

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS INC.,
Petitioner

U.S. Patent No. 7,190,698

**DECLARATION OF INGRID HSIEH-YEE, PH.D,
UNDER 37 C.F.R. § 1.68 IN SUPPORT OF PETITION
FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

I.	Introduction	4
A.	Qualifications and Professional Experience	4
B.	Scope of This Declaration.....	7
C.	Evidence Considered in Forming My Opinions	8
D.	Overview of MARC Records, OCLC, and WorldCat	10
II.	Authenticity and Public Availability of <i>Xiao</i> (Ex.1005).....	15
A.	Authentication.....	15
B.	IