Library Association, 2009. 978-0-8389-0985-0, \$57.00 [*Colorado Libraries*, March 2012]
- Frugal Librarian*. Carol Smallwood, ed. Chicago, IL: American Library Association, 2011. 978-0-83891-075-7, \$42.00 [*Colorado Libraries*, February 2012]

- Dowlin, Ken. *Getting the Money: How to Succeed in Fundraising for Public and Nonprofit Libraries*. Westport, Conn.: Libraries Unlimited, 2008. 978-1-59158-597-X, \$50.00 [*Colorado Libraries*, February 2009]
- Lundahl, Mats. *Bebo de Cuba; Bebo Valdés y Su Mundo*, by Mats Lundahl. Traducción de Linda Oakeshott Dragó. Barcelona: RBA Libros, 2008. 978-849-86-7259-6, \$33.99 [*Críticas*, November 2008]
- Bonelli, Florencia. *Marlene*, by Florencia Bonelli. Buenos Aires: Aguilar, Altea, Taurus, Alfaguara, 2008. 978-987-04-0995-3, \$24.95 [*Críticas*, October 2008]
- Keret, Etgar. Traducción de Ana María Bejarano. *Pizzería Kamikaze y Otras Relatos*. México, D.F.: Editorial Sexto Piso, 2008. 968-5679-29-0, no price given [*Críticas*, September 2008]
- Littauer, Marita, and Florence Littauer. *Enriquece tu Comunicación*. Miami, Fla.: Editorial Unilit, 2008. 978-0-7899-1521-4, \$11.99 [*Críticas*, August 2008]
- Martínez, Guillermo. *La Muerte Lenta de Luciana B*. New York: Rayo Planeta, 2008. 978-0-06-156551-9, \$14.95 [*Críticas*, May 2008]
- Matthews, Joseph R. *The Evaluation and Measurement of Library Services*. Westport, Conn.: Libraries Unlimited, 2007. 978-1-59158-532-9, \$50.00 [*Colorado Libraries*, April 2008]
- Gómez-Jurado, Juan. *A Masacre de Virginia Tech: Anatomía de una Mente Torturada*. Barcelona: Ediciones El Andén, 2007. 978-84-935789-4-7, €16.50 [*Críticas*, February, 2008]
- Matthews, Joseph R. *Library Assessment in Higher Education*. Westport, Conn.: Libraries Unlimited, 2007. 978-1-59158-531-2, \$45.00 [*Colorado Libraries*, February 2008]
- Chacón, Inma. *La Princesa India: Cuando el Viento Azul*. México: Alfaguara, 2006. 970-770398-9, \$19.95 [*Críticas*, June 2006]
- Vargas Llosa, Mario. *Israel/Palestina: Paz o Guerra Santa*. fotogs. by Morgana Vargas Llosa. Spain/U.S.: Aguilar: Santillana, 2006. 84-03-09691-7, paper, \$19.95 [*Críticas*, May 2006]
- Findlay, Diane. *Digging into Dewey*. Fort Atkinson, Wisc.: Upstart Books, 2005. 978-1-932146-18-9, \$16.95 [*Colorado Libraries*, May 2006]
- Pequeñas Resistencias 4: Antología del Nuevo Cuento Norteamericano y Caribeño*. Menéndez, Ronaldo, Ignacio Padilla & Enrique del Risco, eds. Madrid: Páginas de Espuma, 2005. 84-95642-59-X, paper, \$36.95 [*Críticas*, April 2006]
- Rueda, André. *Vengo a Salvar a España: Biografía de Un Franco Desconocido*. Madrid: Nowtilus (Investigación Abierta), 2005. photogs., bibliog. 84-9763-202-8, paper, \$23.95 [*Críticas*, March 2006]
- Becerra, Ángela. *El Penúltimo Sueño*. 1st ed. Barcelona: Planeta, 2005. 84-08-05795-2, \$25.95. [*Críticas*, January 2006]
- Aguilar Camín, Héctor. *La Conspiración de la Fortuna*. 1st ed. México: Planeta, 2005. 970-37-0368-2, \$19.95 [*Críticas*, October 2005]
- México en Sus Libros*. Enrique Florescano with Pablo Mijangos, editors & compilers. 1st ed. México, D.F.: Taurus Aguilar, 2004. 968-19-0783-3, \$16.95 [*Críticas*, September 2005]

- Historia Económica de México*. Coordinator, Enrique Semo. México: Universidad Nacional Autónoma de México, Editorial Oceano de México, 2004. 13 vols. 970-32-0805+3 (obra completa), \$9.25 each [*Críticas*, August 2005]
- Czarnowsky, Christyne A. and Michael H. Williams. *Managing Your Wired Workforce: A Practical Guide*. Denver, Colo.: Bradford Pub., 2003. 1-883726-79-4, \$24.95 [*Colorado Libraries*, 2004]
- Abraham Lincoln*. Colección Grandes Biografías. Madrid: Edimat Libros, 2003. 84-8403-858-0, €4.95 [*Críticas*, 2004]
- Alponte, Juan María. *Colón: el Hombre, el Navegante, la Leyenda*. México, D.F.: Aguilar, 2003. 968-19-1260-8, \$16.95 [*Críticas*, 2004]
- Olcese Salvatecci, Alfieri. *Cómo Estudiar con Éxito: Técnicas y Hábitos Para Aprender Mejor*. México, D.F.: Alfaomega Grupo Editor, 2002. 970-15-0764-9 [*Críticas*, 2003]
- Mujeres Como Islas: Antología de Narradoras Cubanas, Dominicanas, Puertorriqueñas*. 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]
- Puig 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, Life Member

Association for Library Collections & Technical Services (now CORE)
Cataloging, Classification, and Metadata Section
ALCTS Editorial Board, Member, 2015-2019
Committee on Cataloging: Description and Access, Member, 2008-2012
Test Site for RDA, Manager, 2009-2011
Nominating Committee, Member, 2007-2008
Committee for Education and Training of Catalogers, Chair, 2006-2007, Member, 2003-2007
Competencies & Education for a Career in Cataloging Interest Group, Chair, 2010-2012; Member, 2010-
President's Program Committee, 2015-2016
Education Committee, Member, 2005-2007
Task Force on Competencies and Education for a Career in Cataloging, Chair, 2007-2009
2007 Annual Meeting Pre-conference, "What They Don't Teach in Library School: Competencies,
Education and Employer Expectations for a Career in Cataloging," Steering Committee Chair
Fundraising Committee, Member, 2004-2006
Office of Diversity
Committee on Diversity, Chair, 2010-2011
Diversity Research Advisory Committee, Member, 2009-2010
Spectrum Scholar Mentor, 2009
Diversity Grants Review Committee, Member, 2007, 2010, 2012
Office of Statistics and Research
Research and Statistics Committee, Member, 2005-2007; Intern, 2003-2005
Library Research Round Table
Nominating Committee, 2022-2023
Membership Committee, 2019-2021
Board of Directors, Member At-Large, 2009-2012
Office of Accreditation
Accreditation Panel Member (training completed 2004)
Accreditation Review Panel Member (training completed 2004)

REFORMA, Life Member

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

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

American Association of University Women Association for Library and Information Science Education (ALISE)

University of Denver ALISE Representative, 2003-2008; 2010-2011
Membership Advisory Committee, 2007-2010
Technical Services Education Special Interest Group, 2003-
Garfield Doctoral Dissertation Award Reviewer, 2012
Garfield Doctoral Dissertation Scholarship Reviewer, 2014

American Association of University Professors (AAUP), 2006-

American Grant Writers Association, 2014- Grant Professionals Institute, 2016- National Grants Management Association, 2014-

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

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

CERTIFICATION

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

AWARDS AND HONORS

Advanced Practitioner for Service Learning and Community Engagement, University of Denver, 2011
Platinum Star Alumnae, College of Information, Library Science & Technologies, University of North Texas, 2009
Commendation for Integration of Technologies in Teaching & Learning Environment, University of Denver, 2006
Outstanding Adjunct Faculty Member Award, College of Education, University of Denver, 2002
Beta Phi Mu, Pi, University of Pittsburgh, 1985
Alpha Lambda Sigma, University of North Texas, 1972
Albert Nelson Marquis Lifetime Achievement Award, Who's Who in America
Who's Who in American Women
Who's Who of Women Executives
Dictionary of International Biography
Who's Who in the East
Who's Who in the South and Southwest
Who's Who in the World
Who's Who of Online Professionals
Who's Who in Library and Information Science
2,000 Notable Women
Who's Who of Emerging Leaders
Who's Who in Professional and Executive Women
Who's Who in American Education
International Who's Who of Professional and Business Women
International Leaders in Achievement
International Educator of the Year
Who's Who in Finance and Industry
Who's Who in Finance and Business

Invited International and National Conference Presentations

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- Hall-Ellis, Sylvia D. *Employers' Expectations for Technical Services Librarians: What We Don't Know*. Presentation delivered to the ALCTS Research and Publications Committee Program, 133rd Annual Conference, American Library Association, Anaheim, Calif., June 28, 2008.
- Hall-Ellis, Sylvia D., Virginia R. Maloney, and Mary Stansbury. *Institutional Responses to Engaged Scholarship: The Carnegie Foundation Engaged University Classification at Two Universities*. Presentation delivered to the 94th Annual Conference, Association for Library and Information Science Education, Philadelphia, Pa., January 11, 2008.
- Hall-Ellis, Sylvia D. *Puzzles, Problems, and Predicaments*. Presentation delivered to the ALCTS Research Discussion Group, 132nd Annual Conference, American Library Association, Washington, D.C., June 23, 2007.
- Hall-Ellis, Sylvia D. *Cataloging Education: A New Emphasis for the Library and Information Science Curriculum*. Presentation delivered to the ALCTS Pre-conference, 132nd Annual Conference, American Library Association, Washington, D.C., June 22, 2007. <http://www.loc.gov/catdir/cpsa/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.
- Greal, 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.
- 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.

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.

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

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

University of Northern Colorado, Greeley, Colo., November 10, 1998.

United Technology Educational Partnership, Grand Junction, Colo., November 9, 1998.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- Hall-Ellis, Sylvia D. *Grant Writing: An Introduction for Public School Administrators*. Sponsored and hosted by the Office of Administrative Services, Region One Education Service Center, Edinburg, Tex., April 19, 1995.
- Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future*. Sponsored and hosted by the United Independent School District, Laredo, Tex., March 31, 1995.
- Hall-Ellis, Sylvia D. *Using USMARC*. Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., March 29, 1995.
- Hall-Ellis, Sylvia D. *Library Media Center Policies and Guidelines: How to Prepare for School Board Adoption*. Sponsored and hosted by the Office of Technology and Media Services, La Joya Independent School District, La Joya, Tex., March 10, 1995.
- Hall-Ellis, Sylvia D. *Developing District-wide Policies and Guidelines for Library Media Centers*. Sponsored and hosted by the Office of Technology and Media Services, La Joya Independent School District, La Joya, Tex., January 10, 1995.
- Hall-Ellis, Sylvia D. *MARC Cataloging of Audiovisual Materials for Library Media Centers*. Sponsored and hosted by the Office of Library Media and Technology Services, Cypress-Fairbanks Independent School District, Houston, Tex., December 1, 1994.
- Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future*. Sponsored and hosted by United Independent School District, Laredo, Tex., October 14, 1994.
- Hall-Ellis, Sylvia D. *The School Library Media Specialist in the 21st Century: Visions for the Future*. Sponsored and hosted by the Laredo Independent School District, Laredo, Tex., October 7, 1994.
- Hall-Ellis, Sylvia D. *MARC Cataloging for Library Media Centers*. Sponsored by the Office of Library Media and Technology Services, Cypress-Fairbanks Independent School District, Houston, Tex., October 5, 1994.
- Hall-Ellis, Sylvia D. *Jump to the Head of the Class: Undergraduate Library Resources Available at Sam Houston State University*. Sponsored by the Office of the Associate Vice President for Student Services, Sam Houston State University, Huntsville, Tex., October 4, 1994.
- Hall-Ellis, Sylvia D. *CD-ROMs - 1994's Newest and the Best for Secondary Level Media Centers*. Sponsored and hosted by Clear Lake Independent School District, Houston, Tex., June 3, 1994.
- Hall-Ellis, Sylvia D. *Introduction to Classification*. Sponsored by AJ Seminars, Rockville, Maryland; presented at Holiday Inn, Market Center, Dallas, Tex., May 18, 1994.
- Hall-Ellis, Sylvia D. *Basic Descriptive Cataloging*. Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., May 4, 1994.
- Hall-Ellis, Sylvia D. *Shaking the Money Tree - Part II: Writing Successful Grant Applications*. Sponsored by Donna Independent School District, Donna, Texas, and Region One Education Service Center Edinburg, Texas; presented at South Texas Community College Library, McAllen, Tex., April 29, 1994.
- Hall-Ellis, Sylvia D. *Using MARC*. Sponsored by AJ Seminars, Rockville, Maryland; presented at Holiday Inn, Market Center, Dallas, Tex., April 6, 1994.
- Hall-Ellis, Sylvia D. *Automated Authority Control*. Sponsored by AJ Seminars, Rockville, Maryland; presented at University Hilton Hotel, Houston, Tex., March 23, 1994.

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

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

GRADUATE COURSES TAUGHT

San José State University, School of Library and Information Science

INFO 249 – Advanced Cataloging and Classification (Fall 2015, 2016; Summer 2016)
INFO 287 – Special Topics in Cataloging and Classification (Spring 2017, 2018, 2019, 2020, 2021, 2022, 2023; Fall 2020, 2021, 2022, 2023)
LIBR 248 – Beginning Cataloging and Classification (Summer 2002, 2003, 2004, 2005, 2006)
LIBR 249 – Advanced Cataloging and Classification (Summer 2003; Fall 2014; Summer 2015)

University of Denver, University College

Grant Writing Certificate Seminar for Non-Profit Organizations (Fall 2020, 2021; Spring 2021)

University of Denver, Morgridge College of Education HED

5991 – Grant Writing in Higher Education (Spring 2011)
LIS 4010 – Organization of Information (Fall 2002-2009; Winter 2005; Spring 2004-2006)
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 2002-2009; Summer 2005, 2006 – course renumbered LIS 4070, September 2007)
LIS 4401 – Descriptive Cataloging (Winter 2001; Spring 2002, 2003, 2006-2009; Summer 2005)
LIS 4402 – Subject Cataloging (Spring 2001, 2009; Summer 2005, 2006, 2007, 2008)
LIS 4403 – Classification Schemes (Fall 2007, 2008, 2009 (DS)) LIS
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 2006-2009)
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 Moultny, 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 Hessel (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 Salazar (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 Neural 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 Salazar. 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 Ditttrick-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 Svendsen, 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 Svendsen, 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. Greal. 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. Greal (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. Greal (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. Saucedo, 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. Bubbs, 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 Aggie 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 I. 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, 2001-2003
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 June 2023

Attachment 1A

MARC Bibliographic Record

LEADER 01287cam a22004094a 4500
001 9961071333602122
005 20140512163445.0
008 030108t20032003maua b 000 0 eng
010 \$a 2003041603
015 \$aGBA3-Z3254
019 \$a51780596
020 \$a1402073925 (alk. paper)
035 \$a(OCOLC)ocm51477844
035 \$a(WU)6107133-uwmadisondb
035 \$z(OCOLC)51780596
035 \$a(EXLNZ-01UWI_NETWORK)999949359402121
040 \$aDLC\$beng\$cDLC\$dUKM\$dC#P\$dOHX
042 \$apcc
049 \$aGZMA
050 00 \$aTK7895.E42\$bN48 2003
072 _7 \$aTK\$2lcco
082 00 \$a004.6\$b221

245 00 \$aNetworks on chip /\$cedited by Axel Jantsch, Hannu Tenhunen.

264 _1 \$aBoston :\$bKluwer Academic Publishers,\$c[2003]

264 _4 \$c@2003

300 \$aviii, 303 pages :\$billustrations ;\$c25 cm

336 \$atext\$btxt\$2rdacontent

337 \$aunmediated\$bn\$2rdamedia

338 \$avolume\$bnc\$2rdacarrier

504 \$aIncludes bibliographical references.

650 _0 \$aSystems on a chip\$xDesign and construction.

650 _0 \$aComputer networks.

700 1_ \$aJantsch, Axel.

700 1_ \$aTenhunen, Hannu.

856 42 \$3Table of contents\$u<http://www.loc.gov/catdir/toc/fy036/2003041603.html>

938 \$aOtto Harrassowitz\$bHARR\$nhar035011733\$c117.00 EUR

994 \$aX0\$bGZM

997 \$aMARCIVE

MMS IDs

Document ID:

999949359402121

Network Electronic IDs:

Network Physical IDs:

999949359402121

mms_mad_ids:

9961071333602122

Attachment 1B

LIBRARY | CATALOG

[Search Navigation ▼](#)

7 of 7



BOOK

Networks on chip

[Full Record](#)[MARC Tags](#)

000	01238cam a22002894a 4500
001	13051026
005	20080310123318.0
008	030108s2003 maua b 000 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 pc17 2003-01-08 to ASCD; c jf03 2003-01-09 to Subj. d jg12 2003-01-14 e jf12 2003-01-14 to Dewey a aa20 2003-01-17 a ps16 2003-04-07 bk rec'd, to CIP ver. a jg00 2003-04-28; f jg02 2003-05-10 e jg08 2003-05-12 to Dewey
010	___ a 2003041603
020	___ a 1402073925 (alk. paper)
040	___ a DLC c DLC d DLC
042	___ a pcc
050	00 a TK7895.E42 b N48 2003
082	00 a 004.6 2 21
245	00 a Networks on chip / c edited by Axel Jantsch, Hannu Tenhunen.
260	___ a Boston : b Kluwer Academic Publishers, c c2003.
300	___ a viii, 303 p. : b ill. ; c 25 cm.
504	___ a Includes bibliographical references.
650	_0 a Systems on a chip x Design and construction.
650	_0 a Computer networks.
700	1_ a Jantsch, Axel.
700	1_ a Tenhunen, Hannu.

856	41	3 Table of contents u http://www.loc.gov/catdir/toc/fy036/2003041603.html
856	42	3 Publisher description u http://www.loc.gov/catdir/enhancements/fy0820/2003041603-d.html

[Request this Item](#)[!\[\]\(245e28142faa1131d188be32c80447b7_img.jpg\) LC Find It](#)

Item Availability



CALL NUMBER	TK7895.E42 N48 2003 Copy 1
-------------	---

Request in	Jefferson or Adams Building Reading Rooms
------------	---

Status	Not Charged
--------	-------------

Attachment 1C

OCLC 51477844 No holdings in XXX - 127 other holdings

Books		Re		En 200		Re 2022060					
		c		tere 301		place 8184913					
		Stat		d 08		d .7					
Type	a	ELvl	4	Srce		Audn		Ctrl		Lang	eng
BLvl	m	Form		Conf	0	Biog		MRec		Ctry	mau
		Cont	b	GPub		LitF	0	Indx	1		
Desc	a	Ills	a	Fest	0	DtSt	s	Dates	2003		

010 2003041603

040 DLC \$b eng \$c DLC \$d UKM \$d C#P \$d OHX \$d BAKER \$d WEA \$d NLGGC \$d YDXCP \$d BTCTA \$d
STF \$d OCLCQ \$d IG# \$d DEBBG \$d OCLCQ \$d ILU \$d OCLCF \$d OCLCO \$d OCLCQ \$d I8M \$d OCLCQ \$d
ZAP \$d OCLCQ \$d NAG \$d OCLCO

015 GBA3Z3254 \$2 bnb

016 7 B0309109 \$2 bccb

019 51780596

020 1402073925 \$q (alk. paper)

020 9781402073922 \$q (alk. paper)

042 pcc

050 00 TK7895.E42 \$b N48 2003

072 7 TK \$2 lcco

082 00 004.6 \$2 21

084 53.57 \$2 bcl

090 \$b

245 00 Networks on chip / \$c edited by Axel Jantsch, Hannu Tenhunen.

260 Boston : \$b Kluwer Academic Publishers, \$c ©2003.

300 viii, 303 pages : \$b illustrations ; \$c 25 cm

336 text \$b txt \$2 rdacontent

337 unmediated \$b n \$2 rdamedia

338 volume \$b nc \$2 rdacarrier

504 Includes bibliographical references.

505 00 \$g 1. \$t Will Networks on Chip Close the Productivity Gap? / \$r Axel Jantsch and Hannu Tenhunen

-- \$g 2. \$t A Design Methodology for NoC-based Systems / \$r Juha-Pekka Soininen and Hannu Heusala --

\$g 3. \$t Mapping Concurrent Applications onto Architectural Platforms / \$r Andrew Mihal and Kurt

Keutzer -- \$g 4. \$t Guaranteeing The Quality of Services in Networks on Chip / \$r Kees Goossens, John

Dielissen, Jef van Meerbergen, Pete Poplavko, Andrei Radulescu, Edwin Rijpkema, Erwin Waterlander

and Paul Wielage -- \$g 5. \$t On Packet Switched Networks for On-chip Communication / \$r Shashi Kumar

-- \$g 6. \$t Energy-reliability Trade-off for NoCs / \$r Davide Bertozzi, Luca Benini and Giovanni De Micheli -

- \$g 7. \$t Testing Strategies for Networks on Chip / \$r Raimund Ubar and Jaan Raik.

520 1 "Networks on Chip presents a variety of topics, problems and approaches with the common

theme to systematically organize the on-chip communication in the form of a regular, shared

communication network on chip, an NoC for short." "The book is organized in three parts. The first deals

with system design and methodology issues. The second presents problems and solutions concerning

the hardware and the basic communication infrastructure. Finally, the third part covers operating

system, embedded software and application. However, communication from the physical to the

application level is a central theme throughout the book." "The book serves as an excellent reference source and may be used as a text for advanced courses on the subject."--Jacket.

650 0 Systems on a chip \$x Design and construction.

650 0 Computer networks.

650 6 Réseaux d'ordinateurs. \$0 (CaQQLa)201-0007877

650 7 Computer networks. \$2 fast \$0 (OCoLC)fst00872297

650 7 Systems on a chip \$x Design and construction. \$2 fast \$0 (OCoLC)fst01141474

650 7 Chip \$2 gnd \$0 (DE-588)4197163-2

650 7 Netzwerk \$2 gnd \$0 (DE-588)4171529-9

650 7 CHIPS (ELECTRONICS) \$2 nasat

650 7 SYSTEMS ANALYSIS. \$2 nasat

650 7 COMPUTER NETWORKS. \$2 nasat

650 7 TELECOMMUNICATION. \$2 nasat

700 1 Jantsch, Axel.

700 1 Tenhunen, Hannu.

856 41 \$3 ebrary \$u <http://site.ebrary.com/id/10067276>

856 41 \$3 Table of contents \$u <http://catdir.loc.gov/catdir/toc/fy036/2003041603.html>

856 42 \$3 Publisher description \$u <http://catdir.loc.gov/catdir/enhancements/fy0820/2003041603-d.html>

Attachment 1D

[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.: sh2002000568

LCCN Permalink: <https://lcn.loc.gov/sh2002000568>

HEADING: Systems on a chip

000 00973cz a2200265n 450

001 5636724

005 20190825061713.0

008 020425|| anannbabn |a ana

010 __ |a sh2002000568

035 __ |a (DLC)5636724

035 __ |a (DLC)sh2002000568

035 __ |a (DLC)284544

040 __ |a DLC |b eng |c DLC |d DLC

053 _0 |a TK7874.888

150 __ |a Systems on a chip

450 __ |a SOC design

450 __ |a Systems on chip

550 __ |w g |a Embedded computer systems

670 __ |a Work cat.: 2002069402: Test resource partitioning for system-on-a-chip, 2002: |b CIP contents (design philosophy; SOC)

670 __ |a ASTI on FirstSearch, Apr. 25, 2002 |b (descriptor: system-on-a-chip)

670 __ |a WWW.us.design-reuse.com |b (article: Software key to SoC design; text: Systems on chip; presentation at Embedded Systems Conference)

906 __ |t 0221 |u te04 |v 1

910 __ |a Record saved by td04 on 01/29/2019 at 10:56:39

952 __ |a 10 bib. record(s) to be changed

953 __ |a jg12

< [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 1E

[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 87005579**LCCN Permalink:** <https://lcn.loc.gov/sh87005579>**HEADING:** Design and construction**000** 00664cz a2200169n 450**001** 4810729**005** 20110729124740.0**008** 870902n| bnannbbbn |n ann**010** __ |a sh 87005579**035** __ |a (DLC)sh 87005579**035** __ |a (DLC)157414**040** __ |a DLC |c DLC |d DLC**150** __ |a Design and construction**260** __ |i subdivision |a Design and construction |i under individual examinations and types of examinations and under types of structures, machines, equipment, etc., for works discussing their engineering, and/or construction, e.g. |a Automobiles--Design and construction**906** __ |t 8846 |u fk03 |v 0**952** __ |a 0 bib. record(s) to be changed**953** __ |a fj26 |b fj26< [Previous](#)[Next](#) >**Save, Print or Email Records (View [Help](#))**Select Download Format:

Email Text (Labelled Display) to:

 [Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)**Library of Congress**
URL: <https://www.loc.gov/>**Mailing Address:**
101 Independence Ave,**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](#)

 S.E.
Washington, DC 20540

Attachment 1F

[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 85029513

LCCN Permalink: <https://lcn.loc.gov/sh85029513>**HEADING:** Computer networks

000 01442cz a2200409n 450

001 4681787

005 20120323071047.0

008 860211i| anannbabn |b ana

010 __ |a sh 85029513

035 __ |a (DLC)sh 85029513

035 __ |a (DLC)5721786

035 __ |a (DLC)sp 85029513

035 __ |a (DLC)288027

035 __ |a (DLC)7651246

035 __ |a (DLC)409258

040 __ |a DLC |c DLC |d DLC |d DGPO |d DLC

053 _0 |a TK5105.5 |b TK5105.9

150 __ |a Computer networks

360 __ |i subdivision |a Computer networks |i under subjects

450 __ |a Communication systems, Computer

450 __ |a Computer communication systems

450 __ |a Data networks, Computer

450 __ |a ECNs (Electronic communication networks)

450 __ |a Electronic communication networks

450 __ |a Networks, Computer

450 __ |a Teleprocessing networks

550 __ |w g |a Data transmission systems

550 __ |w g |a Digital communications

550 __ |w g |a Electronic systems

550 __ |w g |a Information networks

550 __ |w g |a Telecommunication

550 __ |a Cyberinfrastructure

550 __ |a Electronic data processing |x Distributed processing

550 __ |a Network computers

670 __ |a U.S. Cong. Senate. Comm. on banking, Housing and Urban Affairs. Subcom. on Securities. The changing face of capital markets and the impact of electronic communications networks (ECNs), hearing ... 2000.

906 __ |t 0839 |u te04 |v 1

953 __ |a xx00 |b ta39

◀ Previous Next ▶

Save, Print or Email Records (View Help)	
Select Download Format:	Text (Labelled Display) v 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](#)

Exhibit 1G

Networks on Chip

Axel Jantsch and Hannu Tenhunen (Eds.)



Kluwer Academic Publishers

Networks on Chip

edited by

Axel Jantsch

Royal Institute of Technology, Stockholm

and

Hannu Tenhunen

Royal Institute of Technology, Stockholm



KLUWER ACADEMIC PUBLISHERS

BOSTON / DORDRECHT / LONDON

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN 1-4020-7392-5

Published by Kluwer Academic Publishers,
P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

Sold and distributed in North, Central and South America
by Kluwer Academic Publishers,
101 Philip Drive, Norwell, MA 02061, U.S.A.

In all other countries, sold and distributed
by Kluwer Academic Publishers,
P.O. Box 322, 3300 AH Dordrecht, The Netherlands.

Printed on acid-free paper

All Rights Reserved

© 2003 Kluwer Academic Publishers, Boston
No part of this work may be reproduced, stored in a retrieval system, or transmitted
in any form or by any means, electronic, mechanical, photocopying, microfilming, recording
or otherwise, without written permission from the Publisher, with the exception
of any material supplied specifically for the purpose of being entered
and executed on a computer system, for exclusive use by the purchaser of the work.

Printed in the Netherlands.

Contents

Preface	vii
Part I System Design and Methodology	
1 Will Networks on Chip Close the Productivity Gap? <i>Axel Jantsch and Hannu Tenhunen</i>	3
2 A Design Methodology for NoC-based Systems <i>Juha-Pekka Soininen and Hannu Heusala</i>	19
3 Mapping Concurrent Applications onto Architectural Platforms <i>Andrew Mihal and Kurt Keutzer</i>	39
4 Guaranteeing The Quality of Services in Networks on Chip <i>Kees Goossens, John Dielissen, Jef van Meerbergen, Peter Poplavko, Andrei Rădulescu, Edwin Rijpkema, Erwin Waterlander and Paul Wielage</i>	61
Part II Hardware and Basic Infrastructure	
5 On Packet Switched Networks for On-chip Communication <i>Shashi Kumar</i>	85
6 Energy-reliability Trade-off for NoCs <i>Davide Bertozzi, Luca Benini and Giovanni De Micheli</i>	107
7 Testing Strategies for Networks on Chip <i>Raimund Ubar and Jaan Raik</i>	131

vi		
8		
Clocking Strategies for Networks on Chip		153
<i>Johnny Öberg</i>		
9		
A Parallel Computer as a NoC Region		173
<i>Martti Forsell</i>		
10		
An IP-Based On-Chip Packet-Switched Network		193
<i>Ilkka Saastamoinen, David Sigüenza-Tortosa and Jari Nurmi</i>		
Part III Software and Application Interfaces		
11		
Beyond the von Neumann Machine: Communication as the Driving Design Paradigm for MP-SoC from Software to Hardware		217
<i>Eric Verhulst</i>		
12		
NoC Application Programming Interfaces		239
<i>Zhonghai Lu and Raimo Haukilahti</i>		
13		
Multi-level Software Validation for NoC		261
<i>Sungjoo Yoo, Gabriela Nicolescu, Iuliana Bacivarov, Wassim Youssef, Aimen Bouchhima and Ahmed A. Jerraya</i>		
14		
Software for Multiprocessor Networks on Chip		281
<i>Miltos Grammatikakis, Marcello Coppola and Fabrizio Sensini</i>		

NETWORKS ON CHIP

Preface

During the 1990s more and more processor cores and large reusable components have been integrated on a single silicon die, which has become known under the label *System on Chip (SoC)*. Main difficulties of this era were, and still are, the standardization of the component interfaces and the validation of the entire system with respect to its physical and functional properties. Buses and point to point connections were the main means to connect the components. Buses are attractive because they provide high performance interconnections while they can still be shared by several communication partners. Hence they can be used very cost efficiently.

As silicon technology advances further, several problems related to buses have appeared. Buses can efficiently connect 3-10 communication partners but they do not scale to higher numbers. Even worse, they behave very unpredictably as seen from an individual component, because many other components also use them. A second problem comes from the physics of deep submicron technology. Long, global wires and buses become undesirable due to their low and unpredictable performance, high power consumption and noise phenomenon. A third problem comes from the application perspective. Designing and verifying the inter-task communication in a system is a hard problem per se. Getting it to work and dimensioning communication resources correctly is even harder for large bus based communication networks due to the unpredictability of the communication performance. Moreover, every system has a different communication structure, making reuse difficult.

As a consequence, around 1999 several research groups have started to investigate systematic approaches to the design of the communication part of SoCs. It soon turned out that the problem has to be addressed at all levels from the physical to the architectural to the operating system and application level. Hence, the term *Network on Chip (NoC)* is today used mostly in a very broad meaning, encompassing the hardware communication infra-structure, the middleware and operating system communication services and a design methodology and tools to map applications onto a NoC. All this together can be called a *NoC platform*. The breadth of the topic is also highlighted by the scope of this book which ranges from physical issues to embedded software. Quite natural

for a young and quickly evolving research area, the terminology is not yet uniformly used and different authors use the terms differently, as it will become apparent when reading this book.

It should not come as a surprise that the **first part** deals with design and methodology issues. The infamous design productivity gap is one of the strongest, if not the single most important, driving force towards design, architecture and implementation structures. Only when we succeed to restrict the design space in a sensible way will we be able to exploit technology potential. Several chapters of the first part emphasize predictability of the design process and the guarantee of high level features by all implementations.

The **second part** is concerned with the hardware infrastructure. The network topology, power management, fault tolerance, testing and clocking, among other topics, are all key issues that must be solved satisfactorily to make NoCs feasible.

Software and the application perspective is in the center of **part three**. Not surprisingly, communication services and the role of the operating system in future NoC systems are central in the chapters of this part.

Although this book touches upon most of the important NoC issues, many are only superficially dealt with and some key issues are not addressed. NoC is a young and emerging area and we still have to learn to assess the quality of a particular solution, be it for the topology, switch design, communication or operating system services, with respect to an application or application area. We expect for the near future that NoC specific cost and performance will be developed that may be application sensitive. An interesting question is for instance, how to express the communication performance of a NoC. Raw bandwidth may not be adequate for applications with a highly variable traffic pattern consisting of a mixture of real-time control messages and high throughput video streams. In addition, transient faults may occasionally destroy transmitted information, which may make a fault management and retransmission scheme necessary. But again, this will depend on the sensitivity of the application data. Hence, efficiency of a particular NoC system will be more and more expressed in relation to an application or an application area. Existing benchmarks can be used to address this question but new and modified benchmarks will also be required.

As it is indicated frequently in this book, NoC could lead to a fundamental paradigm shift with respect to the way we develop platforms, we design systems and we model applications. At least it will result in a scalable platform architecture for the billion transistor chip era. We expect in any case the NoC area to flourish and prosper and take unexpected and innovative turns and directions and we hope that this book contributes to this exciting research theme.

AXEL JANTSCH AND HANNU TENHUNEN

I

SYSTEM DESIGN AND METHODOLOGY

Chapter 1

WILL NETWORKS ON CHIP CLOSE THE PRODUCTIVITY GAP?

Axel Jantsch and Hannu Tenhunen

Royal Institute of Technology, Stockholm

axel@imit.kth.se, hannu@imit.kth.se

Abstract We introduce two properties of the design process called the *arbitrary composability* and the *linear effort properties*. We argue that a design paradigm, which has these two properties is scalable and has the potential to keep up with the pace of technology advances. Then we discuss some of the trends that will enforce significant changes on current design methodologies and techniques. Finally, we argue that the emerging Network-on-Chip (NoC) paradigm promises to address these trends and challenges and has all prerequisites to provide the arbitrary composability and the linear effort properties. Consequently we conclude that NoC is a likely basis for future System-on-Chip platforms and methodologies.

Keywords: Networks on chip, Productivity gap, System on chip design methodology

1. Introduction

To boost design productivity it is crucial that the effort to add new parts to a given design does not depend on the size of the existing design but only on the size of the new parts. In other words, the design effort must be a linear function of the size of the new parts. If this is the case, large parts and blocks of previous designs can be reused and the design effort can be invested into the new parts. This is also a necessary prerequisite to provide a solid methodology, architecture, and thus a platform, that are sustainable over several technology generations.

The central thesis of this chapter is that a Network-on-Chip (NoC) has the potential to provide such a sustainable platform and, if successful, will incur such a significant change on the system-on-chip architecture

3

A. Jantsch and H. Tenhunen (eds.), *Networks on Chip*, 3–18.
© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

and design process that it can be called a paradigm change. On the other hand, if it fails to do so, NoC will be just one of several architectures and platforms available to embedded system designers.

Arbitrary composability property: *Given a set of components and a set of combinator operators which allow to connect and integrate the components into larger component assemblages. Components and combinators together are arbitrarily composable if a given component assemblage A can be extended with any component by using any of the combinators without changing the relevant behavior of A.*

Please note, that this and the following property are meant as engineering heuristics, not as mathematical properties. As such they are ideals and can be achieved at higher or lower degrees.

Note further, that this property is defined with respect to what is considered to be a *relevant behavior*. Thus depending on the given objectives and definition of behavior, the same components and combinators may or may not have the arbitrary composability property.

For instance, the standard logic gates NAND, NOR, INV, etc. have this property with respect to their logic level I/O behavior because adding new gates to a netlist of gates will not change the behavior of the original netlist, unless old connections are broken. A given network of gates can be used in any context and will exhibit identical behavior whatever the surrounding netlist may be. New gate netlists can be added to existing ones, using the outputs and results produced by any other part of the circuit without changing the older parts. This is the foundation of our ability to build designs with millions of gates and to reuse large blocks in arbitrary environments.

It should be noted that this nice property of gates is in part due to the implementation process which allows the scaling of transistor sizes, insertion of buffers, and sensible placement and routing by automatic tools.

It is enlightening to see the effects when the arbitrary composition property is violated. Two of the most severe problems in today's designs stem from violations of this property. Timing closure, i.e. the problem to get the timing of the circuit implementation right, is difficult because small changes or addition to the gate netlist may change the timing of the entire system by adding to the critical path or due to an unexpected effect of placement and routing on the timing of seemingly unrelated circuit parts. The system verification problem is so hard because at the system level behaviors are not easily composable and tiny changes in one part may have unexpected effects on seemingly unrelated other parts of

the system. In both cases the design effort grows more than linear with the system size.

If this property is guaranteed, the effort of adding new components to a working system only depends on the new components but not on the size of the reused system (figure 1.1). Thus, a corollary of the arbitrary




Figure 1.1. With the arbitrary composability property the design effort to add new components to an existing system depends on the integration effort and on the new components but not on the size of A because the design effort to build A has already been spent and is reused as well: $\text{Deffort}(S) = \text{Deffort}(B) + \text{Deffort}(C) + \text{Ieffort}(3)$

composability property is the following linear effort property.

Linear Effort Property: *Given is a set of components and a set of combinator operators which allow to connect and integrate the components into larger component assemblages. A design process which builds a system from the components and combinators has the linear effort property if a given set of n assemblages $A_1, \dots, A_n$ can be integrated into a system S by means of the combinators with an effort dependent on n but not on the size of the assemblages: $\text{Ieffort}(n)$. Thus the total design effort for S is*

$$\text{Deffort}(S) = \text{Deffort}(A_1) + \dots + \text{Deffort}(A_n) + \text{Ieffort}(n)$$

Note, that this property implies that the interface complexity of an assemblage does not depend on the size of the assemblage. Obviously, this is not true in practice but it is equally obvious that this is a necessary precondition to build arbitrary large systems. Thus, we must approach this ideal as close as possible to be able to build larger and larger systems. The fact, that we have not been sufficiently close to this ideal is the fundamental reason for the design productivity gap.

We believe that NoC based platforms have a good potential to provide both the arbitrary composition and the linear effort properties to a high degree but they do not automatically guarantee them. We will keep these properties in mind throughout this chapter, but first we review

some of the underlying trends and challenges that lead to NoC and similar architectures.

2. Trends and Challenges

IC manufacturing technology will provide us with a few billion transistors on a single chip within a few years [1]. Assuming that these predictions hold and that the market will continue to absorb ever higher volumes of ICs, the key questions are: how will the future chips be organized and how will future systems, which include these chips, be designed? There are a few trends which, if continued, will bring about a significant change for architecture and design of integrated circuits.

Communication versus computation. Technology scaling works better for transistors than for interconnecting wires. This leads gradually to a domination of performance figures, power consumption and area by wires and make transistors of secondary importance. At the system level it has a profound effect by changing the focus from number crunching and computation to data transport and communication [2, 3, 4]. Communication becomes often the bottleneck because it seems much harder to design and get right.

Deep submicron effects. Cross-coupling, noise and transient errors are only some of the unpleasant side-effects of technology scaling [5, 6]. It requires significant skills, experience, knowledge and time to keep them under control while exploiting the limits of a technology. A digital or system designer with an expected design productivity of millions of transistors per day is not able to deal with these effects properly. Therefore, designers reuse blocks which are carefully designed by experts with the proper skills. However, it is of critical importance that the deep submicron effects don't pop up again when predesigned blocks are combined in arbitrary ways. Consequently, at the physical design level the property of arbitrary composability means that the electrical and physical properties of blocks are not affected when combining them.

Global synchrony. Physical effects of deep sub-micron technology make it increasingly difficult to maintain global synchrony among all parts of the chip [7]. The clock signal will soon need several clock cycles to traverse the chip, clock skew becomes unmanageable, and the clock distribution tree is already today a major source of power consumption and cost. The trends of scaling to smaller geometric dimension and higher clock frequency make these problems more significant every year.

Thus, it is unlikely that large chips will be synchronous designs with only one clock domain.

Design productivity gap. Synthesis and compiler technology development do not keep pace with IC manufacturing technology development [1]. As a consequence we need either exponentially growing design teams or design time to design and implement systems which fit onto a single IC. Since both alternatives are unrealistic we have in the past escaped from the problem by using ever more complex components as primitive design units. These primitive design units have evolved from individual transistors to logic gates to entire ALUs, multipliers, finite state machines and processor cores. In fact *reuse* of ever more complex design elements has been the main device to increase productivity and will likely remain so in the years to come. It has also kept us close enough to the ideal of the linear effort property to manage the increasing number of transistors.

Heterogeneity of functions. Obviously, systems that can be implemented on a single chip become increasingly more complex. As a result different functions and features with vastly different characteristics and history reside on the same chip. Signal processing algorithms, that recover and generate radio signals, will coexist with global control, maintenance and accounting functions as well as with natural language comprehension and generation functions. These functions are developed in different contexts, by different teams, with different design languages and tools. However, they need to be integrated into a single chip.

In order to lead a concrete discussion we describe next a typical NoC architecture and in section 4 we investigate how a NoC approach could address the listed problems. To paint a fair picture section 5 illuminates the price to pay when adopting a NoC based approach. Finally, in section 6 we speculate how the design process will change.

3. Network on Chip

The Network-on-Chip (NoC) architecture, as outlined in figure 1.2, provides the communication infrastructure for the resources. In this way it is possible to develop the hardware of resources independently as stand-alone blocks and create the NoC by connecting the blocks as elements in the network. Moreover, the scalable and configurable network is a flexible platform that can be adapted to the needs of different workloads, while maintaining the generality of application development methods and practices.




Figure 1.2. Each node in the mesh contains a switch and a resource.

A two dimensional mesh interconnection topology is simplest from a layout perspective and the local interconnections between resources and switches are independent of the size of the network. Moreover, routing in a two-dimensional mesh is easy resulting in potentially small switches, high bandwidth, short clock cycle, and overall scalability. A NoC consists of resources and switches that are directly connected such, that resources are able to communicate with each other by sending messages. A resource is a computation or storage unit. Switches route and buffer messages between resources. Each switch is connected to four other neighboring switches through input and output channels. A channel consists of two one-directional point-to-point buses between two switches or a resource and a switch. Switches may have internal queues to handle congestion. We expect that the area of a resource is either the maximal synchronous region in a given technology or a cluster of computing elements and memory connected via a bus. We expect the size of a resource to shrink with every new technology generation. Consequently the number of resources will grow, the switch-to-switch and the switch-to-resource bandwidth will grow, but the network wide communication protocols will be unaffected.

Note, that this is one, fairly simple NoC variant described here only to substantiate the following discussion. Many other more sophisticated architectures have been proposed [8, 9, 10, 11, 12, 13, 14].

4. How does NoC address the problems

Although there are numerous factors and facets to investigate, we focus on two mechanisms that work in favor of a NoC based approach: *Reuse* and *Predictability*. We believe these are by far the most important factors overshadowing other secondary effects.

4.1 Reuse

As mentioned above, reuse has always been the primary means to bridge the technology gap [15, 16, 17]. More and more complex components, from transistors to gates to functional blocks, such as ALUs and multipliers, to microprocessors and DSP cores, have become the primitive building blocks. In this way, the designers could move up the abstraction levels and describe the system's functionality in more and more abstract terms relying on more and more powerful "primitive" components. Curiously, synthesis technology has mostly been used to bridge one, relatively shallow, level of abstraction. This can be observed in both the hardware and the software domain. Technology mapping, logic and RTL synthesis are in principle straight forward steps with a well characterized optimization space. Synthesis techniques that attempted to bridge more abstraction levels and to make more profound design decisions have typically failed. Examples are high level and architectural synthesis. In software the mapping from C to microprocessor instruction sets is more or less direct and the corresponding compiler technology is tremendously successful. In contrast, compilation from functional, logic and other higher level languages lacks efficiency and has never become mainstream. Thus, we believe that reuse by providing more complex components will continue to be the main mechanism to exploit the potential of technology. Synthesis and compilation techniques will provide the surface mapping from more convenient descriptions onto the primitive components.

However, as a difference to the past, communication "components", or better, communication structures and services have also to become primitive design elements. And this is precisely what a NoC based approach is all about: The reuse of communication services.

Let us briefly review what can be reused in a NoC based approach.

Components and resources. Arbitrary computation elements can be connected to the communication network. In fact we expect that typical NoC based systems may contain processor cores, DSP cores, memory banks, specialized I/O blocks such as Ethernet or Bluetooth protocol stack implementations, graphics processors, FPGA blocks, etc. The size of a resource can range from a bare processor core to a local cluster of several processors and memory connected via a local bus. Resources have to comply to the interfaces of the communication network in order to connect to it and use its services.

The reuse of processor cores has been developed during the last ten years by defining bus interfaces. By defining network interfaces, NoC takes this concept further because it allows to integrate an arbitrary number of resources into a network. In a bus-based system, adding a new resource has a profound impact on the performance of the rest of the system because the same communication resource is now shared among more resources. In a NoC adding new resources also means to add new communication capacity by adding new switches and interconnects. This scalability property is a necessary precondition for the arbitrary composability property but it is not sufficient to guarantee it. The communication network must further be able to guarantee allocated bandwidth and to enforce a decent behaviour of the resources to avoid, for instance the monopolization of the entire communication bandwidth by a single resource.

However, it is important to acknowledge, that a NoC based approach has the potential to provide the arbitrary composability property with tremendous benefits for the design productivity.

Communication infrastructure. The main immediate benefit from a NoC based approach is clearly due to the possibility to reuse the communication network. The switches, the interconnects and the lower level communication protocols can be design, optimized, verified and implemented once and reused in a large number of products. If requirements on performance, reliability and cost differ too much in different application domains, domain specific communication networks can be reused at least for all products in the same domain. For instance, it is likely that mobile, hand-held devices have so different demands on power consumption and performance than infra-structure equipment that different NoC platforms for these two domains are well justified.

Apart from the hardwired communication infra-structure, higher level network services can as well be reused. There is in fact a long list of services that would benefit many applications but can impossibly be developed from scratch for each new product. Examples are

- the detection, monitoring and management of faults in the network;
- the allocation and management of network resources and possibly task migration for load balancing and power optimization;
- the management of global and shared memory;
- the provision of sophisticated communication services such as channels with guaranteed bandwidth and quality of service, multi-cast and broadcast communication, etc.

Most of these tasks are typically provided by the operating system in today's uniprocessor applications. Similarly, a NoC operating system will be very generic and can be reused for many products. Due to the increased complexity of future systems as compared to today's systems, the operating system will be much more sophisticated and complex, thus making the case for reuse even stronger.

Application parts and feature reuse. Reuse will not stop with components and generic services. New products can be composed of existing, complete features. Peeking into the future we can envision a traditional mobile phone which is enhanced by speech analysis subsystem, a speech synthesizer and a language analysis and processing sub-system to provide a spoken language interface to the phone. The same modules or features are apparently useful in a wide range of products and should therefore be reused as much as possible. Obviously, the main challenge will be to define and standardize the high level interfaces between these features to allow for an efficient communication and sharing of information. However, we can observe that a NoC provides an excellent ground for this kind of feature reuse. For optimized implementation a feature may come fully implemented either in software or partially as a dedicated hardware block. Either way, the feature can be plugged into one or several resource slots and the NoC provides at least the low level communication and network services for free for the interaction between the feature and the rest of the system.

Design, simulation and prototype environment. A significant part of system development costs are typically spent in setting up simulation environments and building prototypes. Since many products are based on the same NoC platform much of this investment can be shared by many products. Furthermore, even if different domain specific NoC platforms vary in their performance and power characteristics, they are

sufficiently similar to allow the reuse of much of the design and verification environment across very different application domains. In contrast, application specific platforms which are not derived from the same principle concepts but are developed in an ad-hoc way to suit a particular application domain, will not provide as much potential for cross-domain reuse.

Verification effort. The system verification effort is frequently as high as the design effort itself. Hence, by reusing predesigned and preverified parts and services, verification time can be as drastically reduced as design time. But reuse is even more important for the verification than for the design activity because the uncertainty about the required verification time is much higher and more difficult to plan. Moreover, the uncertainty about the resulting product quality is high and the potential cost of undetected errors in the final product can be enormous. Since risk and uncertainties around verification and verification effects are much higher, verification benefits more from reuse than the design activities. Reused components are typically much better verified, because they have already been used in other products. Moreover, system verification can be done much more effectively when correctness and reliability of the components can be assumed, because errors are identified faster.

In summary, the usage of a NoC based platform boosts the potential for reuse in many ways.

4.2 Predictability

The second main aspect of our focus is predictability and it is in fact closely related to reuse.

Communication performance. Due to its regular geometry and communication network, communication performance becomes potentially much more predictable. "Potentially" and not "necessarily" because the regular communication hardware will significantly help to analyze and assess performance but measures at higher protocol levels and network services have to realize this potential. Since the communication hardware is shared by many resources, the activity of one resource can delay the communication of other resources. One mis-behaving resource can monopolize part of the communication hardware and indefinitely block other resources from accessing it. Therefore, the network has not only to provide the raw bandwidth but also a policy for bandwidth allocation. There are many ways to do this with different advantages and disadvantages. One possibility is to allocate channels with a maxi-

mum bandwidth and fixed latency to two communicating resources. By properly allocating and managing these channels many communication activities can coexist peacefully and all of them know exactly the available bandwidth and what latency each data packet will exhibit. This makes systemwide performance analysis feasible and accurate.

As indicated above, a reliable communication resource allocation policy resulting in predictable communication performance for all tasks and applications in the NoC is a central part for establishing the property of arbitrary composability. If a task or a feature can request, obtain and use communication bandwidth independent of all other tasks and features in the NoC, then the cost of reusing, mixing and matching existing tasks and features is relatively small, since they have not to be modified internally.

Electrical properties. Due to the regularity of the layout the electrical and physical properties of the network are well known. The switch-to-switch wires, most likely the longest on the chip except for clock, power and ground wires, have all exactly the same well defined length. Potential uncertainties and irregularities are confined to the resource slots and have most likely no global impact. Thus, the regular and known geometry leads to the accurate analysis of electrical properties. Furthermore, since the NoC platform is reused in many products, elaborate electrical and power models can be developed e.g. to model the dynamic power consumption caused by the communication in the network. High level traffic pattern models can be combined with electrical models for the better assessment of noise and power consumption.

Design and verification time. As already elaborated above, reuse will decrease uncertainties and risk in particular for the verification tasks. This naturally makes design and verification time more predictable. Another reason for increased predictability of the design process is reduced design freedom. Since the network structure and basic services are fixed as well as the size of resource slots, the design space is significantly reduced and the remaining tasks are well separated resulting in more smaller and independent tasks. The remaining system level architectural decisions are the selection of the platform, the network size and the resource allocation. Then all the resources can be designed and implemented separately. Hence, the main challenge is the development of the overall system functionality, its partitioning into processes and their mapping onto resources. However, due to the relatively fixed architecture and the predictable performance and electrical properties (see above) the behavioural design is to a large extent independent from

the implementation phase. This increased modularization of the design tasks makes the overall process more predictable.

5. Disadvantages

Obviously, all these nice properties, outlined above, do not come for free. Essentially we pay for each of them by losing optimality. Reusing components always means that we use something more general, thus less optimal, than necessary for a particular task. Adopting a fixed and relatively inflexible network topology means that we exclude all the other topologies that may be more suitable for the problem at hand. A particular protocol stack with associated services will not be ideal for all embedded systems and any NoC operating system will often be too general with unnecessary overhead and lacking important services at the same time. Only time will tell if it is possible to find a sensible balance between generality and efficiency which fits sufficiently many applications.

One way to alleviate the problem may be offered by a layered protocol stack and a layered set of services. An application may select a NoC platform only up to a particular level of the stack, thus avoiding the overhead of the upper layers. This flexibility however comes at a price. A strict layering of protocols and services may not be the optimal solution. An integrated implementation of several layers may result in a more efficient design at the expense of the possibility to have direct access to the lower layers.

6. Effect of NoC on the design process

Let us briefly review a design process based on a NoC platform. The main activities are:

Configuring the platform. Depending on the available configuration options this phase may be extensive or insignificant. A more general platform suitable for a wider range of applications will have more configuration options.

Every platform will allow to select the size and most likely offer a selection of communication services that can be included. Every design will need the core service for transporting raw data from sending to receiving resources. This core service may be packet oriented or based on virtual channels. More sophisticated communication means with varying levels of fault tolerance, bandwidth and latency control can be optional. Still higher level services could support transparent task migration from one resource to another or a virtual global shared memory.

Selecting resources. Since every resource slot can receive an arbitrary resource, all the resources have to be chosen by the designer. There may be some constraints because network services like a NoC operating system have to be implemented also. Part of the resource selection task is the design of the memory architecture and I/O architecture, which require specific resources distributed in a particular way.

Reuse of features. The selection of resources is of course intimately connected to the reuse of features. In our terminology a feature is a particular functionality together with its implementation. E.g. a speech analysis feature can be modeled as a task graph and implemented on one or several DSP processors or custom hardware blocks. Thus, a feature can occupy several resource slots but it can also share a resource, e.g. a DSP, with other features. Resource sharing of features has to be considered carefully because it will most likely compromise the property of arbitrary composability unless precautions are taken.

Evaluation and integration. The integration of all features and resources into the final system, the evaluation of performance and the verification of functionality is the final and perhaps most challenging task. An efficient and elaborate simulation and prototyping environment, that can be shared by many design projects, will significantly aid the successful system integration. But above all, the simplicity and predictability of the interfaces of resources, the communication network and features at all levels of the protocol stack will determine how well the property of linear effort can be approximated.

7. Conclusion

Developing a system with several dozens or hundreds of processor like resources is a formidable task. It can only be accomplished if features and components are aggressively reused and if they are *arbitrarily composable* at the physical level, at the architectural and structural level and at the application level (figure 1.3). Or, in other words, the protocol stack from the physical layer to the network and transport layer to the application layer must be well defined such that arbitrary components and features can be connected to the NoC backbone via the protocol stack without affecting the rest of the system.

We have focused on two issues which are particularly crucial for providing the properties of arbitrary composition and linear effort. *Reuse* from blocks to services and features allows designers to productively combine existing parts to new configurations and innovative products. *Predictability* makes reuse efficient because adding new components and




Figure 1.3. Features and components must be arbitrarily composable at the physical, the architectural and the application level to facilitate seamless feature integration.

features does not require redesigning and reverifying the existing parts, thus avoiding a more than linear growth of the design effort.

If NoC based platforms are able to provide the arbitrary composability and the linear effort properties while still allowing for sufficiently efficient product implementations, the design of systems-on-chip will be revolutionized, eventually leading to a "design by feature combination". The NoC paradigm is highly suited to provide SoC platforms scalable and adaptable over several technology generations because it opens the opportunity to define and standardize a communication service infrastructure and a protocol stack from the physical to the application layer. If these are well defined and facilitate the arbitrary composability and the linear effort properties, it will allow the efficient reuse of large resources, communication and network services, and application features. The TCP/IP protocol stack and the Internet, which also revolutionized the use of computers, can be an inspiring example. Indeed, due to their inherent scalability NoC platforms may allow the design productivity to grow as fast as technology capabilities and may eventually close the design productivity gap.

References

- [1] Semiconductor Industry Association. International technology roadmap for semiconductors. Technical report, World Semiconductor Council, 1999. Edition 1999.
- [2] James A. Rowson and Alberto Sangiovanni-Vincentelli. Interface-based design. In *Proc. of the 34th Design Automation Conference*, 1997.
- [3] Kurt Keutzer, Sharad Malik, Richard Newton, Jan Rabaey, and Alberto Sangiovanni-Vincentelli. System-level design: Orthogonalization of concerns and platform-based design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(12):1523–1543, December 2000.
- [4] Marco Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, and A. Sangiovanni-Vincentelli. Addressing the system-on-a-chip interconnect woes through communication-based design. In *Proceedings of the 38th Design Automation Conference*, June 2001.
- [5] Dennis Sylvester and Kurt Keutzer. Getting to the bottom of deep submicron. In *Proceedings of the 1998 IEEE/ACM international conference on Computer-aided design*, pages 203–211, 1998.
- [6] Dennis Sylvester and Kurt Keutzer. Getting to the bottom of deep submicron ii: a global wiring paradigm. In *Proceedings of the 1999 International Symposium on Physical Design*, pages 193–200, 1999.
- [7] Thomas Meincke, Ahmed Hemani, S. Kumar, P. Ellervee, J. Öberg, T. Olsson, P. Nilsson, D. Lindqvist, and H. Tenhunen. Globally asynchronous locally synchronous architecture for large high performance ASICs. In *Proc. of IEEE Int. Symp. on Circuits and Systems (ISCAS)*, volume II, pages 512–515, Orlando, USA, May 1999.
- [8] Pierre Guerrier and Alain Greiner. A generic architecture for on-chip packet-switched interconnections. In *Proceedings of Design, Automation and test in Europe*, pages 250–256, 2000.
- [9] Shashi Kumar, Axel Jantsch, Juha-Pekka Soininen, Martti Forsell, Mikael Millberg, Johnny Öberg, Kari Tiensyrjä, and Ahmed Hemani. A network on chip architecture and design methodology. In *Proceedings of IEEE Computer Society Annual Symposium on VLSI*, April 2002.
- [10] William J. Dally and Brian Towles. Route packets, not wires: On-chip interconnection networks. In *Proceedings of the 38th Design Automation Conference*, June 2001.

- [11] Edwin Rijpkema, Kees Goossens, , and Paul Wielage. A router architecture for networks on silicon. In *Proceedings of Progress 2001, 2nd Workshop on Embedded Systems*, October 2001.
- [12] Drew Wingard. MicroNetwork-based integration of SOCs. In *Proceedings of the 38th Design Automation Conference*, June 2001.
- [13] K. Goossens, J. van Meerbergen, A. Peeters, and P. Wielage. Networks on silicon: Combining best-effort and guaranteed services. In *Proceedings of the Design Automation and Test Conference*, March 2002.
- [14] Michael Bedford Taylor, Jason Kim, Jason Miller, David Wentzlaff, Fae Ghodrat, Ben Greenwald, Paul Johnson Henry Hoffman, Jae-Wook Lee, Walter Lee, Albert Ma, Arvind Saraf, Mark Seneski, Nathan Shnidman, Matt Frank Volker Strumpfen, Saman Amarasinghe, and Anant Agarwal. The Raw microprocessor: A computational fabric for software circuits and general-purpose programs. *IEEE Micro*, 22(2):25–35, March/April 2002.
- [15] Michael Keating and Pierre Bricaud. *Reuse Methodology Manual for System-on-Chip Designs*. Kluwer Academic Publishers, 1998.
- [16] Terry Thomas. Technology for ip reuse and portability. *IEEE Design & Test of Computers*, 16(4):6–15, October 1999.
- [17] Henry Chang, Larry Cooke, Merrill Hunt, Grant Martin, Andrew McNelly, and Lee Todd. *Surviving the SOC Revolution - A Guide to Platform-Based Design*. Kluwer Academic Publishers, 1999.

Chapter 2

A DESIGN METHODOLOGY FOR NOC-BASED SYSTEMS

Juha-Pekka Soininen and Hannu Heusala

VTT Electronics (Technical Research Centre of Finland), P.O. Box 1100, Kaitoväylä 1, FIN-90571 Oulu, FINLAND

University of Oulu, Department of Electrical and Information Engineering, P.O. Box 4500, FIN-90014 University of Oulu, FINLAND

Abstract: Diversity of computational requirements of information technology products will increase. The products will be based on computers, but the full exploitation of silicon capacity will require improvements in design productivity and system architectures. The chapter presents a backbone-platform-system methodology for development of products that are based on a network-on-chip concept. Encapsulation of dedicated subsystems into a NOC system is the key idea of an "integrated distributed embedded system" approach. The cornerstones of the methodology are the reuse of a NOC backbone and a platform architecture. The NOC backbone encapsulates the physical integration problems into communication solution for system architect. The NOC platform provides application area optimized computation and storage resources for system developers. Essential parts of the methodology are decision-support methods for both NOC platform and NOC system designers.

Key words: Design methodology, platform-based design, reuse, decision-support

1. INTRODUCTION

Exploitation of a billion-transistor capacity of a single ASIC requires new system paradigms and significant improvements to design productivity [1].

19

A. Jantsch and H. Tenhunen (eds.), *Networks on Chip*, 19–38.
© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

The Network-on-Chip types of approaches presented in earlier chapters try to divide and conquer the use of available silicon surface. Integration of tens or hundreds of computers into a single chip together with asynchronous message passing network partitions the design into more manageable units and allows to optimize the subsystems according to application requirements [2]. Such scalable and modular architectures are natural steps in system architecture evolution especially when physical implementations are considered.

Structural complexity and functional diversity of such systems are the challenges for the design teams. Simple extrapolations of design productivity roadmaps reveal that already in the near future only the very largest companies can complete their ASIC projects in case the full silicon capacity is used. In the longer period, it is clear that design capacity is going to be the limiting factor for the system complexity.

Structural complexity can be increased by having more productive design methods and by putting more resources in design work. Design method improvement using higher abstraction levels leads to the more complicated and expensive design space exploration and synthesis tasks. When increasing the number of designers, the communication overheads consume almost all the benefits. The most attractive approach is the reuse of earlier designs, because if the design is reused at least three times the cost of making design reusable is justified. In platform-based design, the reuse is extended to architectures [3] and recent reconfigurable platforms duplicate regular structures of computational units [4].

Increasing number of gates can be and will be used for increasing the functionality of the products. The billion-transistor capacity on silicon and operation frequency at gigahertz range means that we shall have teraoperations per second capacity for applications. In addition to software and computer hardware, the same system can have DSP functions, analog and radio frequency functions, advanced controlling functions, artificial intelligence, etc. The system level decisions must take into account all the different needs of application types. It will require common language, concepts and terminology. The optimization of the implementations of applications will mean the integration of a variety of memory technologies, analog electronics, reconfigurable technologies and computer technologies into same platform.

The structural complexity will force us to operate at the system level when developing NOC-based systems. It will not be possible to go into details when making decisions concerning reuse or application mappings, for example. It will also force us to develop platform-type of architectures and resources. We have to replace the target product requirements with more abstract concepts. We have to use requirements of the product area,

characteristics of typical applications, forecasts on near future applications, etc. This will mean the shift from synthesis or mapping based methods to analysis and estimation based approaches.

Design methodology is a coherent set of theories, methods, techniques and/or principles used to analyze and/or develop methods for a particular domain. The main purpose of the methodology is the partitioning of the design problem into manageable tasks and definition of the tools and practices for those tasks. When considering NOC-based systems, the first problem arises from the diversity of functionality. Most of the existing design methodologies are application area specific. However, these methodologies are results from long evolution and they have built-in knowledge and expertise that should not be ignored. The challenge is therefore to combine these different approaches.

This chapter presents the principles of a NOC design methodology. Since the NOC-based systems have not been designed or manufactured yet, our approach is based on the extrapolations and forecasts. Section 2 gives a brief review of most important existing and evolving design methods and tools. Section 3 describes the requirements of NOC design methodology in rapidly changing technological and economical environment. Section 4 describes the NOC architecture and quality characteristics of NOC. The backbone-platform-system design methodology for NOC-based systems is presented in Section 5. The novelties of the approach are in the layered structure of the methodology, in the integration of physical issues into architecture design, and attempt to manage system complexity by emphasizing the evolutionary nature of products and product development. The decision-making support required at early phases of design is the most challenging part in the presented approach. Finally, Section 6 summarizes this chapter.

2. DESIGN METHODOLOGIES FOR IC BASED SYSTEMS

The motivation behind technological development is the need to produce products that increase our welfare. Such products are needed, used, and profitable. The same basic need is also behind the technological development, which in turn enables us to introduce increasingly complex products. The problem is that the increasing complexity requires more advanced design methodologies, because complexity increases the number of steps necessary for building them.

Design methodology is dependent on the underlying technology, but it is also a cultural issue and a business issue. Sexes, professions, companies, nationalities have their background, history and differences, and they may

need different type of operational environments in order to perform optimally. Inside the companies, the cumulated knowledge, the existing products, and the strategic decisions affect to the design methodologies [5].

The design methodology is responsible for product success in a sense that it must guarantee that the company performs the activities that are needed in quality validation. The main quality factors are performance, cost and variability, but these are very generic and cannot be directly used in product development. Important features of methodology are the transformation of user needs to physical requirements for different abstraction layers and partitioning of those needs and requirements to subsystems and technologies. The quality function deployment (QFD) for example, tries to solve these problems by introducing correlation matrixes for quality characteristics and technical requirements [6].

2.1 Evolution of ASIC techniques and design methods

Since its invention over forty years ago, IC-technology has offered for designers increasing number of switching elements in planar form. Basic logic operations needed to be performed in the digital systems on the surface of a silicon chip are still AND, OR and NOT. So far, the functions of digital systems are based on the Boolean algebra and combinational logic blocks are constructed of those basic logic gates. Memory function is brought to the system by adding feedback connections to a combinational block.

Switching elements are manufactured on the surface of a silicon chip in a silicon foundry applying different available semiconductor manufacturing technologies. Nowadays, and most likely for next ten to fifteen years, it is usually CMOS technology, where every switching element consists of two complementary connected transistors. The function and elementary structure of SOC and NOC systems are still, like in the first digital ICs from early 60's, based on the CMOS switching elements implementing AND, OR and NOT Boolean functions.

Over forty years, the semiconductor industry has been able to develop new manufacturing technology generations. The increasing number of switching elements is offered for designers in different implementation formats. Therefore, we have general-purpose logic (GPL) devices, memories, microprocessors, micro-controllers, programmable logic devices (PLD), field programmable gate arrays (FPGA), ASICs, SOC and NOCs. Implementation format can be understood to be a form specifying how switching elements are grouped and connected on the chip surface. The format specifies also the customer's and the manufacturer's responsibility in circuit design process.

The GPL devices, e.g. 74/54 and 4000 families, were the first implementation format offered by the semiconductor industry in early 60's. Billions units of them are still used by the electronics industry every year for every kind of products. Schematic entry is most efficient design method with GPL format. GPL format is an example of reuse of silicon layout blocks in the form of SSI/MSI level designs of 10 to 100 gates each. If 50 to 100 pages of schematics is a manageable amount of documentation then a digital system consisting of 500 to 1000 SSI/MSI devices could be possible to be designed in proper time.

Soon after the introduction of GPL format, the semiconductor industry implemented also the first memory devices and then in 1971 the first microprocessor with 2300 transistor switches on a single chip. The combination of microprocessor, memory and GPL devices on the Printed Circuit Board has formed the heart of digital systems known as embedded systems [7]. Embedded Systems are an example of implementation format where the HW reuse is in maximum. Many different systems may have the same HW architecture. The designer personalizes the system by writing the SW code and compiling it into a binary pattern of memory devices. The same amount of man-hours is needed to write 50 to 100 pages of SW code than 50 to 100 pages of schematics. One line of the SW code corresponds to a GPL device in schematic documentation. If one line of SW code could implement more functionality than one SSI/MSI device then SW designers productivity is higher.

Next step was the micro-controller device (MCU) in 70's. MCUs are whole microprocessor systems on a single chip. The microprocessor of a micro controller system cannot have the leading edge complexity because a large share of the silicon area has to be reserved for memory on peripherals. Since late 70's, MCU devices have been the most popular general-purpose implementation format of digital systems.

Pure HW formats, after the introduction of GPL devices, have followed the development from Simple PLD devices, Complex PLDs and FPGAs to different ASIC formats [8]. The problem with schematic entry of PLD/FPGA and ASIC designs is the management of complexity. Schematic entry is applicable to the point where the number of pages is less than 50 to 100. If the symbols used in schematics are normal library components then maximum number of gates in design could be in the range from 5K to 10K. HW description languages like VHDL and Verilog, and logic synthesis tools were the solution for complexity management problems in the 90's [9,10]. The VHDL model of a digital system includes specification of both the HW structure and behavior of the system. Synthesis tool compiles a VHDL model to logic schematics. VHDL improves designer's productivity if synthesis of 50 to 100 pages of VHDL model produces more schematic

sheets. Complexity of HW modeled ASICs is then in the range from 10K to 1M gates and every page of the VHDL model specifies 250 to 10K gates of the implemented logic.

System-on-Chip is an ASIC or PLD/FPGA implementation format where one single silicon chip contain mostly reused IP based logic blocks like whole microprocessor systems. Manageable logic complexity of SOC's is somewhere between 1M gates to 100M gates. The new product specific logic of a SOC could be no more than 1M gates unless more productive design methodology is available.

Networks-on-Chip implementation formats are needed when the complexity available on a single chip requires a network of SOC's. Design documentation of a 1G gate NoC could be 100 pages of schematics where one page describes a structure of 10 million gates. Basic element, schematic symbol, could be 1 million gates IP block taken from the IP library of the design system.

2.2 System-on-chip design methodologies

In the late 1980's, it became possible to develop and implement computer organizations and architectures using ASIC design methods. It was realized that this introduces new possibilities for system optimization, because hardware/software partitioning is more flexible and not constrained by the costs as it was earlier. The first approaches to *software/hardware partitioning* were based either on software or hardware specifications, but very soon the need for functional specification and functional verification were identified, which resulted to *SW/HW codesign approaches* [11,12]. Functional specification was done using either a single language for the complete system [13] or by combining subsystem models that were each specified using application-specific languages [14]. Granularity at which the partitioning was done varied in different approaches from fine-grained instructions and medium-grained coprocessors to coarse-grained architecture level objects. The verification of the partitioned system required *cosimulation* tools, which could simulate both software and hardware [15].

The *system synthesis* approaches generate both the hardware implementation and software generation from the characteristics of functional specification. The design space exploration and evaluation of cost function are critical phases [16]. In the mapping based approaches the idea is to evaluate different mappings of functions to architectural elements. These *function/architecture codesign* approaches allowed using existing architectures or separately modeled architectures and the role of quality evaluation increased [17]. The intellectual property (IP) blocks and virtual components were proposed as a solution for design productivity problems in

late 1990's. The IP-block concept contains both the functionality and implementation. In *IP-based design*, the main problems are related to the evaluation of IP-block quality and integration issues. *Platform-based design* integrated and extended the earlier methods by reusing also system architectures and topologies in addition to components. Platform can also contain the software layer that helps in application development [18].

2.3 System design methodologies

The *waterfall model* divides the software development process into requirement analysis, architecture design, coding, testing and manufacturing phases that are executed sequentially in a top-down fashion [19]. In the *V-model*, the role of testing and verification is emphasized. Each phase of the waterfall model is accompanied by respective validation or verification phase [20]. The *spiral model* is based on the idea that the phases of waterfall model are repeated for different versions of the system [21]. These sequences of versions include conceptual models, specifications, prototypes, and product versions. This successive refinement approach is also called as *incremental development* approach, because the idea is to bring in more details during the development. The *meet-in-the-middle* principle combines the top-down and bottom-up approaches. The idea is to define the primitive components and objects and to develop them and to system structure simultaneously, so that in the middle of design process the final system is constructed from these primitives. ASIC synthesis tools with libraries and IP-based design are derivatives of this principle. In the *hierarchical design* methods, the system is partitioned to subsystems that are designed separately. *Concurrent engineering* attempts to take into account all the aspects of system design from the very beginning of the development process. It aims in developing subsystems simultaneously and increasing the communication between different design teams. Cross-functional teams and information sharing are the main enablers [22]. In *product family development*, the idea is to develop simultaneously a complete set of product variants [23]. These approaches are based on the ideas of mass customization, development of platform product that is easy to modify, or development of core product that can be extended by product features.

3. NOC METHODOLOGY REQUIREMENTS

NOC technology is a solution for complex systems, e.g. systems that can not be implemented otherwise. NOC-based products will have enormous amount of computational capacity and storage capacity in a very small

volume. NOCs will not be feasible unless the number of implementation is large. In consumer products, the mass production quantities can be achieved for example in personal communication terminals, such as mobile phones. Combination of diverse functionality and optimized product characteristics is the main challenge. In infrastructure products, such as information servers and networking products, the NOC technology can provide performance and capacity, but the success depends on how usable the NOC-based platform will be, when different types of products are considered. Third alternatives are generic computation platforms that are implemented using NOC platforms. The challenge is the programming model. The methodology presented in this chapter is mainly targeted to the consumer product segment.

3.1 Technology capacity effects

According to technology forecasts, the number of transistors is going to increase for the next ten years. When we use the 65 nm technology, it is possible to have over 2.4 billion transistors in a single chip. Current system architectures and design styles do not scale up to such dimensions and complexities.

If we consider the available silicon surface and the characteristics of gates and wires from the digital design view, it is obvious that we have to partition our chip area to several clock domains that communicate with each other asynchronously [24]. Similar partitioning is required for testability, energy distribution, and communication reasons. We need to reserve a part of the chip area to this kind of infrastructure and to partition the rest for application purposes, where each part has its own synchronization. This clearly speaks for joint optimization of physical architecture and design methodology.

When we take the computer science view, then the silicon surface is a high-capacity computation platform. Sequential processor can not exploit such a capacity effectively due to earlier mentioned reasons. Operation level parallelism leads to extremely complex design of dedicated hardware. Instruction-level parallelism (ILP) leads to extremely complex control of execution and besides, it is difficult to extract enough ILP from large applications. Thread-level parallelism could be viable alternative, but it requires significantly more threads than in current processors. Task and process level parallelism leads to parallel computer architectures, which have problems related to shared memories and communication, but these architectures could exploit the area more effectively. Besides, the programming model would be simple [25]. Application level parallelism means integration of embedded systems into a single chip and design

approach that resemble distributed system design. Such architecture could exploit all the possible area, but the feasibility requires that the ratio of local and global computations is high [2].

A NOC design methodology has to support both integration and effective use of different implementation technologies. Technology embedding means additional costs in ASIC manufacturing process. On the other hand, the use of different technologies should be justified by final product characteristics. This linkage between the costs and benefits of technologies and applications must be created in the design process.

3.2 Product requirements and economics

Already today, the complex SOC projects stretch the capabilities of companies and all the forecasts imply that product development capacity is lagging behind the technology capacity. Managing functional diversity and complexity are product requirement related issues that must be solved by NOC architecture and methodology.

For example, personal communication devices will have both local and mobile communication capabilities, various types of multimedia features, intelligent user interfaces, etc. Good implementation will require various types of technologies and computing architectures, which means that the implementation of total functionality may consist of a set of optimized subsystems that are integrated by common functional "glue".

Commercial products must meet both the market window and target cost constraints. Design productivity and design reusability are the critical parameters in this respect. The cost/object can be used as measure of effectiveness. If we analyze the cost/transistor in different types of designs it is easy to see that dedicated logic is more costly than IP-block based design or memory design, which benefits from regular and repeatable blocks. If we use cost/function, the software approaches are better than reconfigurable or hardware approaches. This means that in NOC systems we must favor flexibility over fixed solutions and regularity over ad hoc solutions.

The cost of design of NOC will be such that design reuse must be exploited at all levels and reusable units must be very large. In addition to IP-block reuse, we have to reuse the computation architectures, network architectures, software platforms, existing chips, and design methods. This means that we have to have very clearly defined interfaces for connecting resources to the network and for using the network resources with hardware or software. The reuse of the architecture requires means to modify and configure the architecture model so that application area specific platforms can be designed. The reuse of complete chip needs standardized application

mapping procedures that hide the architecture and hardware details from application designers.

3.3 Diversity and complexity

The product development consists of several areas of expertise and in the decision-making process, we have to be aware of effects to all aspects of design. When we start to develop our ideas into the products, the number of details is increased rapidly and the design methods and practices diverse. The design process is actually a sequence of design decisions that limit the design space. In the beginning, the number of possible implementations is largest and each design decision actually removes possible implementations until we have only one left. Working simultaneously in multiple technologies requires firstly that we partition our system functionality to technologies. Partitioning requires that we have portable models that can be transferred between technologies and subsystems. Secondly, we must derive quality criteria for each abstraction level from the quality criteria of the product. Thirdly, we have to make trade-off between different technologies and optimize the complete system instead of subsystems. This means that we need analysis and estimation methods that can be used at early phases of design and that take into account the coupling effects between technologies. Fourthly, we have to validate our decisions at various levels of abstraction. This needs an ability to combine different models. Finally, we have to be able to use technology dependent methods and tools. This requires encapsulation and interfacing capabilities.

In addition to management of work and implementation dependent issues, also business and economical issues significantly affect on how we should develop our products. Economical constraints such as time to profit require the reduction of design effort and design time, which can be achieved with platform-based approach with reuse both at implementation level and design level. The problem is how to guarantee adequate performance and variability. The market success depends on product's capability to satisfy users. User needs are diverse and the trend is to customize the products and services. If it is done with software, we sacrifice the performance features. If we do it in hardware, it is not cost effective. Naturally, the problem is more complex and we should consider the competitive strategy and market situation too.

In NOC design methodology, we must find out how to balance the system characteristics. Most important are complexity vs. feasibility, flexibility vs. suitability for purpose, and generic vs. dedicated solutions. These relate to the problem when and how we should apply reuse instead of design. Complexity/feasibility relates to the principles on how we design our

NOCs. It relates to both hardware architectures and software architectures. Flexibility issues relate to how we identify and encapsulate those parts of the system that must be optimized without losing the programmability and configurability of our system. The generic/dedicated trade-off is mainly a design effort issue. The problem is the separation of common or generic problems from those that need special methods. The generic problems should be solved so that they become as a part of NOC infrastructure.

4. NOC-BASED SYSTEMS

Networks can be applied to very different kinds of integrated systems and the success of a NOC depends on how effectively it can satisfy the often-contradictory requirements of applicability and performance. Applicability requires flexibility and generality, simple programming models, and simple development of final functionality. Performance requires efficiency, power awareness, dedicated structures, and optimal combinations of resources and functions.

The NOC design methodology presented in this chapter is meant for NOC systems that are based on product family type of product development presented in Figure 7-1. The NOC systems also use a specific architectural approach that is called "*integrated distributed embedded system*". The target application domain is telecommunication systems, but the ideas can be applied to other types of systems as well.




Figure 2-1. NOC-based product family development

We believe that NOC-based systems are economically feasible, if implemented chips can be used in several product variants, and if we can reuse the design in various application areas. However, we also believe that successful product must provide well enough quality for all users, which often requires solutions that are designed to specific needs. Achieving this requires balancing between flexibility and optimality in platforms.

If we compare the different implementation formats for ASIC systems, it is obvious that good performance characteristics need dedicated solutions. The configurable hardware or software based systems are lagging behind in area efficiency and power consumption, for example. If we compare the design effort and design cost the situation is opposite. We combine the benefits of both by providing hardware optimized, customizable computation *platforms* for each application area. The hardware optimization is done by embedding computation resources that are application domain specific into the chip. The design effort is reduced by providing a common basis for the integration, e.g. the *backbone*. The role of the backbone is to provide mechanisms for "distribution" and "integration". Infrastructure services, such as communication and interfaces are the best examples.

4.1 Architecture of NOC-based system

The NOC architecture is presented in Figure 7-2. The architecture is 2-D mesh, which consists of *resources* and *network*. The network elements are *switches*, *channels* and *resource-network interfaces*. The resources are embedded systems integrated into this architecture. The idea is to separate infrastructure area and application area also in the physical layout.

The only limitations to resources are related to the sizes of resources and to the interfaces. Any type of resource that can be implemented in synchronous area with the silicon technology is possible. From system perspective, the resources are independent embedded systems that have standard interfaces to the infrastructure services. The whole network of these embedded systems can be considered as a distributed system, where the network elements provide communication services.

The *regions* are used for embedding different technologies such as large memory systems, parallel computers, or reconfigurable arrays into our NOC. The size of the region can be larger than the size of resource, and the communication between the region and rest of the NOC can be arranged via specific switches only.




Figure 2-2. Example of a network on chip architecture. S = switch, rni = resource-network interface, P = processor core, C = cache, M = memory, D = DSP core, re = reconfigurable logic, L = dedicated hardware.

4.2 Quality characteristics

NOC based systems have similar quality characteristics as any other system with respect to performance, cost and variability. However, in addition to generic quality characteristics, we can distinguish the quality characteristics for the NOC platform, which are different from product characteristics. Identification of these is important, because they affect our design decisions and the quality assurance procedures.

When considering the NOC platform quality characteristics, we have to consider communication systems, flexibility, and system integration issues. Communication system has to provide performance guarantees for delay, bandwidth, power consumption, and communication reliability. The basic communication functionality has to be accompanied by performance figures to allow application engineers to build reliable applications with predictable performance.

Flexibility is a major concern, since the main point of a platform is to support a variety of concrete products and applications. Hence, the platform has to support different traffic types with widely varying requirements. Safety critical systems will put higher demands on predictability than multi-

media applications. Power management and efficiency is of highest concern in hand-held devices. The matter is complicated further by the prospect that various application types will co-exist in future NOC based products. Thus, a NOC platform has to be careful at offering a sufficient range of choices to application engineers.

In order to support system integration and application mapping, a methodology has to provide sophisticated means to for system analysis. Since the final product will contain a large number of different resources, it will be a formidable task to integrate the performance assessments of the resources with the performance guarantees provided by the network into a reliable system performance assessment.

5. BACKBONE-PLATFORM-SYSTEM METHODOLOGY

The organization of work in our backbone-platform-system NOC design methodology is based on heavy reuse of existing designs, reuse of software systems and reuse of platforms. Our NOC concept encapsulates embedded systems into a network, computers, reconfigurable fabrics, or embedded memories into resources, and intellectual property blocks, virtual components and software into embedded computer systems. The hierarchical encapsulation introduces well-defined interfaces between different layers of NOC system and allows constructing the complete system from black-box type of units.

We have partitioned our design methodology in three ways. Firstly, we have divided the NOC system development in three main layers that are backbone, platform and system development. Secondly, we have divided the methodology according to what kind of methods we need for the development of subsystems. Thirdly, we have divided each design activity into phases according to the purposes of those phases.

5.1 NOC layers

The layered NOC development separates technology specific e.g. backbone issues, application area specific e.g. product platform issues and product instance specific e.g. system issues from each other. The layering enables the separation of infrastructure design from resource design and application design.

The *NOC backbone layer* consists of infrastructure and especially communication services. Infrastructure part consists of physical components and wiring for clocking, energy distribution, and testing, for example. The

communication services consist of physical components such as channels, switches, and network interfaces, and basic services such as protocols for physical and data link layers. The purpose of the backbone is to provide an *integration framework* for platform and resource designers and to hide physical level details from them. It also provides a *communication model* for application developer in a form of protocols. The physical design issues have an important role, because the system-level communication challenges the limits of technology.

The *NOC platform layer* describes the computation platform for target application area. Its main purpose is to hide the hardware from application designer and to provide a *programming model* for them. Platform design requires the understanding of how the target systems operate and what kind of computation they have. Scaling of the network, definition of regions, design of the resource nodes, and definition of the system control are the activities that must be based on abstract models of applications. The platform encapsulates the hardware design problems and serves as a manufacturing integration platform for system developers.

The *NOC system layer* describes programmed chip in the final product. The resource allocation, optimization of network usage and verification of performance and correctness are the main problems that are basically similar to what distributed and parallel system designers have to face.

5.2 NOC design methods

The integrated distributed embedded system concept does not define the embedded resources. The platform designer can use all his creativity to improve the applicability of the platform in the application area domain.

Freedom of choices challenges the design methodology, because consistent methodology must provide tools and techniques that allow designing and modifying everything that are not fixed. We have organized our methodology as a hierarchy of subsystem specific methodologies. The lower level methodologies can be used in implementing the subsystems of upper levels, as long as upper level constraints are fulfilled. The methods at each level are following:

1. System-level design and networked system's design methods. System-level design deals with application modeling and analysis, while network design methods are applied for mapping and load balancing.
2. Parallel computer design, distributed computer system methods, embedded system design, etc. The choice of the method depends on the type of region that we are developing.
3. Computer system design, software system design, codesign, or configuration design depending on the resource, are examples.

4. Basic software, logic and synthesis methods and design flows.

The feasibility of hierarchical methodology depends on how easy it is to model and present the constraints for lower levels and how much these constraints actually effect on the final system. Application programming interfaces and virtual component interfaces are essential elements in successful implementation.

5.3 NOC design phases

The complete design flow consists of a sequence of design activities. The exact set and order of activities is always product specific. In our methodology, a generic design activity is divided into analysis, estimation, decision, and validation phases. *Analysis* phase is needed for understanding what we have as input for our design decision. For example, when software system is designed the first task is requirement analysis, where the objective is to find out what are the user needs that should be implemented but implementation considerations are strictly forbidden. *Estimation* is part of design space exploration where the objective is to forecast the outcome of possible decision without putting too much effort into the implementation of the decision. *Decision* is the phase when selections are implemented as a new system model. Decision includes the creation of a new model, which typically includes the refinements and transformations. In the *validation* phase, the effects of decisions are measured using the new system model as input. The validation should consider all the quality characteristics of the system.

5.4 NOC design flow

The overview of the design flow for NOC-based systems is presented in Figure 7-3. In this figure, the design flow is presented as path through the design plane, where the x-axis direction represent the development of system resources and the y-axis the development of system functionality. In the bottom left corner of the plane, the system does not exist, and in the top right corner, the final NOC-based system is ready. In NOC system development, we divide the problem according to NOC layers, e.g. the backbone development, platform development, and system development.

The *backbone design* aims at providing infrastructure services for different types of NOC platforms. In backbone design, we have to combine the technology dependent implementation and very generic architecture considerations. The analysis phase must consider the silicon characteristics and general principles of system construction. During the estimation phase, the focus must be on production yield, power management and

2. A Design Methodology for NoC-Based Systems

infrastructure/application area efficiency. In the decision phase, network topology, switches, channels, resource-network interfaces and protocols must be designed as well as the principles of resource integration e.g. how the resource are connected to the infrastructure. The quality of these services should be measured during validation phase as simplicity of resource integration, cost and performance.




Figure 2-3. NOC design flow

The *platform design* aims at architecture and set of system services that is optimized for a certain product area. The three major decisions concerning the platform architecture are the dimensions of the network, types and sizes of technology regions, and the contents of resources. Understanding the characteristics of typical application workload is the main requirement for analysis phase. The definition of network size and shape requires the analysis of the total complexity of computation, communication and storage. The characteristics of applications affect to technology regions. For example, do we need memory capacity that does not fit into synchronous resource slot, or large areas of reconfigurable arrays or processors, or parallel computer level capacity? Finally, we have to analyze the characteristics of individual functions in order to see what kind of computing resources are needed inside the resource areas. Because of complexity, we have to be able to do these analyses very rapidly and at very abstract level, and current methods and tools are not capable to deal with this problem.

The platform architecture decisions must be supported by feasibility estimations that take into account the application characteristics, resource characteristics and function/architecture mappings. These estimations are

needed both in network level and resource level, and they have to support the evaluation of relative costs of different implementation formats. Performance and efficiency are the behavioral criteria to be considered.

In the decision phase, we must design all the regions and resources. Applying existing design methods and tools is essential. Integration of the models of different blocks both for implementation and estimation purposes requires a *NOC description language* for the modeling of the complete structure and how resources connect to infrastructure services.

The validation at platform level must cover the quality criteria above. Since the platform architecture is a programmable and configurable system that can be used for various types of products, the performance validation must be done using the methods of distributed and computer system design. Analytical performance estimation, workload model based performance simulations, and benchmarks are therefore feasible alternatives.

The final programming and configuration of a NOC system is called the *application mapping*. Application mapping consists of four activities. First, we have to model the whole system. Secondly, we have to partition the functionality into network resources. Thirdly, we have to implement the behavior with the resources. Finally, we have to download the complete system description into our NOC platform and verify it. In various phases of application mapping, we need to construct a complete description of a system although the description itself may consist of several different descriptions. A *NOC assembly language* is therefore needed for presenting the application's mapping to resources and interfacing to platform's system services.

During the application mapping the modeling of the behavior of the system and its workload for the platform elements are main tasks. Modeling of complete functionality so that we can reuse existing subsystem implementations requires high-abstraction level methods and means to add interfacing functions that are platform and resource dependent. Workload models are required during mapping phase for optimization purposes.

Mapping functions to resources is the decision that must be supported by estimates on system quality. Feasibility of platform for chosen set of functions, performances of functions with chosen mappings and utilization of platform resources are figures of merits that must be considered before committing to implementation design.

Validation and verification of complete NOC-based system are problems in which we do not have answers. Built-in self-test structures are most likely alternatives for verification of the platform, but also hierarchical testing approaches with separation of communication, data storage and computation must be developed.

The presented NOC design flow requires new languages for representing the NOC specific structures and services and new methods for supporting NOC related decisions. Proactive decision support methods are needed for efficient reduction of design space. Complexity and capacity estimations are needed for sizing and scaling so that we know the quantities of objects. Mappability estimates are needed for selections of physical and functional objects, so that we know the suitability of objects. Performance and workload estimation is needed for allocation and validation, so that we know the effectiveness of our design. Cost and quality estimation is needed for feasibility analysis. We want to know as early as possible if the work is worth the effort.

Validation methods are needed after decision for analysis of consequences. Validation methods can be more detailed because only a small number of alternatives are designed in more detail. However, the validation methods must also be very effective. Validation must cover all the quality criteria, but if we consider the performance, for example, then the quality validation procedure may consist of network simulation, transaction-level architecture simulations, and instruction-level processor simulations.

6. SUMMARY

NOC-based products will be extremely complex and difficult to design. Management of structural complexity and functional diversity requires that the physical issues, architecture issues and methodology issues are considered and optimized together. Most likely economically feasible products can be developed only if both designs and implementations are reused, and if product functionality and system resources can be separated during the design.

Principles of a backbone-platform-system design methodology have been presented. The BPS methodology is based on the manufacturing platform-based product family development. The application-area specific platforms are based on the NOC backbone that is a generic NOC architecture allowing integration of dedicated computation and storage resources and a set of supporting infrastructure services.

The BPS methodology supports the encapsulation of the design of network elements, e.g. resources and their functionality, and the capability to support design decisions. The most important decisions concern technology and computing architecture selections, system dimensions and mappings at different abstraction levels.

REFERENCES

- [1] Allan, A. et al, 2001 Technology Roadmap for Semiconductors, Computer, Vol. 35, No. 1, 2002, pp. 42-53
- [2] Kumar, S., et al, A Network on Chip Architecture and Design Methodology, Proc. of 2002 IEEE Computer Society Annual Symposium on VLSI, 2002, pp. 117-124
- [3] Chang, H. et al, Surviving the SOC Revolution – A Guide to platform based design, Kluwer, 1999, 235 pp.
- [4] Compton, K. and Hauck, S., Reconfigurable Computing: A Survey of Systems and Software, ACM Computing Surveys, Vol. 34, No. 2, 2002, pp. 171-210
- [5] Porter, M., Competitive Strategy, Free Press, 1980, 365 pp.
- [6] Day, R.G. Quality Function Deployment – linking a company with its customers, ASQC Quality Press, 1993, 245 pp.
- [7] Wolf, W. Computers and Components - Principles of Embedded Computing System Design, Morgan-Kaufman Publishers, 2001, 662 pp.
- [8] Smith, M. J. S., Application-specific integrated circuits, Addison-Wesley, 1997, 1040 pp.
- [9] Gajski, D. (ed.), Silicon Compilation, Addison-Wesley, 1988, 450 pp.
- [10] MacMillen, D. et al, An Industrial View of Electronic Design Automation, IEEE Trans. on Computer Aided Design of Integrated Circuits and Systems, Vol. 19, No. 12, 2000, pp. 1428-1448
- [11] Ernst, R. et al, Hardware-Software cosynthesis for microcontrollers, IEEE Design and Test of Computers, Vol. 10, No. 4, 1993, pp. 64-75
- [12] De Micheli, G., Computer-aided hardware-software codesign, IEEE Micro, Vol. 14, No. 4, 1994, pp. 10-16
- [13] Kumar, S. et al, The codesign of embedded systems: a unified hardware/software representation, Kluwer, 1996, 274 pp.
- [14] Edwards, S. et al, Design of Embedded Systems: Formal Models, Validation and Synthesis, Proceeding of IEEE, Vol. 85, No. 3, 1997, pp. 366-390
- [15] Staunstrup, J. and Wolf, W., Hardware/software co-design: principles and practice, Kluwer, 1997, 395 pp.
- [16] Eles, P. et al, System Synthesis with VHDL, Kluwer, 1998, 370 pp.
- [17] Balarin, F., et al, Hardware-Software Co-Design of Embedded Systems – The POLIS Approach, Kluwer, 1997, 297 pp.
- [18] Keutzer K. et al, System Level Design: Orthogonalization of Concerns and Platform-Based Design, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol. 19, No. 12, 2000, pp. 1523-1543
- [19] Royce, W., Managing the Development of Large Software Systems, IEEE WESCON, 1970, pp. 1-9
- [20] McDermid, J.A. (ed.), Software engineer's reference book, Butterworth, 1991, 1500 pp.
- [21] Boehm, B.W., A Spiral Model of Software Development and Enhancement, IEEE Computer, Vol. 21, No. 5, 1988, pp. 61-72
- [22] Linton, L. et al, First principles of concurrent engineering: A competitive strategy for product development, Report of the CALS/Concurrent Engineering Electronic System Task Group, 1992
- [23] Wheelwright S.C. and Clark, K. B., Revolutionizing Product Development, Free Press, 1992, 364 pp.
- [24] Hemani, A. et al, Network on chip: an architecture for billion transistor era, Proc. of IEEE NorChip Conference, 2000, 8 p.
- [25] Forsell, M., A Scalable High-Performance Computing Solution for Network-on-Chip, IEEE Micro, Vol. 22, No. 5, 2002, pp. 46-55

Chapter 3

MAPPING CONCURRENT APPLICATIONS ONTO ARCHITECTURAL PLATFORMS

Andrew Mihal

Kurt Keutzer

University of California, Berkeley

mihal@eecs.berkeley.edu

keutzer@eecs.berkeley.edu

Abstract

Embedded system designers are faced with an expanding array of challenges in both application and architecture design. One challenge is the task of modelling heterogeneous concurrent applications. Another is the task of finding a programming model for heterogeneous multiprocessor architectural platforms. Compounding each of these challenges is the task of implementing heterogeneous applications on heterogeneous architectures. We believe that the key problem underlying each of these challenges is the modelling of concurrency, and the key to modelling concurrency is to capture concurrent communication formally in models of computation. This chapter broadly outlines a disciplined approach to the design and implementation of communication structures in embedded applications. Our approach combines the Network-on-Chip paradigm with the models of computation paradigm. We use models of computation to capture the communication requirements of an application as well as to abstract the capabilities of a communication architecture. Then, application requirements and architectural capabilities are matched using a discipline based on Network-on-Chip principles. In this chapter we describe this approach and present a case study where a Click network routing application is implemented on a multiprocessor architecture using this discipline.

Keywords:

model of computation, network on chip, application development environment, programming model, programmable platform, heterogeneous concurrency, multiprocessor architecture, design space exploration, mapping

1. Introduction

Embedded system designers are faced with an expanding array of challenges in both application and architecture design. One challenge is the task of mod-

39

A. Jantsch and H. Tenhunen (eds.), *Networks on Chip*, 39–59.
© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

elling heterogeneous concurrent applications. The burden of this task is typically born by an *application development environment*. An application development environment is a framework that is tailored to modelling (and principally simulating) system applications. Successful application development environments have considerable specialized support for application-domain-specific modelling including domain-specific libraries and domain-specific models of concurrency.

Examples include environments for Internet packet processing[1], signal processing[2, 3], and reactive systems[4, 5] among many others. Although there are many issues that must be addressed in application development environments, the critical issue is allowing the system developer to naturally organize, partition, and model the concurrency of the application.

Another set of challenges arises with each successive generation of architectural platforms. Moore's law enables an exponentially growing number of architectural features in these platforms. In this chapter we focus on programmable architectural platforms or simply *programmable platforms*[6]. Features in the current generation of programmable platforms include: processing parallelism, special-purpose functional units, intelligent memory structures, networks-on-chip, and peripherals[7]. Processing parallelism comes in three varieties. Process-level parallelism is associated with independent programs and memory spaces. Instruction-level parallelism involves the parallel execution of operations from a single thread of program control. Gate-level parallelism encompasses the varieties of parallelism offered by small special-purpose execution units and reconfigurable structures. The principle challenge of programmable platforms is finding a way to harvest their capabilities in a programming environment.

Many programming environments for programmable platforms consist simply of an assembler. Some go as far as to offer a C compiler and a debugger. Sequential languages like C and assembler do not help the software developer deal with process-level concurrency. Beyond the issue of merely offering primitive software development tools there is the deeper issue of allowing developers to effectively access and exploit the important architectural features of the programmable platform. Developers need a *programming model* of the programmable platform. A programming model is a complete abstraction of the key features of a programmable platform. Its goal is to enable developers to exploit the capabilities of the device while hiding less essential details.

A specific example is the challenge of developing a programming model for the Intel IXP1200 processor[8]. A successful programming model for this device must allow programmers to efficiently exploit the parallelism of six microengines with four thread contexts each, a special-purpose hash engine, and a complex communication architecture with Ethernet peripherals. Although

there are many issues involved in the development of a programming model, the central issue is exposing and managing parallelism.

1.1 The Implementation Gap

Suppose that we have been fortunate enough to find an application development environment that naturally fits our system application. Further suppose that we have found a programming model that adequately abstracts and exports the key features of our target programmable platform. The remaining challenge is to *correctly* map from our application development environment to the programming model of our programmable platform in a way that *efficiently* implements our application. We denote this challenge the *implementation gap*.

Depending on the application domain, there may be a number of problems in crossing the implementation gap. We believe that the fundamental issue is mapping the concurrency captured in the application development environment to the parallelism in the programming model of the programmable platform. This transition needs to be made correctly and efficiently.

Concurrent communication is one aspect of this issue. For example, a signal processing application may be designed using an application development environment like YAPI[3] that is natural for signal processing applications. YAPI contains an abstraction for concurrency based on Kahn Process Networks[9]. This abstraction defines a system of queues that concurrent processes use for communication. The target architectural platform may provide only a single shared memory for communication between all processing elements. The way that the application wants to perform concurrent communication does not match the structures that the architecture provides for performing concurrent communication. This implementation gap problem makes the task of mapping from application development environment to programming model difficult.

This example shows that the implementation gap exists even when designers start with a good application development environment that provides formal abstractions for concurrency. The implementation gap problem can be much worse. If designers choose poor abstractions for concurrency, or design heterogeneous applications that use multiple different abstractions for concurrent communication, the implementation gap grows in complexity.

1.2 A Disciplined Design Methodology

We maintain that modelling concurrency is the core problem behind the design of a successful application development environment, the exportation of a successful programming model, and the correct and efficient mapping from the application development environment to the programming model. The key to the embedded systems design challenge is modelling concurrency, and the key to modelling concurrency is to capture concurrent communication formally




Figure 3.1. Y-chart Design Methodology

in *models of computation*. A model of computation is a mathematical formalism for describing the interaction between components in a concurrent system. Space does not permit a detailed discussion of models of computation, but to understand the principles that are referenced in this chapter the reader should see [10]. Our design methodology uses models of computation to capture the concurrency contained in an application as well as the parallelism provided by a programmable platform. This allows us to compare application requirements and architectural capabilities in a common theoretical framework. Then, we combine models of computation with Network-on-Chip principles to form a discipline for implementation where application requirements are matched with architectural capabilities. The fundamental Network-on-Chip principle we use is the concept of viewing a programmable platform as a composition of computational modules that communicate over an interconnect topology using network protocols.

Figure 3.1 graphically describes the steps in our methodology. The arrangement of steps forms a Y-chart[11]. The Y-chart is an iterative strategy for design space exploration that enables the co-evolution of hardware and software.

In this approach, applications are designed with a high-level application development environment that uses models of computation. This provides software programmers with the right abstractions for modelling heterogeneous,

concurrent applications. Section 2 describes how application environments use models of computation to define the concurrency in an application. In particular, we present the *Communication Implementation View*, an abstraction of the application and the model of computation that formally captures the communication between concurrent components.

At the same time, models of computation are used to export the functionality of the target programmable platform to the developer. The programming model for a specific architecture is that of a machine which performs concurrent communication according to the rules of a particular model of computation. This style of programming model is an appropriate match for application development environments that use models of computation themselves. The concurrent communication implemented by the programmable platform is also captured in the Communication Implementation View. This technique is described in Section 3.

Next, the application development environment and programming model abstractions are combined in a mapping process. In this step, designers match the concurrent communication requirements of an application with the concurrent communication capabilities of an architectural platform using the Communication Implementation View. Section 4 describes a discipline for this task based on the Network-on-Chip principle of viewing the architecture as a communication network. The result of the mapping process is a formal model that describes the implementation of an application on an architectural platform. A code generation process is then used to produce a concrete implementation of the system.

The use of our design methodology and mapping discipline is demonstrated with a network routing application in Section 5. We summarize and draw conclusions in Section 6.

2. Application Development Environments

The goal of an application development environment is to provide programmers with abstractions that are useful for a particular application domain. This is accomplished by separating the concerns of computation and concurrency. Designers focus on implementing the functionality of the application using a domain-specific library of computational components. A model of computation provides a model of concurrency that is natural for the application. It does the work of defining the interaction between computational components. Without this separation of concerns, designers would have to use *ad hoc* techniques to implement the application's computation and concurrency together. Then the definition of the application's concurrency would be implicit and informal. Models of computation define the application's concurrency explicitly and formally.

To implement an application that uses models of computation on an architectural platform, we must implement not only the computation defined by the application's computational components but also the concurrency imposed by the application's model of computation. The Ptolemy II system provides a general strategy for implementing models of computation[12]. Although Ptolemy primarily targets simulation on desktop workstations, the separation of concerns and modular design principles that it uses have broad utility. These underlying theories enable several features that are important for application development environments, especially the ability to design heterogeneous application models that use several different models of computation in combination. In this work we consider applications designed in application development environments that use Ptolemy principles to separate the concerns of computation and concurrency. We extend Ptolemy's general strategy for implementing models of computation to target execution on programmable platforms. Next, we describe the division of labor in the Ptolemy system.

Ptolemy works by separating the functionality of a model of computation into two parts. First, the *communication* portion of the model of computation describes how concurrent application components share data (see Section 2.1). Second, the *control* portion of the model of computation describes how application components execute relative to each other (see Section 2.2). These aspects of a model of computation are implemented with modular components that perform separate communication and control tasks. The *model of computation components* work together with domain-specific library components to make complete, functional application models.

2.1 Communication

Ptolemy library components, called *actors*, communicate through consistent interfaces called *ports*. Ports are connected to other ports with *relations*. The semantics of communication of data *tokens* over a relation is implemented by a model of computation component called a *receiver*. A receiver implements the communication aspect of a model of computation.

Each model of computation has its own receiver, but all receivers implement the same modular interface. The primary functions in the receiver interface are *put*, *get*, *hasToken*, and *hasRoom*. The communication semantics of any model of computation can be implemented through these functions.

2.2 Control

The *director* model of computation component implements the control aspect of a model of computation. Directors tell actors when to execute by calling functions through the *executable* modular interface. The primary functions in this interface are *initialize*, *prefire*, *fire*, *postfire*, and *wrapup*. Directors can

implement any combination of sequential or parallel firing semantics that a model of computation may require.

The director is a centralized control component. Since Ptolemy primarily targets simulation, it is not inefficient to implement even parallel execution semantics with a centralized controller. However, programmable platforms must make a firm distinction between distributed and centralized control. Processes that can execute in parallel should not have to rely on a centralized controller unless there is an explicit execution dependency. We therefore extend Ptolemy's system to include a model of computation component called a *subdirector*. One subdirector is instantiated for each actor in an application model. The director and subdirectors cooperate to implement the control semantics of a model of computation. The director implements the control functionality that is centralized in nature, and the subdirectors implement the control functionality that is distributed in nature.

2.3 The Communication Implementation View

In an application development environment like Ptolemy, the details of how models of computation are implemented are purposefully hidden. Developers do not use directors, subdirectors and receivers directly. Rather, a model of computation is simply specified as a property of the application model. This is exactly the abstraction of concurrency that makes application development environments useful for modelling concurrent applications. However, the details of how models of computation are implemented become important once designers are ready to target a programmable architectural platform.

Our design methodology provides a useful abstraction for capturing these details called the Communication Implementation View. The Communication Implementation View is part of the *multiple views methodology*[13]. In the multiple views methodology, designers use specialized abstractions for different aspects of the embedded system design process. In this case, the abstraction is used for the task of mapping concurrent applications onto programmable platforms.

The Communication Implementation View provides an abstraction of a concurrent application that graphically depicts the details of how the application's concurrency is implemented. Here we make use of the separation of concerns described above that powers Ptolemy's general strategy for implementing models of computation. The Communication Implementation View shows designers an abstraction of an application model that contains the application's computational components together with the model of computation components defined by Ptolemy. This view captures the way in which the application is divided into computation, communication and control aspects.




Figure 3.2. Simple Ptolemy Application




Figure 3.3. Communication Implementation View

For a concrete example, consider the basic Ptolemy application shown in Figure 3.2. This application model consists of a producer actor and a consumer actor that communicate over a single relation. The model has an annotation that indicates the model of computation used. The Communication Implementation View generalizes to all models of computation, so in this example we refer to the application's model of computation with the variable *X*.

Figure 3.3 shows the Communication Implementation View for this application. The components in the Communication Implementation View are derived from the original application actors and their underlying model of computation components. There is one director that implements the central control for the model, and one receiver that implements the semantics of communication for the relation between the application actors. Each application actor has a subdirector that implements the distributed control of the model.

The Communication Implementation View is useful in particular because it formally captures the concurrent communication of the application. It divides the application into components that *implement* concurrent communication (receivers) and components that *utilize* concurrent communication (actors). We believe that capturing concurrent communication in this manner is the key to




Figure 3.4. General Architecture

mapping concurrent applications onto programmable platforms. In the next sections we elaborate on how the Communication Implementation View is used in our design methodology.

3. Programming Models

Formally capturing an application's concurrency is only part of the design challenge. There is still the issue of formally capturing the capabilities of an architectural platform in a programming model. In particular, we are interested in capturing the concurrent communication capabilities of the target platform's communication architecture. In this section, we describe how models of computation can be used to export this functionality to the programmer.

What characteristics of a communication architecture are important to capture in a programming model? To answer this question, consider the architecture shown in Figure 3.4. In this figure, two processing elements are connected with a bus. This example is a generalization of a more complicated communication architecture. It illustrates the level of granularity that is used to identify the important features of a communication architecture. Each processing element has a hardware bus interface. The bus and bus interfaces together implement a communication link that transfers data between the processing elements. The programming model for the device must describe how software running on the processing elements can utilize this communication link. Examples of the important characteristics of the communication link include:

- **Memory:** The communication link may exhibit the behavior of a shared storage location. Transmitting and receiving data is abstracted as reading or writing to the shared storage. An example is a bus that uses a FIFO for the intermediate storage of data. An important characteristic of this type of communication link is the amount of space available for data.

- **Control:** How does the communication link influence the flow of control of the software running on the processing elements? Transmit and receive operations may exhibit blocking or non-blocking behavior.
- **Data:** A communication link may be designed to handle continuous streams of data, or it may be appropriate only for shorter, intermittent messages. Are data items guaranteed to arrive at the receiver in the same order they were transmitted? Does the communication link guarantee that no data items will be lost, or that the data will arrive without bit errors?
- **Time:** Are there any guarantees on the time it takes to transfer data? Programmers must also know what happens when there is a mismatch in the transmit and receive rates. If a processing element does not read incoming data in a timely fashion, it may be overwritten with new data or the new data may be lost.

The characteristics of a communication link that are used to define a programming model are identical to the characteristics used to describe a receiver, the component that implements the semantics of concurrent communication in the application development environment. Therefore we can use the concept of a receiver as an abstraction for the communication link. The functionality of the communication link is exported to the programmer as a component in the Communication Implementation View called a *virtual receiver*. This is shown in Figure 3.5. This abstraction applies to a specific communication link in the communication architecture. In general, a complex communication architecture will export a set of these virtual receiver components. Each can implement a different communication semantics.

The virtual receiver abstraction tells the programmer that the architecture implements the communication semantics of a particular model of computation. This is a good abstraction for a programming model because it enables a direct comparison between the communication semantics required by an application and those implemented by the architecture. Application requirements appear in the Communication Implementation View in the form of components that *want to use* a particular type of receiver. Architectural capabilities appear in the Communication Implementation View in the form of components that *implement* a particular type of receiver. The next step in the design methodology is to match the application requirements with the architectural capabilities.

4. Bridging the Implementation Gap

In the mapping stage of a design methodology, designers make a transition between the application development environment and the programming model. The challenge in this step is to *correctly* map the concurrent communication




Figure 3.5. Exporting Virtual Receivers

requirements of the application to the concurrent communication capabilities of the programmable platform in a way that *efficiently* implements the application. In our methodology, designers tackle this challenge with the Communication Implementation View abstraction and a discipline based on Network-on-Chip principles.

In the previous section we placed application components and components exported from the architecture together in the Communication Implementation View. The goal of the mapping process is to replace receivers from the application with virtual receivers from the architectural platform. Then, the Communication Implementation View will contain a model that describes how the computational components of the application interact with each other using the communication architecture of the system. This model is called the *communication implementation model*.

Due to the implementation gap, the communication semantics implemented by the communication architecture may or may not match the communication semantics specified by the application's model of computation. Therefore the architectural virtual receivers may or may not be acceptable substitutes for the original application receivers. It is a special case when the application and the architectural platform match naturally. Then, the mapping problem is simplified considerably and high performance implementations can be achieved. Examples of this "natural fit" principle are given in [11]. Designers can make connections in the Communication Implementation View such that the com-




Figure 3.6. Natural Fit Mapping

putational components of the application are directly connected to the virtual receivers exported from the architectural platform. This is shown in Figure 3.6.

This communication implementation model represents a complete solution to the mapping problem. It is a formal model of the interaction between the application and the architectural platform. A designer can then run a code generation process on this model to obtain a concrete implementation of the application that executes on the target programmable platform. This is further described in Section 5.4.

If the application and the architectural platform are not a perfect match for each other, additional work is required to solve the mapping problem. Consider a communication architecture that implements non-blocking transmit and receive operations. The communication semantics implemented by this communication architecture are not an exact match for an application designed using the Communicating Sequential Processes model of computation, which calls for a rendezvous style of communication[14]. Alternatively, consider a communication architecture that does not guarantee that data items will be transmitted without error. These semantics are probably not a good match for any high-level application model of computation. The functionality of the architectural platform and the application can differ in all of the ways that characterize communication semantics - memory, control, data and time.

We can leverage Network-on-Chip principles to close these semantic gaps. The Network-on-Chip paradigm introduces the concept of viewing the architectural platform as a network of communicating computational modules. In our discipline, we extend this idea to formulate the implementation gap as a communication network problem. By definition, a networking protocol provides a particular communication abstraction on top of a set of lower level abstractions. A protocol can be designed to implement rendezvous-style communication on top of a non-blocking infrastructure, or to guarantee error-free transmission on top of a lossy bus. Designers solve the implementation gap problem by creating protocol stacks to match the communication semantics implemented by the architectural platform with those required by the application.

Figure 3.7 shows a model of a typical protocol. In this model, two instances of the protocol communicate with each other using a communication link that




Figure 3.7. Protocol Model




Figure 3.8. Exporting Protocols as Virtual Receivers

implements a particular communication semantics. The combination of the two protocol instances and their communication link implement a new communication link that provides a higher level communication semantics. The functionality of a protocol-based communication link can be exported into the Communication Implementation View in much the same way as an architectural communication link. The difference between exporting a protocol-based link and an architectural link stems from the fact that the protocol-based link not only *implements* the receiver interface, it also *requires* a receiver for the lower-level connection. Therefore the virtual receiver for the protocol-based link is split into two halves, as shown in Figure 3.8. The gap in the middle indicates the protocol's requirement for a lower-level connection. It is meant to be filled in with a lower-level virtual receiver.

An architectural virtual receiver that was not directly appropriate for the application components may be a suitable receiver for the communication link between the protocol components. Figure 3.9 shows this relationship graphi-

cally. The bottom layer in this figure represents the architectural platform. The communication link shown in that layer provides communication functionality for the protocol components in the middle layer. The protocol builds on top of this functionality to provide a higher level of communication semantics for the application actors in the top layer. Figure 3.10 shows how these communication abstractions are described in the Communication Implementation View.

The use of layers of abstraction to bridge the implementation gap is the core of our discipline for solving the mapping problem. Designers can chain any number of protocols together to implement the desired communication abstractions, using both top-down and bottom-up techniques. Protocol stacks reminiscent of the OSI model follow directly[15]. Behavioral type checking is used to determine that all connections between components that require the receiver interface (application actors and protocols) and components that implement the receiver interface (protocols and architectural virtual receivers) are valid[16]. This is used to prove that the protocol stacks designers create in the Communication Implementation View correctly implement the semantics of concurrent communication required by the application's model of computation.

In summary, this discipline is important and useful because it gives designers the necessary framework for understanding the implementation gap. Instead of using *ad hoc* techniques to match the concurrent communication required by an application with that provided by an architectural platform, designers can think about the problem as a communication network problem and use networking principles to reach a solution. Network-on-Chip protocols implement a transition between the application development environment and the programming model. The Communication Implementation View provides a useful high-level abstraction for making this transition. The result is a formal model that describes how the computational components of the application interact with each other using the communication architecture of the system. This technique quickly reveals the roots of the mapping problem and provides the right abstractions for a disciplined solution.

5. Design Example

In this section we will apply this design methodology to a network processing application. This application domain is interesting because there exists a remarkable implementation gap between network processing application development environments and network processor programming models.

It is important to clearly distinguish between the network processing application area and the Network-on-Chip paradigm for system design. The network processing application area is concerned with traditional computer network problems such as IP routing, network address translation, quality of service, and others. A network processor is an architecture designed specifi-




Figure 3.9. The Layers of Abstractions Behind a Protocol Stack




Figure 3.10. A Protocol Stack in the Communication Implementation View

cally for network processing applications[7]. In this example, we consider a Network-on-Chip architectural platform that implements a network processing application. Network-on-Chip and network processing share many of the same concepts and concerns. The similarity between the application that the system performs (the handling of Internet packets) and the internal workings of the system (communication via Network-on-Chip packets) is merely coincidental.

An example of a network processing application development environment is the Click Modular Router[1]. The goal of Click is to provide designers with the right abstractions for describing the inherent characteristics of packet processing applications. These characteristics include manipulating bits in packets, buffering packets, and controlling the flow of packets (both spatially and temporally). Click combines an extendable library of components with a model of computation and a tool set for producing executable implementations. Users assemble dataflow-like graphs of library components to describe a packet processing task. The Click model of computation defines the concurrency in these graphs. It abstracts away the details of how parallel streams of packets flow through the graph.

High-level Click models can be executed with reasonable performance on desktop machines, but there is no current methodology to target any of the numerous network processor architectures that have been proposed[7]. These heterogeneous multiprocessor architectures are supposed to excel at packet processing. Unfortunately, they are extremely difficult to program. Programming models for current network processors do not give developers the abstractions they need to implement concurrent applications like Click routers efficiently and correctly. The challenges here are twofold. The first challenge is to export a programming model from a network processor architecture that captures the parallelism implemented by the machine. The second challenge is to map the concurrency in a Click application to the concurrency captured in the programming model. Our disciplined design methodology can tackle these challenges.

5.1 Application Model

Figure 3.11 shows a graphical model of a Click router application based on the design in [1]. To simplify this example, we have eliminated the error handling functionality of the router and have shown only two channels. This application model was designed in Teepee, one of the core tools of the Mescal project[13]. Teepee is based on Ptolemy II and provides an implementation of the Click model of computation following the principles described in Section 2. This Click model can be simulated within Teepee to test for functional correctness using Ptolemy's standard simulation capabilities.




Figure 3.11. Simplified Click Router Application

5.2 Target Architectural Platform

The target architectural platform consists of six processing elements that are connected with two shared buses as shown in Figure 3.12. This architecture is similar to the microengines and internal buses of the IXP1200 processor. Two external interface units connect the on-chip shared buses to off-chip buses. One external bus is connected to Ethernet MACs and the other is an external memory bus. Teepee extends Ptolemy's Discrete Event modelling capabilities to describe and simulate architectural platforms such as this one.

5.3 Mapping

The first step of the mapping process is to choose a partitioning of the computational components of the application onto the processing elements in the architectural platform. Currently this is done manually. Here we choose to dedicate one processing element to each FromDevice packet source element. These elements perform the work of reading in packets from the Ethernet MACs and separating packet headers from data payloads. A third processing element will perform the CheckIPHeader and LookupIPRoute operations. Two final processing elements are used for the output portions of the router application.




Figure 3.12. Target Architectural Platform

The second step is to implement the connections between the application components in the Communication Implementation View. The partitioning we have chosen requires that some of the application components communicate with application components on different processing elements. These connections will be made using stacks of virtual receivers.

The communication semantics specified by the Click model of computation state that the flow of control follows packet headers as they are passed from actor to actor. Compared to other models of computation, these semantics are not restrictive and are therefore compatible with a broad range of behavioral types. However, the broadcast nature of the shared buses makes the hardware virtual receivers unusable directly. A protocol is required to provide the abstraction of multiple point-to-point communication links on top of this broadcast medium. Teepee provides such a protocol, called Address. The Address protocol works by prepending Network-on-Chip messages with a data field that contains a destination address. Instances of the Address protocol only keep messages if they match the destination address.

The communication semantics implemented by the Address protocol actors are an acceptable match for the Click application actors. A portion of the communication implementation model is shown in Figure 3.13. This model describes how the two FromDevice actors communicate with the CheckIPHeader actor using a shared bus and the Address protocol. The shared bus is exported into the Communication Implementation View as a SharedBusInterface virtual receiver with six ports (one for each processing element that the shared bus connects to). Instances of the Address protocol break the shared bus into separate point-to-point links. The remainder of the connections in the Communication Implementation View are made in a similar fashion.

5.4 Code Generation

A complete communication implementation model is a formal model of the interaction between the application and the architectural platform. Together with the model of the target architectural platform, it contains all of the information necessary to produce a concrete implementation of the system. The




Figure 3.13. Portion of Click Router Communication Implementation View

mathematical formalisms behind these models allow us to produce implementations in a correct-by-construction fashion. This process is called code generation. Examples of code generation targets include compiled-code simulators, synthesizable HDL, and executable machine code.

To produce executable machine code for the programmable platform, the code generation algorithm implements each object in the communication implementation model as a software object on the programmable platform. Application actors such as FromDevice and LookupIPRoute are realized as software processes. The same is true for protocol components such as the Address actors. Architectural virtual receivers (e.g. SharedBusInterface) are synthesized as device drivers for the programmable platform's communication architecture. Each software object is assigned to a particular processing element according to the partitioning made in the mapping step.

Our current code generation algorithm produces an individual C++ program for each processing element. These programs initialize software objects and manage threads of control. The initialization and execution semantics for all software objects are specified by the mathematical formalism that powers the Communication Implementation View. No additional operating system functionality is required. Standard C++ compilation techniques are used to produce machine code for each processing element.

6. Conclusion

Embedded system designers are faced with three important challenges. First, there is the challenge of modelling a heterogeneous concurrent application in an application development environment. Second, there is the challenge of capturing the important capabilities of a programmable platform architecture in a programming model. The third challenge is to correctly map from the application development environment to the programming model of the target platform in a way that efficiently implements the application.

We believe that the key to overcoming these challenges is to properly model concurrency in both the application and the architectural platform. In this chapter, we have introduced a new disciplined design methodology that focuses on formally modelling concurrency using models of computation. In this approach, application developers use the abstractions provided by models of computation to model heterogeneous, concurrent applications. At the same time, models of computation are used to abstract the parallelism of an architectural platform in a programming model. The use of models of computation in both the application development environment and the programming model gives designers a common theoretical framework for comparing application concurrency and architectural concurrency. This enables a disciplined approach to the solution of the mapping problem. Designers use the Communication Implementation View to abstract the mapping problem as a communication networking problem. Then, Network-on-Chip principles are employed to close the implementation gap efficiently and correctly.

Moore's law drives the exponential growth in the complexity of embedded system architectures. The latest generation of heterogeneous, concurrent architectures are accompanied by the latest generation of heterogeneous, concurrent applications. Without a disciplined design methodology, designers will have to resort to *ad hoc* techniques to implement concurrent applications on complex architectural platforms - a doubtful proposition. Using the appropriate abstractions to correctly model concurrency, designers can tackle the design challenges put forward by these systems.

References

- [1] E. Kohler et al. The Click Modular Router. *ACM Transactions on Computer Systems*. 18(3), pg. 263-297, August 2000.
- [2] The MathWorks, Inc. SIMULINK User's Guide. 1993.
- [3] E.A. de Kock et al. YAPI: Application Modeling for Signal Processing Systems. *Proc. of the Design Automation Conference*. June 2000.
- [4] G. Berry and G. Gonthier. The Esterel synchronous programming language: Design, semantics, implementation. *Science of Computer Pro-*

- gramming*. 19(2), 1992.
- [5] D. Harel. Statecharts: A Visual Formalism for Complex Systems. *Science of Computer Programming*. 8(3), pg. 231-274, 1987.
- [6] K. Keutzer, S. Malik et al. System-level design: orthogonalization of concerns and platform-based design. *IEEE Transactions on CAD*. vol. 19, pg. 1523-1543. 2000.
- [7] N. Shah. Understanding Network Processors. *Master's Thesis*, Dept. of Electrical Engineering and Computer Science, Univ. of California, Berkeley. 2001.
- [8] Intel Corporation. IXP1200 Network Processor Datasheet.
- [9] G. Kahn. The semantics of a simple language for parallel programming. *Proc. of IFIP Congress*. August 1974.
- [10] E. A. Lee. Computing for Embedded Systems. *IEEE Instrumentation and Measurement Technology Conference*. May 2001.
- [11] B. Kienhuis, E. Deprettere et al. A Methodology to Design Programmable Embedded Systems. In LNCS series of Springer-Verlag, Volume 2268, *SAMOS: Systems, Architectures, Modelling and Simulation*. Eds. F. Deprettere et al. November 2001.
- [12] J. Davis, R. Galicia et al. Ptolemy II: Heterogeneous Concurrent Modeling and Design in Java. *Technical Report UCB/ERL No. M99/40, University of California, Berkeley*. 1999.
- [13] A. Mihal, C. Kulkarni et al. A Disciplined Approach to the Development of Architectural Platforms. *IEEE Design and Test of Computers*. November/December 2002.
- [14] C. A. R. Hoare. Communicating Sequential Processes. *Prentice Hall International Series in Computer Science*. 1985.
- [15] M. Sgroi, M. Sheets et al. Addressing the System-on-a-Chip Interconnect Woes Through Communication-Based Design. *Proc. of the Design Automation Conference*. pages 667-672, June 2001.
- [16] E. A. Lee and Y. Xiong. System-Level Types for Component-Based Design. *First Workshop on Embedded Software, EMSOFT2001*. October 2001.

Chapter 4

GUARANTEEING THE QUALITY OF SERVICES IN NETWORKS ON CHIP

Kees Goossens, John Dielissen, Jef van Meerbergen,
Peter Poplavko[†], Andrei Rădulescu, Edwin Rijpkema,
Erwin Waterlander, and Paul Wielage

Philips Research Laboratories, Eindhoven, The Netherlands

[†] *Technical University of Eindhoven, Eindhoven, The Netherlands*

Kees.Goossens@philips.com

Abstract

Users expect a *predictable quality of service* (QoS) of embedded systems, even for future, more dynamic, applications. System-on-chip designers use networks on chip (NOC) to solve deep submicron problems, and to divide global problems into local, decoupled problems. NOCs provide services through protocol stacks, and introducing *guaranteed* services enables IP re-use and platform-based design. It also provides globally predictable behaviour, as required by the user, when combining local, decoupled solutions. There are several levels of QoS commitment (correctness, completion, completion bounds), with increasing cost. A combination of guaranteed and best-effort (no commitment) services combines their respective attractive features: predictable behaviour, and good average resource utilisation. The *ÆTHEREAL* NOC is an example of this approach, and forms the basis of a QoS-based design style, as advocated in this chapter.

1. Future applications and systems on chip

In this section we raise the following question: *why and how must systems implementing future applications guarantee the quality of service* (QoS)? We look at each of the components in turn, in the context of embedded systems and systems on a chip (SOC). We then observe that new SOC architectures and design methods divide global problems into local ones, and rely on networks on chip (NOC) to compose local solutions. We answer the question, in Section 2, by presenting a synthesis of a QoS-based design style and networks on chip. Section 3 explains that QoS can be decomposed into several levels of

61

A. Jantsch and H. Tenhunen (eds.), *Networks on Chip*, 61–82.
© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

commitment, with different resource requirements. In Section 4 we introduce the *ÆTHEREAL* NOC, and compare it to other NOCs, in Section 5.

1.1 Characteristics of future applications

We are witnessing the *convergence* of previously unrelated application domains: computation (personal digital assistants, mobile computing), communication (telephone, videophone, networking), and multimedia (audio, photography, video, augmented and virtual reality). Convergence leads to increased functionality and heterogeneity, as previously unrelated functions are combined. Systems become more dynamic, or even unpredictable, as new algorithms seek to take advantage of the disparity between average and worst-case processing. Compressed audio and video data of MPEG2 is an example. Algorithms are also shifting to higher semantic levels, and the semantic complexity of data is less predictable than its volume. An example is the shift from constant pixel processing (scaling, edge enhancement, and so on), to MPEG4 object detection and synthesis, to face recognition and scenario detection. Finally, increased interaction with the environment also makes systems more dynamic. System functionality can be influenced by the current location (affecting communication or computation capabilities), environmental conditions (precipitation and multipath interference can affect mobile communication), and other systems in the proximity (e.g. car guidance, ad hoc networking).

Trends such as ambient intelligence and the networked home make future applications *embedded* and *pervasive*. Moreover, applications acquire a responsibility for the control of physical objects, such as home heating systems, cars, and so on. Real-time performance and safety are critical in many of these applications.

Below, we examine how users interact with these applications, and then what is required to design these systems.

1.2 Quality of service: a user view

Users expect a certain behaviour of applications; in other words, they must be *predictable*. While those expectations may be low, as is often the case for personal computers, a certain fitness for purpose is always assumed. Consumer electronics are subject to higher demands: a television must have a robust user interface and is not allowed to crash or be unresponsive. Expectations are stricter yet for real-time applications (e.g. involving audio and video, or control systems); a television must display at least 50 pictures of a constant quality per second, for example. The essence of quality of service (QoS) is therefore the offering of a predictable system behaviour to the user.

A central question is how to reconcile the dynamic, unpredictable nature of future applications with the requirements for predictable services.

1.3 Implementing future applications on chip

Having identified the characteristics of future applications and their QoS requirements, we now turn to the question of their implementation in SOCs. Moore's law describes the *exponential growth* over time of the number of transistors that can be integrated in an IC. It predicts that chips in 2010 will count over 4 billion transistors, operating in the multi-GHz range [1]. It is this abundance of computational power that has fuelled the convergence of application domains described above. However, Moore's law is only a prediction, and two major obstacles must be solved to make it a reality [2]. First, to use future VLSI technologies, several *deep-submicron problems* must be solved: the increasing disparity between transistor and wire speeds, power delivery and dissipation, and signal integrity. Second, the intrinsic computational power of an IC must not only be used efficiently and effectively, but the time and effort to design a system containing both hardware and software must also remain acceptable. The so-called *design productivity gap* states that the increase in our ability to design SOCs does not match Moore's law. To close the gap, system design methods to implement applications with the latest VLSI technology (including the definition of architectures, mapping applications to architectures, and programming) must harness exponential hardware resource growth in a scalable and modular manner.

The following two examples show that until recently, architectures and design methods at both the deep-submicron and system levels have been often been *global* in nature. This hampers scalability, and in the next section we show that approaches that compose local solutions are becoming popular.

The physical communication between intellectual property blocks (IP) has made use of a mix of local and global wires. When timing constraints of the design must be verified, IPs cannot be checked independently because they are interconnected; correcting a timing violation in one IP may invalidate the timing of another. This process does not necessarily converge to a solution, and is design specific (not re-usable). This global timing closure problem results from a tight (timing) coupling [3].

The second example considers the model of time. Until now, the dominant design style has been synchronous; a global notion of time has been implemented by globally synchronous clocking. Increasing processing variations on a single IC make this style untenable in the future [4]. This will impact clocking regimes, but also the programming model. When using shared memory, which has been the dominant method to communicate between tasks, if tasks are not tightly synchronised (requiring a common, global time frame) then task scheduling may lead to interference and unpredictable behaviour.

1.4 Composing local subsolutions

To avoid the exponential complexity of global methods and solutions, there is increasing interest in subdividing global problems into *local, decoupled* problems, and then composing the local solutions. This requires, foremost, *compositionality* to assemble a global solution from local ones (Figure 4.1). But solutions must also be *scalable* (and likely hierarchical, Figure 4.1(c)) because there will be an exponential number of local solutions. This applies at all




Figure 4.1. Divide global problem (a) into local problems that are composed flat (b) or hierarchically (c). Examples are GALS, local timing closure, local shared memory with global fifo communication or message passing, local busses or switches with global packet switching.

levels of SOC design: lay-out, timing verification, clocking, and programming. This trend commenced some time ago with IP-based re-use, platform-based design [5]. The interest in decoupling and composing local solutions can be observed in hierarchical lay-out and wiring, local clocking strategies such as GALS (globally asynchronous, locally synchronous) [6], software-programmable fixed IPs or tiles [7, 8], and synthesisable tiles for dedicated silicon [9] or FPGAs [10, 11]. At the task level, these scalable system architectures (including e.g. chip multiprocessing [12, 13, 14]) use tasks with local, independent address spaces and time frames that are composed by means of timing-independent fifo channels (for example, Kahn process networks for stream-based processing).

1.5 Networks on chip

All the approaches that advocate local solutions have a common reliance on a scalable and compositional communication medium to efficiently combine the large number of (hardware and software) IPs or subsystems in a working system. This challenge is addressed by networks on chip (NOC), which therefore play a pivotal role in future SOCs. NOCs help to solve both deep-submicron problems (timing closure, wiring, lay-out, etc.) and compositional design methods (services and protocol stacks). More detailed argumentation for the use of NOCs in future SOCs can be found in [15, 9, 16, 17, 2], and elsewhere in this volume.

In the following section we show that NOCs and QoS naturally combine to solve the two problems stated in this section: NOCs combine decoupled local solutions, and a service-based design style makes both the design process and the QoS of the resulting SOC more predictable.

2. NoCs and QoS: a synthesis

We advocate system design centred on NOCs and QoS for three reasons. First, services relieve the inherent tension between the *dynamic nature of future applications* and the user requirement for predictable services. This addresses concerns raised in Sections 1.1 and 1.2. Second, IP re-use and platform-based design aim to re-use applications and architectures by decoupling them. The use of NOCs and their associated network protocol stacks are a way to achieve this (cf. Section 1.5). Finally, it is our tenet that QoS, in particular *guaranteed services*, allow us to move towards *globally predictable, locally predictable* methods and solutions for SOC design. This, in combination with NOCs solves a need of recent architectures and design methods, introduced in Section 1.4. We discuss each point in turn.

2.1 Dynamic applications and predictable QoS

A QoS-based design style alleviates the demanding task for the system *designer*. He or she must plan systems that reliably provide users with the expected QoS for future applications. Moreover, systems must always be cost effective. Higher costs are acceptable when safety is critical, but they are always under pressure for consumer products, whether they be television sets or cars. Systems, especially SOCs, therefore have limited resources, and they must be shared and managed as the application unfolds its dynamic behaviour. Resource management depends on two phases: negotiation to obtain resources, followed by a steady state in which allocated resources are used. As applications become more dynamic, the first phase, *renegotiation*, will become more frequent. Thus, users can be given a predictable QoS by giving resource management and QoS a prominent place in system design.

2.2 Platform-based design, NoCs, and QoS

The aim of *platform*-based design [5, 18] is to reduce the cost of system design through re-use of applications and architectures. A platform decouples applications and system architectures, by defining a template architecture and programming model. In other words, by limiting the freedom in which an application can be implemented on an architecture, the interdependence of application and architecture is concentrated and reduced (Figure 4.2(a)). However, the convergence of applications entails an increasing diversity and dynamics in resource usage (such as communication and computation patterns), and this results in an increasing need for *differentiated services* [15]. It is therefore important that the platform offers the appropriate services (for communication, computation, power management, etc.). As a result, the communication infrastructure is a critical component of a platform, because it must solve the




Figure 4.2. Decouple applications and architectures by means of a platform (a), or network (e.g. NOC, Internet) services (b), or network protocol stacks (c).

apparent contradiction of implementing diverse application behaviours with application-dependent IPs in an application-independent manner. Using a NOC for the platform interconnect tackles both problems: it integrates heterogeneous IPs in a standard fashion (in other words, the NOC services largely define the platform, Figure 4.2(a)); and it naturally provides differentiated services by means of a (partially application-dependent) protocol stack [17].

Figure 4.2(b) shows that NOCs offer a small set of services ("NoC services"), on top of which different kinds of communication can be implemented (e.g. shared-memory cache-coherent traffic, message passing). The OSI transport layer [19] is a natural place to separate the application and architecture, because it is the first network-independent layer. The set of services should be small, yet allow the upper layers in the protocol stack to offer differentiated services [20]. This method has been used successfully in computer networks [21], where the internet protocol (IP) plays a similar role. Higher layers offer more specialised protocols such as TCP and UDP, and FTP and HTTP. For NOCs, there are several proposals [22, 23, 24], but convergence has yet to occur.

Concluding, NOCs help IP re-use and platform-based design to combine IPs, and NOC services aid in decoupling applications and architectures.

2.3 QoS enables predictable composition

As we have seen in Section 1.3, to fully exploit Moore's law, it is key to combine and control many local, perhaps autonomous, components in an efficient and flexible manner resulting in the required QoS. To localise timing, communication, data, scheduling, and so on, means that components (subsystems, IPs, tasks, threads, etc.) must first be identified. After composition, their interaction is controlled by means of arbitration, scheduling, or resource management. But we must address the apparent contradiction of providing *globally predictable* QoS (cf. Section 2.1) while current approaches tend to *globally unpredictable, locally predictable* (GULP!) regimes. How can global

timing guarantees be given when locally synchronous components are combined asynchronously [6]? How can global performance be quantified when tasks are combined in a latency-insensitive fashion [14, 7]? It is our tenet that QoS, in particular guaranteed services, allow us to move towards *globally predictable, locally predictable* methods and solutions. We elaborate this claim below, by first defining what role QoS plays, then stating how the local solutions (e.g. IP design) benefit, and finally how QoS eases the composition of local solutions.

2.3.1 Services. Effective steering of the limited resources in a NOC (cf. Section 2.1), and hence effective QoS, is contingent on a reliable re-activation of resources to instructions. In other words, a predictable or guaranteed behaviour of system components is a prerequisite to providing the user with expected QoS, at whatever level. The essence of QoS-based system design is to restrict the interaction between components to well-defined services. Two phases can be distinguished: (a) negotiation, followed by (b) resource steering or scheduling, and observation or inspection. As an example, consider a microprocessor which negotiates a connection to a memory with guaranteed bandwidth, but without a constraint on transaction ordering. The NOC that offers communication services will honour the request if has sufficient resources available, and will reject it otherwise. Later, the microprocessor may renegotiate its connection, e.g. to a higher bandwidth. The NOC will release the resources of the connection for use by other connections, when it is closed by the microprocessor.

2.3.2 Advantages for local solutions. QoS-based interaction has a number of advantages for IP design. By limiting component interaction to a set of well-defined services, their interfaces are simplified because there are fewer eventualities to take into account. Moreover, failures of service users (IPs) are concentrated at the reconfiguration points: after a service provider (such as a NOC) has committed to the request (such as a connection with bounded jitter), its provision can be relied upon. Similarly, IPs do not interfere with each other because the service provider and user have a local contract. This allows components to be designed and implemented in isolation. An advantageous side effect of negotiation is that the requirements of IPs must be stated explicitly, aiding the design process.

2.3.3 Advantages in composing local solutions. We now consider the impact of QoS-based interaction when combining IPs to obtain a SOC. First, services that are guaranteed to an IP are not affected by other IPs in the network, making reasoning about the IP in isolation possible. This is essential for a *compositional construction* (design and programming) of SOCs.

Moreover, SOCs can be more *robust*, because rather than relying on cooperation to share resources, resource management enforces the contracts between service providers and users. An IP that (maliciously or erroneously) does not adhere to a cooperative protocol, such as TCP/IP, cannot, therefore, disturb the system as a whole [25]. Failure is therefore local in space (one IP) and time (at one of its reconfiguration points).

Next, when components are combined, the performance (i.e. the QoS) of their composition must be validated. This can be done by analysing the complete system implementation. However, this analysis may be too hard, because the behaviour of individual components or their interactions are complex to model accurately (e.g. traffic analysis of NOCs [26], cache behaviour [27]). *Statistical* approaches are therefore frequently used [28]. Unfortunately resource requests often do not fit models (e.g. bursty traffic versus fractal or normal traffic distributions [26]) invalidating the verification. The oxymoron "statistical guarantee" does therefore not guarantee a QoS, but implies a (usually post hoc) analysis relying on a statistical model of the resources and their usage. Instead, we propose to validate the composition of the local solutions at the level of services (their interface), instead of having to open them up, and consider their (global) combination. This *abstraction* is essential in reducing the verification state space.

Finally, services can make QoS provision *architecture independent*, because how a component, such as a NOC, offers its services is not relevant. This removes the need to second-guess the inner workings of a component, because its services describe all that is required to know. Interaction becomes prescriptive (state what is required, i.e. what must happen), rather than prognostic (try to predict the service provider's behaviour) or reactive (act on how the service provider behaves).

We can therefore conclude that by abstracting local IP behaviours to their service requirements or provision, makes the global QoS of their composition more predictable and hence easier to reason about. This simplifies the design of SOCs.

3. On the cost of guaranteeing QoS

We have seen that a QoS-based approach is required to implement future applications, and to offer a predictable performance to users. In essence, offering a QoS requires a commitment. In this section, we present different levels of commitment, and their effect on predictability and cost in terms of resource usage. Although guaranteed services, which offer commitment, have many advantages over so-called best-effort services, which offer no commitment, we show that their combination is beneficial. The *ÆTHEREAL* NOC, described in the next section, is an example of that approach.

3.1 Different levels of commitment

As has been explained before, to offer a certain QoS with finite resources, the service provider and user negotiate to arrive at a contract, i.e. a *commitment* by the service provider to honour the request. If a commitment has been given, the service is *guaranteed*, otherwise it is a *best-effort* service. Commitment exists at several levels: 1) *correctness* and integrity of the result, if and when it is delivered. Examples are parity checking and error correcting codes for uncorrupted data transmission, and redundant computation with majority voting for safety-critical systems. 2) Promise of *completion* or delivery. This involves the cumulative availability (over time) of sufficient resources (e.g. memory, cpu cycles), and their performance (e.g. absence of deadlock or livelock). Note that a minimum number of resources may be required: for example, an in-place FFT needs to store all samples simultaneously, and a mobile communication may require a minimum battery charge to generate a sufficiently strong signal. 3) *Bounds* on the performance. Examples are the completion time (when is the result available), cumulative time to completion and its variations (e.g. bounded latency and jitter), and (peak and average) energy consumption.

Almost any form of commitment to progress leads to resource allocation (e.g. memory space, cpu cycles, communication bandwidth, battery power) and hence requires resource management. Moreover, levels of commitment depend on those below them. We illustrate this with two examples.

An IP connected to a NOC may not always be able to accept incoming data. Assuming its input buffers are finite, several solutions are possible, with direct consequences for the service level that can be offered. Packets that arrive at a full buffer are dropped, instantly precluding completion bounds such as latency and jitter guarantees. General computer networks typically use this approach. Alternatively, packets are not dropped and are left waiting in the network. Care then has to be taken to not introduce deadlock (e.g. if re-ordered packets that wait for free buffers do not overtake each other), or livelock (e.g. in deflection routing [29] packets keep moving, but may never arrive at their destination). In both cases, again, completion bounds may be in jeopardy. An approach to offer completion bounds (e.g. for a bounded communication latency) is to ensure packets never wait in the network by reserving buffer resources at the receiver, and by using end-to-end flow control to constrain the sender to never send more than the available space.

The second example concerns data transmission in unreliable media. This is a topic of importance in NOCs because wires suffer increasingly from interference, such as cross-talk and voltage drops. To ensure data is transported unchanged, it can be retransmitted when corrupted, or error correction can be used. The energy efficiency of both approaches is investigated by [30] for unreliable busses. What concerns us here, is that using retransmission takes a

variable, possibly unbounded, amount of time, whereas error correction takes place in a constant time. Therefore, time-related guarantees, such as minimum throughput, can only be given by the latter.

We conclude that the strictest guarantees, namely those involving performance bounds, pervade the system (all protocol layers): they cannot be grafted on as an afterthought [17, 31]. In NOCs we should and can choose communication protocols and styles to offer the required services, but, as the examples show, this requires vigilance.

3.2 Commitment and resource usage

The effect of the level of commitment on resource usage is illustrated in Figure 4.3. Suppose that the resource plotted vertically is bandwidth. For




Figure 4.3. (a,b) Commitment to completion or bounded completion is reflected in resource usage. (c) Combining guaranteed services with best-effort services, discussed in Section 3.3.

real-time performance, the requested bandwidth must be offered in the same interval, see Figure 4.3(a). Thus, as many resources must be available as the largest request (r_{RT} at time t_{RT}). With resources dimensioned for the average resource usage (r_{AVG} in Figure 4.3(b)), completion times may shift to the future: the peak resource request at time t_{RT} is only completed at time t_{AVG} , in the example. Depending on the required completion bound and variability in resource requests more or fewer resources can be added. However, more resources than r_{RT} do not improve performance, while fewer resources than r_{AVG} mean that no completion commitment can be given, because the backlog of requests keeps growing.

Architectures offering only best-effort services do not reserve resources, and hence can have a *better average resource utilisation*, at the cost of *unpredictable or unbounded worst-case behaviour*.

3.3 Both best-effort and guaranteed services

A service is guaranteed if a commitment is given, and best effort otherwise. This holds for individual services, not their ensemble. For example, data transport may be uncorrupted (commitment to correctness), and lossless (commit-

ment to delivery), and without throughput guarantees (best-effort throughput, i.e. no commitment to a completion bound). Moreover, a given service can be offered both with and without commitment, to flexibly use the available resources. For example, a data transport service can be offered both with and without completion bounds, to serve different users, in a single SOC. In Figure 4.3(c), the critical dark traffic, uses the guaranteed service, for real-time performance. The remaining $r_{RT} - r_{AVG}$ resources are, on average, available to other traffic, for example with less strict completion bounds, or with no completion bound (best-effort completion).

Figure 4.3(c) shows that when resources must be dimensioned for the worst-case for a given service commitment (e.g. guaranteed latency), the resources can also be used to give less stringent commitments (e.g. guaranteed delivery), or even a best-effort service. A combination of best-effort and guaranteed services gives the advantages of guaranteed services to only part of the system, but the available resources are used more efficiently. In the next section we show how this can be put in practice, in the *ÆTHEREAL* NOC.

4. The *ÆTHEREAL* network on chip

In this section we introduce the *ÆTHEREAL* NOC [17, 2]. It offers guaranteed services to obtain the advantages in composability, robustness, and so on of QoS-based design. These services are used for real-time and critical functions. The *ÆTHEREAL* NOC also provides best-effort services, to take advantage of their lower resource requirements and potentially better average performance. We describe the NOC services [24] in the next section.

In Section 4.2, we show how the services can be efficiently implemented, using a mix of time-division-multiplexed circuit switching, and packet switching [32, 31]. The programming model (connection creation and closing) and its advantages are also explained.

4.1 *ÆTHEREAL* Services

The *ÆTHEREAL* NOC offers differentiated services with *connections*. A connection describes communication between one master and one or more slaves, with an associated service level, such as fifo transaction ordering, and maximum latency. Connections must be *created* stating the requested service level; this is the negotiation phase. The NOC either accepts or rejects the request for the connection. Connection acceptance may lead to resource reservations in the NOC, e.g. buffers or a link bandwidth percentage. After usage, *closing* a connection frees the resources. Different connections are created and closed independently, possibly at different points in time. Configurations can be computed at compile time (i.e. off-line), or at run time. To ensure that

the programming model scales as NOCs become larger, negotiations can be distributed.

4.1.1 Transactions. Once a connection has been created, the master initiates *transactions* by means of *requests* which zero or more slaves *execute*, perhaps leading to a *response*. Examples of transactions are read, write, acknowledged write, test and set, and flush. By offering these transactions the *ÆTHEREAL* transaction model is similar to existing bus protocols, to ease migration of IP from current interconnects to NOCs. However, to be able to take full advantage of increased NOC performance, transactions can also be pipelined, split, and posted [24].

4.1.2 Connection Types. The *ÆTHEREAL* NOC can offer three kinds of connections (Figure 4.4). A *simple* connection contains a master and




Figure 4.4. Connection types.




Figure 4.5. A narrowcast connection.

a single slave. The master initiates transaction which the slave executes, possibly resulting in a response (e.g. for a read). On a *narrowcast* connection, containing a master and one or more slaves, the request from the master is sent to and executed by exactly one slave. An example of the narrowcast connection is shown in Figure 4.5, where the master performs transactions on an address space that is mapped on two memory modules. Depending on the transaction address, a transaction is executed on one of the two memories. A *multicast* connection is a connection between one master and one or more slaves, in which requests are duplicated and each slave receives a copy of those requests. Currently, in a multicast connection no return messages are allowed, because of the volume of traffic they may generate (i.e., one response per destination), and the increased complexity in the master (because individual responses from slaves must be merged into a single response).

4.1.3 Connection Properties. We distinguish the following services, or connection properties: 1) data integrity, 2) transaction ordering, 3) transaction completion, 4) connection flow control, and 5) connection

throughput, latency, and jitter. A connection can request any combination of these properties (e.g. a throughput guarantee, flow control, but no transaction ordering). Recalling the levels of commitment, of Section 3.1, properties 1 and 2, above, commit to correctness (including order) of results. Completion is guaranteed by property 3 and 4 (4 implies 3). Connection latency and jitter give completion bounds. We discuss each property in turn.

1) *Data integrity* means that data is transported unchanged. We assume that data integrity is solved at the data link layer; every connection therefore offers data integrity. However, as noted in Section 3.1, this must be done in a way that supports higher-level commitments, in particular bounds on transport completion (latency).

2) *Transaction ordering*. Transaction orders are only defined on a single connection; transactions on different connections may be transported and executed in any order. In general, responses and requests may be reordered during their transport in the network. This means that requests (responses) may arrive in a different order at a single slave (master) than they were sent. (Note that we refer to requests and responses of different transactions. Within a single transaction, requests and responses are always ordered as follows: the master sends a request, a slave receives the request and executes it, the slave sends the response, the master receives the response.) We refer to [24] for a detailed analysis of different orderings. Here it suffices to state that three connections orderings are useful. *Unordered* connections, in which no order is assumed between any request and response. *Locally ordered connections*, where requests sent by the master for each slave are delivered to that slave in order they were sent. Requests for different slaves are unordered. Responses are delivered to the master in the order the requests were generated. For example, on a narrowcast connection with multiple single-port memory slaves, the read and write order to a single memory is important, but not between memories. *Globally ordered connections*, in which requests for (responses from) all slaves are delivered to the slaves (master) in the order the requests were sent. This last ordering is not offered as a service, because local ordering is usually sufficient, and because the user can emulate global ordering by means of acknowledged transactions. Global ordering may be used when the order in which different devices (slaves) are programmed is of importance.

3) *Transaction completion* stipulates that transactions requesting a response (e.g. acknowledged write) guarantee that the master always receives a response. The response contains a) a report that the request was not delivered to the slave, or b) the response of the slave, after it successfully executed the request, c) a notification that the slave successfully executed the request, but the response was dropped, d) a report that the slave failed to execute the request. For connections with flow control, no data will be dropped, and transaction completion is automatic. Without flow control fewer resources are required in

the network, but data may be dropped. For example, it may be fine to drop a transaction, as long as you know it has been lost, so that you can resend it.

4) *Connection flow control* guarantees that data that is sent will fit in the buffers at the receiving end. End-to-end flow control is one of main techniques to avoid network congestion, if data cannot be dropped (cf. Section 3.1). If a slave is slower than its master, and the slave buffers fill up, then the master will be blocked until there is sufficient free space for a transaction. Similarly, a slave may be blocked by a slower master on the return path.

5) *Bounds on connection throughput, latency, and jitter.* The *ÆTHEREAL* NOC provides connections that guarantee a bandwidth per fixed time interval. Thus, a combined throughput, latency, and jitter guarantee is given. Depending the time interval, the latency and jitter bounds may be rather large. This is acceptable in many applications, such as audio and video streaming, where throughput is more important than latency. In the next section give the reasons for giving the throughput guarantee in this form.

The *ÆTHEREAL* services include transactions of bus protocols, to ease migration from current bus-based systems to NOCs. But where in OCP [23] and VCI [22] connections are used only to relax transaction ordering, we offer a more general connection property model. Not only are more properties considered, but the request and response communications can be independently configured (e.g. for flow control, and throughput guarantees), allowing more fine-grained resource management. Connection-based service provision therefore allows better differentiation of services, and allows users to make full use of NOC performance.

4.2 *Æthereal architecture*

The challenge of designing a NOC lies in finding a balance between the NOC services and their implementation complexity and cost. Moreover, as has been mentioned in Section 3.1, different services (levels of commitment) are dependent on each other. Offering time-related guarantees (a completion bound) influences the NOC to the core. The *ÆTHEREAL* NOC has both guaranteed and best-effort services, and an architectural challenge is how to combine these efficiently. In other words, how can *guaranteed worst-case behaviour be joined with good average resource usage*. In the following sections we describe two conceptually disjoint networks that each solve one part. An efficient combination is then presented. Both networks contain two components: routers and network interfaces. Routers transport data and can be described at the OSI network layer. As the *ÆTHEREAL* services are offered to IP at the transport layer and are end to end (master to slave and vice versa), network interfaces are required to bridge the network layer and transport layer views on communication.

4.2.1 A guaranteed-throughput router. To give time-related guarantees on a connection, such as throughput guarantees (on a finite time scale) or latency bounds, the interference of other traffic in the NOC must be limited and characterised. Circuit switching gives strong guarantees but at a high cost: connections have to be dimensioned for the worst case traffic. Time-division multiplexed circuit switching reduces this cost, but creating and closing of circuits still takes much time, and grows with the size of the network. The life-time of circuits must grow to amortise this cost [33]. Rate- and deadline-based scheduling [34, 35] can characterise the worst-case contention and offer bounds on latency, by regulating the inflow of data, but at the a high buffer cost.

The *ÆTHEREAL* NOC uses *contention-free routing*, which is based on a time-division-multiplexed circuit-switching approach, where one or more circuits are set up for a connection, which is assume to be relatively long-lived. Guaranteed throughput (GT) packets never use the same link at the same time, i.e. all contention is avoided. This can be achieved by controlling both the time GT packets enter the network, and their speed in the network. All routers logically have a common notion of time, embodied in a slot counter (see Figure 4.6). GT packets propagate at the rate of one router per slot counter increment. By regulating the time a GT packet is injected in the network, it is in effect scheduled to use each successive link in its path in a successive slot, see Figure 4.6. This method avoids the emptying and filling of circuits between




Figure 4.6. Three routers, all in slot $s = 3$, switch packets p_i of circuit c_i to outputs o_i , for $T(s, o_i) = c_i$. For each table, only the relevant columns (outputs) are shown. Input/output link pairs are labelled io_i .

switching of circuits. Conceptually it resembles input queuing with store-and-forward routing, which results in low buffering costs because GT packets are small. The end-to-end (network interface to network interface) latency is the number of hops multiplied by the size of the GT packet. This is minimal, once the packet has entered its slot. This model also allows multicast circuits. GT packets are headerless, and are routed by means of slot tables in every network interface and router. Section 4.2.3 explains how the slot tables are programmed.

4.2.2 A best-effort router. The best-effort router uses packet switching and has a more conventional structure. Our experiments [31] indicate that both input queueing with worm-hole routing or virtual-cut-through routing, and virtual output queueing with worm-hole routing are feasible, in terms of buffering costs. Input queueing uses fewer buffers, but suffers from head-of-line blocking. Virtual output queueing has a higher performance but at the cost of more buffers.

4.2.3 A combined GT-BE router. The logically separate guaranteed (GT) and best-effort (BE) routers are combined (Figure 4.7(a)) to share the router resources (e.g. switch and data path, Figure 4.7(b)), and to obtain the advantages of both. The GT router offers a fixed end-to-end latency for its traffic, which has the highest priority, enforced by the arbiter. The BE router




Figure 4.7. Two views of the combined GT-BE router.

uses all the bandwidth (slots) that has not been reserved or is not used by GT traffic. This allows the sharing of links and data path. Resources are therefore never left unused, when there is data, cf. Figure 4.3(b). They are either used for critical traffic with real-time requirements (for which a completion bound has been given), or for best-effort traffic (without a completion bound).

To allow distributed and scalable programming of connections, the GT router slot tables (cf. Figure 4.6), are programmed by means of BE packets, see the arrow "program" in Figure 4.7(a&b). This can be done in a pipelined and concurrent (multiple simultaneous negotiations, also from the same the source), and distributed (active in multiple routers) fashion. Negotiations, resulting in slot allocations, can be done at compile time, and be configured deterministically at run time. Negotiations can also be done at run time, centrally or distributed.

4.2.4 A network interface. *ÆTHEREAL* network interfaces convert the OSI network layer services of the routers to transport layer services for the user. All connection properties (cf. Section 4.1) that are end-to-end are implemented by the network interfaces. These are: reordering, transaction

completion, and flow control. Unordered connections require that transaction identifiers are used, these are added at the network interfaces. For locally ordered connections, reordering buffers are required in the network interfaces. Since our routers do not reorder traffic in a connection, this is only required for narrowcast connections at the master for responses. Transaction completion also requires resource reservations in the network interfaces. Flow control serves to prevent overflow of buffers at the slave or master network interface. In a credit-based approach, this requires state (credits for the amount of space available), and additional communication (when data is consumed credits are returned to the producer). This is taken care of by the network interface.

IPs negotiate with network interfaces to obtain connections with certain properties. For this network interfaces may reserve resources, such as network interface buffers and credit counters, and slots in router tables.

5. Related work

Differentiated services have received attention in general computer networks in the context of ATM [36] and the Internet [25]. Real-time traffic schemes have been described using rate-based and deadline-based scheduling [34, 35].

The differences between single-hop on-chip communication such as busses and switches and NOCs are described in [24]. Bus protocols such as [22] often have time-division multiplexing and/or priorities added, such as MicroNetwork [37], to give throughput guarantees. To reduce the guaranteed but average high latency, statistical approaches can be used (e.g. Lotterybus [38]). The bound on completion time is then lost (cf. Section 3.1). Switches [39, 40], also single-hop interconnects, can also offer performance guarantees.

On-chip busses with bridges [41, 42, 43], but also [44], are potentially full-blown NOCs. By judiciously placing restrictions on topology, bridging, and transaction models, many problems that arise in general networks can be avoided. For example, an appropriate topology simplifies routing to a single path, and avoids transaction reordering. Limiting a master to a single outstanding transaction has the same effect. By not buffering in bridges, effectively an end-to-end circuit is set up per transaction, avoiding reordering and the need for end-to-end flow control. In both cases, the increasing latency of larger networks make these simplifying assumptions untenable in the future [33, 45].

The octagon interconnect [46] is an interesting combination of packet and circuit switching. An octagon corresponds to a single 8×8 router that uses circuit switching per transaction (or equivalently, packets of unlimited size), with a high performance for a connection in an octagon. An individual octagon can be made larger (the so-called core and edge node strategy), but this does not scale to larger networks for the reasons mentioned above. Packet switching can then be used instead, by using so-called bridge and member nodes strategy.

At this point all general NOC issues appear. This NOC seems a best-effort architecture, at least when multiple octagons are used.

The Spin [28, 29] NOC uses packet switching with worm-hole routing and input queuing in a fat tree topology. It is a scalable network for data transport, but uses a parallel network (bus) for control. It is a best-effort network, and is optimised for average performance (e.g. by the use of optimistic flow control coupled with deflection routing). Commitment is given for packet delivery, but latency bounds are only given statistically.

Dally et al [9] describes a preplaced mesh NOCs, with IPs that are synthesised in the tiles. It uses packet switching on lossless virtual channels with end-to-end flow control. No details are given on the programming model, but guaranteed throughput (and latency) services are supported.

In the context of FPGAs using bit-level circuit switching to implement inter-IP communication is expensive. A first optimisation is to use coarser circuits [47] to reduce the cost. A packet-switched interconnect with a predefined set of services allows sharing of communication resources, just like for ASICs, and enables dynamic reconfiguration [10, 11]. The NOC can be synthesised [10], but a preplaced hard-wired NOC is the next logical step.

New SOC architectures that rely on NOCs include chip multiprocessing [12, 13, 14, 7, 8, 48], to interconnect the homogeneous or heterogeneous tiles, and network processors [49].

6. Conclusions

We observe that, although future applications will be more dynamic, users expect embedded systems to behave predictably. A design style based on guaranteed quality of services (QoS) can solve this apparent contradiction. *Designers* of systems on chip (SOC) use networks on chip (NOC) to keep up with Moore's law. NOCs solve both deep-submicron problems (e.g. signal integrity), and narrow the design productivity gap (the efficiency with which we design SOCs) by dividing global problems into local, decoupled problems, e.g. GALS.

The combination of NOCs and QoS is natural, through network protocol stacks, and beneficial for several reasons. It enables IP re-use and platform-based design to decouple applications and architectures. QoS-based design also ensures that when global problems, such as clocking, are solved by combining local, decoupled solutions (e.g. GALS), the combination has a globally predictable behaviour, as required by the user.

There are several levels of QoS *commitment*, building on each other: correctness (e.g. uncorrupted transport), completion (e.g. packet delivery), and completion bounds (e.g. maximum latency). A service without commitment (a *best-effort* service, such as unbounded latency), can have a better average

resource utilisation than a *guaranteed* service, but at the cost of unpredictable or unbounded worst-case behaviour. The *combination of best-effort and guaranteed services is advantageous*: critical parts of the system are predictable, while resources are used more efficiently.

The ÆTHEREAL NOC combines best-effort and guaranteed services at different levels. Its programming model and architecture are also described. The ÆTHEREAL NOC is at the basis of a QoS-based design style, as advocated in this chapter.

References

- [1] Semiconductor Industry Association. *The International Technology Roadmap for Semiconductors*. 2001.
- [2] Paul Wielage and Kees Goossens. Networks on silicon: Blessing or nightmare? In *Euromicro Symposium On Digital System Design*, Dortmund, Germany, September 2002. Keynote speech.
- [3] T. N. Theis. The future of interconnection technology. *IBM journal of research development*, 44(3):379–390, May 2000.
- [4] Marcel J. Pelgrom, Hans P. Tuinhout, and Maarten Vertregt. Transistor matching in analog CMOS applications. In *IEDM*, pages 915–918, 1998.
- [5] K. Keutzer, S. Malik, A. Richard Newton, Jan M. Rabaey, and A. Sangiovanni-Vincentelli. System-level design: Orthogonalization of concerns and platform-based design. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 19(12):1523–1543, 2000.
- [6] Jens Mutersbach, Thomas Villiger, and Wolfgang Fichtner. Practical design of globally-asynchronous locally-synchronous systems. In *6th International Symposium on Advanced Research in Asynchronous Circuits and Systems (ASYNC)*, April 2000.
- [7] Paul Stravers and Jan Hoogerbrugge. Homogeneous multiprocessing and the future of silicon design paradigms. In *VLSI-TSA*, 2001.
- [8] Ken Mai, Tim Paaske, Nuwan Jayasena, Ron Ho, William J. Dally, and Mark Horowitz. Smart memories: A modular reconfigurable architecture. In *ISCA*, June 2000.
- [9] William J. Dally and Brian Towles. Route packets, not wires: On-chip interconnection networks. In *Design Automation Conference*, pages 684–689, June 2001.
- [10] Théodore Marescaux, Andrei Bartic, Dideriek Verkest, Serge Vernalde, and Rudy Lauwereins. Interconnection networks enable fine-grain dynamic multitasking on FPGAs. *FPL*, 2002. LNCS 2438.

- [11] Edson L. Horta, John W. Lockwood, David E. Taylor, and David Parlour. Dynamic hardware plugins in an FPGA with partial run-time configuration. In *Design Automation Conference*, June 2002.
- [12] Lance Hammond, Basam A. Hayfeh, and Kunle Olokotun. A single-chip multiprocessor. *IEEE Computer*, pages 79–85, September 1997.
- [13] Jaehyuk Huh, Stephen W. Keckler, and Doug Burger. Exploring the design space of future CMPs. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2001.
- [14] Eylon Caspi, André DeHon, and John Wawrzynek. A streaming multi-threaded model. In *Third Workshop on Media and Stream Processors (MSP-3)*, December 2001.
- [15] M. Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, and A. Sangiovanni-Vincentelli. Addressing the system-on-a-chip interconnect woes through communication-based design. In *Design Automation Conference*, pages 667–672, June 2001.
- [16] Luca Benini and Giovanni De Micheli. Networks on chips: A new SoC paradigm. *IEEE Computer*, 35(1):70–80, 2002.
- [17] K. Goossens, J. van Meerbergen, A. Peeters, and P. Wielage. Networks on silicon: Combining best-effort and guaranteed services. In *Proceedings of Design Automation and Test Conference in Europe*, pages 423–425, March 2002.
- [18] A. Ferrari and A. Sangiovanni-Vincentelli. System design: traditional concepts and new paradigms. In *International Conference on Computer Design*, pages 2–12, 1999.
- [19] J. D. Day and H. Zimmerman. The OSI reference model. In *Proceedings of the IEEE*, volume 71, pages 1334–1340, 1983.
- [20] K. G. W. Goossens and O. P. Gangwal. The cost of communication protocols and coordination languages in embedded systems. In Farhad Arbab and Carolyn Talcott, editors, *Coordination languages and models*, number 2315 in Lecture notes in computer science, pages 174–190. Springer Verlag, April 2002.
- [21] Steve Deering. Watching the waist of the protocol hourglass. In *6th IEEE International Conference on Network Protocols*, October 1998. Keynote speech.
- [22] VSI Alliance. Virtual component interface standard, 2000.
- [23] OCP International Partnership. Open core protocol specification, 2001.
- [24] Andrei Rădulescu and Kees Goossens. Communication services for networks on silicon. In Shuvra Bhattacharyya, Ed Deprettere, and Juer-

- gen Teich, editors, *Domain-Specific Processors: Systems, Architectures, Modeling, and Simulation*. Marcel Dekker, December 2002.
- [25] Vijay P. Kumar, T. V. Lashman, and Dimitrios Stiliadis. Beyond best effort: Router architectures for the differentiated services of tomorrow's internet. *IEEE Communications Magazine*, pages 152–164, May 1998.
- [26] Girish Varatkar. Traffic analysis for on-chip networks design of multimedia applications. In *Design Automation Conference*, June 2002.
- [27] Sungjoo Yoo, Kyoungseok Rha, Youngchul Cho, Jinyong Kung, and Kiy-oung Choi. Performance estimation of multiple-cache IP-based systems: Case study of an interdependency problem and application of an extended shared memory model. In *International Workshop on Hardware/Software Codesign*, May 2002.
- [28] Pierre Guerrier and Alain Greiner. A generic architecture for on-chip packet-switched interconnections. In *Proceedings of Design Automation and Test Conference in Europe*, pages 250–256, 2000.
- [29] Pierre Guérrier. *Un Réseau D'Interconnexion pour Systèmes Intégrés*. PhD thesis, Université Paris VI, March 2000.
- [30] Davide Bertozzi, Luca Benini, and Giovanni De Micheli. Low power error resilient encoding for on-chip data buses. In *Proceedings of Design Automation and Test Conference in Europe*, March 2002.
- [31] E. Rijpkema, K. Goossens, A. Rădulescu, J. van Meerbergen, P. Wielage, and E. Waterlander. Trade offs in the design of a router with both guaranteed and best-effort services for networks on chip. In *Proceedings of Design Automation and Test Conference in Europe*, March 2003.
- [32] Edwin Rijpkema, Kees Goossens, and Paul Wielage. A router architecture for networks on silicon. In *Proceedings of Progress 2001, 2nd Workshop on Embedded Systems*, Veldhoven, the Netherlands, October 2001.
- [33] André DeHon. Robust, high-speed network design for large-scale multiprocessing. A.I. Technical report 1445, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, September 1993.
- [34] Hui Zhang. Service disciplines for guaranteed performance service in packet-switching networks. *Proceedings of the IEEE*, 83(10):1374–96, October 1995.
- [35] Jennifer Rexford, John Hall, and Kang G. Shin. A router architecture for real-time communication in multicomputer networks. *IEEE Transactions on Computers*, 47(10):1088–1101, October 1998.
- [36] ATM Forum. *ATM User-Network Interface Specification*. Prentice Hall, July 1994. Version 3.1.
- [37] Drew Wingard. MicroNetworks-based integration for SOC's. In *Design Automation Conference*, 2001.

- [38] Kanishka Lahari, Anand Raghunathan, and Ganesh Laskhminarayana. Lotterybus: A new high-performance communication architecture for system-on-chip designs. In *Design Automation Conference*, June 2001.
- [39] Jeroen A.J. Leijten, Jef L. van Meerbergen, Adwin H. Timmer, and Jochen A.G. Jess. Stream communication between real-time tasks in a high-performance multiprocessor. In *Proceedings of Design Automation and Test Conference in Europe*, pages 125–131, 1998.
- [40] Paul J.M. Havinga. *Mobile Multimedia Systems*. PhD thesis, University of Twente, The Netherlands, February 2000.
- [41] Kyeong Keol Ryu, Eung Shin, and Vincent J Mooney. A comparison of five different multiprocessor SoC bus architectures. In *Euromicro*, 2001.
- [42] Tycho van Meeuwen, Arnout Vandecappelle, Allert van Zelst, Francky Catthoor, and Diederik Verkest. System-level interconnect architectures exploration for custom memory organizations. In *International Symposium on System Synthesis*, pages 13–18, October 2001.
- [43] Milenko Drinic, Darko Kirovski, Seapahn Meguerdichian, and Miodrag Potknojak. Latency-guided on-chip bus network design. In *Proc. of IEEE/ACM International Conference on Computer Aided Design*, pages 420–423, November 2000.
- [44] John Bainbridge and Steve Furber. CHAIN: A delay-insensitive chip area interconnect. *IEEE Micro*, 22(5), 2002.
- [45] Frederic Chong, Henry Minsky, André deHon, Matthew Becker, Samuel Peretz, Eran Egozy, and Frank F. Knight, Jr. Metro: A router architecture for high-performance, short-haul routing networks. In *International Symposium on Computer Architecture*, April 1994.
- [46] Faraydon Karim, Anh Nguyen, Sujit Dey, and Ramesh Rao. On-chip communication architecture for OC-768 network processors. In *Design Automation Conference*, June 2001.
- [47] André DeHon. Rent's rule based switching requirements. In *SLIP*, April 2001. Extended abstract.
- [48] Shashi Kumar, Axel Jantsch, Juha-Pekka Soininen, Martti Forsell, Mikael Millberg, Johnny Öberg, Tiensyrjä, and Ahmed Hemani. A network on chip architecture and design methodology. In *ISVLSI*, 2002.
- [49] David Whelihan and Herman Schmit. Memory optimization in single chip network switch fabrics. In *Design Automation Conference*, June 2002.

II

HARDWARE AND BASIC INFRASTRUCTURE

Chapter 5

ON PACKET SWITCHED NETWORKS FOR ON-CHIP COMMUNICATION

Shashi Kumar

Department of Electronics and Computer Engineering, School of Engineering, Jönköping University, Sweden

Abstract:

Designing a system on a chip with large number of cores poses many challenging problems. Designing a flexible on-chip communication network, which can provide the desired bandwidth and can be reused across many applications, is the key problem. In this chapter, we discuss how we can borrow and adapt the concept of packet switched communication from computer networks to design on-chip networks. In this new paradigm, the cores on the chip communicate among themselves by sending packets using a network of switches. We briefly describe the existing proposals, which have come up in the last couple of years. The network is built using a set of identical switches connected in regular topology. The important issues for design of the network are switch design, design of network access by cores and communication protocols. A particular protocol stack, called Nostrum, is described as a case study of protocols for on-chip communication. Ideas and tools can be borrowed from computer networks for evaluating on-chip networks.

Keywords: System on Chip (SoC), interconnection network, communication, protocols, IP cores

85

A. Jantsch and H. Tenhunen (eds.), *Networks on Chip*, 85–106.
© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

1. INTRODUCTION

In a few years time, it will be possible for designers to integrate more than fifty IP cores, each of the size of a small processor, along with many memory blocks of different types on a single chip. To avoid synchronization and clock skew problems, these multi-core SoCs will operate using Globally Asynchronous Locally Synchronous (GALS) paradigm, where each core will operate in a separate clock domain. The major challenge the SoC researchers face today is to come up with **structured, scalable, reusable and high performance** interconnection architectures/platforms for integrating a large number of heterogeneous cores on a single chip.

An interconnection platform is **structured**, if the underlying interconnection graph has nice properties for its layout and for mathematical analysis. A **scalable** architecture, in this context, implies that there is no inherent technical limitation in the architecture to build system with large number of cores. Reusability in all aspects of system design is key to fast time to design and test of complex and large SoCs. A **reusable** interconnection platform will require that the designer only concentrates on selection and specialization of cores and a new SoC is configured by plugging in the cores in a pre-designed interconnection network. The interconnection architecture must provide facilities for a large number of cores to efficiently exchange information. This implies that the interconnection architecture should allow high communication bandwidth between pairs of cores as well as concurrent communication among many pairs of cores. Many commercially important applications, especially in the area of multi-media communication and processing, require multi-core SoCs with high communication bandwidth among cores. Therefore, **high performance** in this context implies high communication bandwidth among cores. All these properties are related but one property in architecture does not imply all other. In fact, it is difficult to ensure all these desirable properties in an interconnection network simultaneously.

To meet these challenges, many research groups have concurrently proposed the idea of using a packet switched communication network, similar to one used in computer networks, for on chip communication. In this chapter, we discuss the development of this idea and various aspects of building on-chip networks.

The rest of this chapter is organized as follows. In section 2, we discuss the natural evolution of packet switched communication from earlier

methods used for connecting on chip components. In section 3, we define the necessary terms and describe concepts related to packet switched communication. In section 4, we briefly describe some important proposals for on chip communication and summarize the state of the art by providing a brief comparison among them. Section 5, describes in some details a concrete proposal for communication protocols for on-chip communication. Section 6, briefly discuss an attempt for evaluation of packet switched architecture and protocols. Section 7 summarizes the advantages and limitations of packet switched network

2. EVOLUTION OF ON-CHIP PACKET SWITCHED NETWORKS

The idea of using direct dedicated wires to connect pairs of communicating cores is not even considered, let alone used, since it will lead to large number of pins for every core, large routing time, large routing area, unpredictable delays in signals and unpredictable signal quality. Such a system is not scalable, leads to very low utilization of routing resources and hardly any part of it is reusable. In spite of these limitations, such a scheme to build SoCs has a potential of highest possible performance.

Bus-based interconnection is the most commonly used architecture for inter-core communication in present day SoCs. Buses help to time-share the wires among many communicating cores and lead to reduction of I/O pins in cores and simplifies wiring. Buses are more scalable and can achieve almost the same performance as compared to direct connection architecture. Multiple and hierarchical bus architectures have been proposed for providing efficient and high performance communication among many communicating units. Buses have been used to implement single board computers as well as small multi-processor systems [13]. Unfortunately, buses are also not scalable beyond some limit and may not provide required performance in large systems because the available communication bandwidth is shared among all the units connected to the bus.

Building bus based SoCs has another problem that it require interface designs or wrappers for every core before it can be connected to on-chip buses. The Virtual Socket Interface (VSI) Alliance [9] has tried to remove this problem by proposing a standard, called Virtual Component (VC) interface. Cores and buses need to be wrapped using this standard for easy interfacing. Figure 1 illustrates connection of a core to an on-chip bus. The main focus in VSI Alliance is to enhance the level of reuse from IP cores to SoC integration platform, which also include on-chip communication buses, software and SoC development methodology. This platform based

methodology has been strongly motivated by the success of reusing mother boards and templates for fast development of customized personal computers and application specific systems. Since the underlying interconnection architecture in these platforms is one or multiple buses, they are inherently unscalable in terms of size and performance beyond some limit.




Figure 1. Evolution of packet switched networks

The computer architecture researchers who have been searching for scalable architectures to build chip multiprocessor [19,20] systems and now they have met the SoC designers at a crucial cross-road (Figure 2). The scalability and success of internet is beckoning them to borrow the idea of using routers (or switches) based networks and packet based communication from computer network design for SoC design and multi-processor systems on chip.

Packet switched communication, besides providing theoretically unlimited scalability also provides possibility of standardization and reuse of communication infrastructure. These features are crucial to chip designers to lower time to design and time to market new products. Recently many researchers have made proposal of packet switched on-chip networks for connecting cores [1-6], [10]. Most of the researchers call the proposals as Network on Chips (NoC). NoC research is still in its infancy and no real SoC has been built using NoC ideas to the best of our knowledge. However, it is expected that future generations of products will have NoC based components.




Figure 2: A cross-road for SoC and Chip-Multiprocessor designers

3. PACKET SWITCHED NETWORKS: BASIC CONCEPTS AND COMPONENTS

3.1 Definitions

Most of the terminology used for on-chip packet switched communication is adapted, naturally, from computer network and multi-processor areas.

3.1.1 Switch

The switch is the main component of the interconnection architecture for routing information from a source to a destination. The switch has a functionality of transferring information from one of its input ports to any one/more of its output ports. Figure 3 shows a schematic of a switch. The number of input and output ports is generally small. The time gap between when the information enters the input port to the time when the information leaves the switch through an output port is called switch-delay. Information may need to be buffered before going out through the output port. Figure 3 shows a schematic diagram of a generic switch required to implement a packet switched communication. The switch has buffers to receive incoming packets and buffers to store packets before they are put on the output port. A crossbar switch helps to connect any output of an input buffer to any output buffer. A **switch control** helps to configure the crossbar and management

of the buffers. It may be noted that a specific design of the switch may include both input and output buffers or only one type of buffers. Researchers have also proposed buffer-less switch for packet switched communication [21].




Figure3: A Generic Switch

3.1.2 Switch Network

A switch network consists of an interconnection of many switches to enable a large number of units (cores) to communicate with each other. In this chapter, we will call the communicating cores as resources. A switch network is characterized by its topology, which is basically the underlying graph, in which a switch is represented by a node and a link between two switches is represented by an edge. Various topologies have been used/proposed for building switch networks for multiprocessor systems. Various types of meshes, hypercubes of various dimensions, shuffle network, butterfly network, binary trees and fat trees are a few examples of the topologies used for switching networks for inter-processor communication or communication between processors and memory modules [13]. Figure 4 shows some important topologies. A topology is more suitable for SoC design if it has an area efficient layout on a two dimensional surface. Therefore, mesh and binary tree topologies are more suitable than hypercube and butterfly topologies.




Figure 4: Typical Network Topologies

The inter-related parameters which characterize a switching network and which determine its performance are the network **diameter**, **connectivity**, **bandwidth** and **latency** [13]. Network **diameter** is the maximum number of intermediate nodes between a source and any destination pair. **Connectivity** refers to the number of direct neighbors of any switch node in the network. **Bandwidth** is a measure of maximum rate (in bits/second) of information flow in the network. It is possible to define bandwidth for a single link, single node or the aggregate bandwidth of the whole interconnection network. **Latency** is the time taken by a message to travel from the source to the destination.

3.1.3 Connection-oriented and Connectionless Communication systems

Communication is called connection-oriented if a route is established between the source and the destination for the information transmission. In the connectionless case information is encapsulated in entities, which travel through the network from source to the destination. In the traditional circuit switched communication, where a physical path is set up and is used for telecommunication, falls in connection-oriented communication. Packet based communication is possible on both types of systems. If it is possible to set up the network in such a way that the packets belonging to the same message reach the destination in proper sequence, then the communication scheme is called virtual channel or virtual circuit. Virtual circuit is obviously

implementable on connection-oriented systems; it is also implementable on connection less systems.

3.2 Communication Protocols

A set of rules and methods are required for transfer of information from one resource to another resource in any system. These rules are generally referred as protocols. Hardware handshaking with request and acknowledgement signals is a simple example of control signal protocols for communication between two units connected through direct wires. A large number of bus protocols exist which allow reliable communication among many connected units. The bus protocols provide rules for usage of shared communication wires for information exchange and for resolving conflicts among users. Besides specifying the timing relationship among various control and data signals, it also provides upper limits on the physical length and transfer rate or bandwidth [12]. Units can communicate with each other if they have the desired interface to the bus, which implement the bus protocols.

In a packet switched network, communication protocols determine how a resource is connected to the network as well as how the information flows from source to destination. The protocols used for packet switched communication are much more complex than bus protocols, and are generally partitioned and organized in many layers. The architecture defining the protocol layers is referred as a protocol stack. 7-layered OSI model has commonly been used as a standard and as a reference to study various protocol stacks used for communication in computer networks [12]. Most of the NoC researchers have also proposed layered stacks similar to OSI reference model.

3.2.1 Why layered communication?

In a layered protocol, various functions required for communication are divided hierarchically among various layers. Besides helping to manage complexity due to variations in requirements of various heterogeneous resources, it also provides mechanism to separate out various concerns. For example, in the case of a SoC in which many different processors and other cores have been integrated two cores can communicate messages, for example consisting of an array of pixels representing a picture, at a higher layer of protocol stack. At this layer of protocol the two cores need not worry about how the message will be broken up into packets and communicated through the physical network.

3.2.2 Functions of protocol layers in SoC context

The lower three layers of the OSI reference model, namely physical layer, data link layer and network layer, can be used with very little change in their functionality for on-chip packet switched communication. Because of the limited scope of on-chip communication, the top four layers are normally compressed into two layers. Below, we define the functionality of a five layered protocol stack for on-chip communication [2], [16].

- **Physical Layer:** This layer is concerned with physical characteristics of the physical medium used for connecting switches and resources with each other. In the context of SoC, it specifies voltage levels, length and width of wires, signal timings, number of wires connecting two units etc.
- **Data Link Layer:** This layer has the responsibility of reliably transferring information across the physical link. The layer provides functionality of transferring one word of information from one node to a connected node without error. Since the two connected units may be working asynchronously, the data link layer takes care of hardware synchronization besides error detection and correction. Data link layer may also include data encoding or data rate management for controlling power consumption etc.
- **Network Layer:** The network layer provides the service of communicating a packet from one resource to another using the network of switches. Its functionality includes buffering of packets and taking routing decisions in the intermediate switches.
- **Transport Layer:** This layer has the responsibility of establishing end-to-end connection and delivery of messages using the lower layers. Therefore its functionality includes packetization of a message at the source and depacketization or assembly of packets into a message at the destination node.
- **Application Layer:** For on chip communication, the relevant functionality of upper three layers of OSI reference model can be merged into this layer. Important services in this layer include message synchronization and management, conversion of data formats at receiver and application dependent functionality.

3.3 Communication backbone

In the context of on-chip communication, the communication backbone is the communication infrastructure provided for integrating cores for

developing SoCs. The infrastructure consists of the following three components:

- Network of switches
- Three lower layers of communication protocol stack
- Communication services

The communication services corresponding to the lower three levels of protocol stack need to be implemented in every switch of the network. The services corresponding to the upper layers of protocol stack need to be implemented in the resources.

3.4 Network Access

The other important aspects of on-chip communication are the methods required to prepare resources to be integrated to the system. This includes methods to check whether the resource meets the physical, layout, electrical and timing constraints for integration. A wrapper or an interface needs to be added to the resource so that its internal working becomes transparent to the network and other resources in the network. We call this wrapper as Resource Network Interface (RNI).

3.4.1 Resources

The resources in the network could have any functionality. Therefore, a resource could be a standard processor or DSP core, memory of any type, an ASIC block implementing a sequential or parallel behavior, or a reconfigurable hardware block. The obvious requirement is that the resource should be implementable using the same technology as the switch network and its implementation should meet the constraints on metal layers and voltage levels. Different NoC architectural proposals may add a few more constraints on its physical size and layout and its clock frequency.

3.4.2 Resource Network Interface

As mentioned above, every resource requires a wrapper for interfacing to a given packet switched network. Figure 6 shows a generic resource network interface (RNI). This interface has the same purpose as a network card in a computer, which must be connected to the internet. RNI must implement at least the lower three layers of protocol stack. This means that RNI enables a resource to communicate with the network using packets. The internal design of an RNI can be partitioned into two parts. The part connected to

the network is resource independent and the other part connected is resource dependent. The resource independent part can be reused for interfacing any new resource. It depends on its design that is on the three lower level protocol layers implemented by the switch. In an ideal situation, a switch should not need to distinguish an RNI from other switches to which it is connected in the network.

The resource dependent part will depend on the signals available on I/O pins, like data and address bus widths, control signals etc. If a set of cores follows the same standard, like having standard I/O buses and control signals, then the same RNI can be used for all the resources in this set. The major functionality of resource specific part of RNI includes:

- Packetization and depacketization of information.
- Data encoding for error detection and correction

The resource independent part takes care of timing, buffering and synchronization aspects during data communication.




Figure 5: Resource Network Interface

4. EXISTING PROPOSALS

In this section we describe some proposals for building packet switched communication for on-chip communication. As the area of NoC is in its infancy, our aim is to illustrate the direction and the state of thinking in this area rather than completeness. There are four aspects of the packet switched network that are important. These are: the link, topology of the network, the packet structure and the communication protocol stack.

4.1 Scalable, Programmable, Integrated Network (SPIN)

Guerrier and Greiner have proposed a generic architecture for on-chip packet switched communication, called SPIN[6]. Their aim was to demonstrate that this idea is already feasible in the existing technology and can soon be adapted for building SoCs. They propose two one directional 32-bit data paths between two switches in the network. They also propose a Fat-tree topology for the switch network because of its lower diameter as compared to a mesh, and because of its efficient VLSI layout. A packet in SPIN consists of a sequence of variable number of 32-bit words, with the first word acting as a header and following words containing the payload. The packet ends with a special word, which includes checksum. It allows datagram approach implying that the packets belonging to the same message can follow different route in the network.

The protocol stack used in SPIN implements the message passing communication model. The protocol stack is divided in layers such that the lowest layer has functionality similar to the three lower layers of OSI reference model. The lowest layer is fully hardwired and is implemented in wrappers and switches. Other layers can be built over this layer to support models for data flow streams and address space [6].

They have demonstrated the implementation feasibility of building a router in 0.13 μ process, which can be used to connect 128 units (resources) and will provide a peak bandwidth of 1.82 Tbits/sec. The SPIN research group is actively engaged in solving the problem of higher latency in handling memory access using packet switched network.

4.2 On-Chip Micro-Networks [1]

Benini and De Micheli [1] in their conceptual paper on NoC, predict that packet switched micro-networks will be essential to handle complexity of SoC designs of the future. They classify the protocol stack layers for micro-networks in three parts. The lowest part of the protocol deals with issues related to physical wiring. The second part deals with the architectural and logical design of the micro-network and will correspond to data link, network and transport layers. The highest part of the protocol stack layers will take care of application and system issues. The researchers rightly conclude that the major effort in the micro-network design will go in specializing the protocols to suit the application and its performance requirement. No concrete details are specified for the network, packet or protocols.

4.3 Network on Silicon [5]

Philips Research Laboratory has proposed a communication network, called Network on Silicon, for integrating IP blocks and memory modules on the same chip which provides predictable and guaranteed performance [5]. The research group in Philips realized that any proposed communication network should be flexible and efficient. The network must support guaranteed-throughput (GT) for real time applications and best-effort (BT) communication where throughput is not guaranteed, but no data is lost. They propose a connection oriented crossbar based routers to build the network for GT routing. The data rate of the link is assumed to be same as that of the routers. Every router maintains a routing table for crossbar configuration. Every entry in the table, called slot, corresponds to an input to output connection.

For GT communication pairs, slots are reserved in the routing tables of the required routers. The un-reserved slots are used for BE traffic using packet switched communication. Therefore, it is possible to provide high performance for a few communicating pairs in a flexible general architecture.

4.4 Linköping SoCBUS [23]

The SoCBUS tries to add the benefits of circuit switched communication to packet switched communication. They define the concept of Packet Connected Circuit (PCC). The idea is similar to a virtual circuit. A path is initially searched and set up between the communicating IPs using datagram based packet switching and then the path is locked. After that the path works in a circuit switched mode for the rest of the communication period. Such a scheme can lead to high throughput, short latency and deterministic communication behavior. But such a scheme may lead to restricted routing flexibility for the rest of the communication traffic. A demonstrator chip is planned to be built to show the effectiveness of the scheme.

4.5 KTH-VTT Network on Chip

Royal Institute of Technology (KTH), Stockholm and VTT Electronics, Oulu have jointly developed a NoC platform [10] for developing large and complex systems on a chip. This platform consists of a two dimensional mesh of switches. Resources are placed in the slots formed by the switches. A direct 2-D layout of switches and resources is assumed providing physical-architectural design integration. Each switch is connected to one resource and four neighboring switches, and each resource is connected to

one switch. A resource could be any computing or memory core, or it could be a configurable or fixed hardware block as long as it fits physically into the slot. The communication network supports packet switched communication using both datagram as well as virtual path based approach.

In order to reduce the packet delay in the network and reduce the hardware cost of the switch, researchers have designed a buffer-less switch with a specialized efficient dynamic routing algorithm. The algorithm uses the awareness about local traffic around the switch to take a good routing decision [21]. A five-layered protocol stack, called Nostrum, is proposed to implement layered communication. The switch in the network implements the lower three layers of protocol stack. We will discuss more details of the protocol stack in the next section.

4.6 Discussion

All the proposals discussed above have the common thread that they support packet based communication, are scalable and are capable of handling heterogeneity in cores. Various proposals differ in network topologies. Network on Silicon [5] and SoCBus [23] provide special mechanisms in the network to ensure that it is possible to set up high performance communication paths between, at least, a few pairs of resources in the network. These proposals make packet switched communication suitable for data intensive real-time applications. KTH-VTT [10] emphasizes the importance of physical level-architectural level design integration. Although all resources in their proposal must fit physically in a fixed size slot, they propose the concept of a **Region** to accommodate larger resources like shared memories [10]. All proposals assume a processor type resource has direct connection to access its program or data memory. It will be too slow to access such memories through packet switched network. However, a large memory shared among many resources will be treated as a separate resource, which is accessed, by other resources using general packet switched network.

At this stage nobody has reported a comparison of these proposal regarding cost or performance.

5. NOSTRUM: A CASE STUDY OF A PROTOCOL STACK FOR NOC

5.1 Protocol Layers

In this section, we discuss some details of the protocol stack proposed for on chip communication network proposed for KTH-VTT [17]. The ideas proposed in Nostrum are directly applicable, or easily adaptable, to other on-chip communication networks. Like the OSI reference model, the communication in Nostrum is layered with different responsibility for different layers. A clear separation of functionality of the layers in the protocol stack allows independent development and reuse of resources, on-chip network and applications.

Nostrum has three compulsory layers, namely, physical, data link and network layers. A transport layer is also defined to establish and maintain communication between source and destination resources. These lower three layers put together provide the service of delivering a packet of data from one resource to any other resource in the network. Nostrum supports both datagram as well as virtual circuit approach of packet delivery. Below we define the functionality of the lower four layers [17] for on chip communication.

Physical Layer

The physical layer models the physical link between the switches and provides the service of moving a word of data from output of one switch to the input of the connected switch. The signal degradation, data error are modeled by this layer.

Data Link Layer

The data link layer is responsible for reliable communication of data through a link with the specified synchronization, error detection/correction, and flow control.

Network Layer

The network layer is responsible for delivery of packets from the source resource through the network of switches to the destination resource. Since in a datagram approach, different packets belonging to a message can follow different routes and can reach out of order, the network layer provides possibility of dynamic routing algorithm as well as buffering of packets in the switches. For virtual circuit approach, the network layer sets up and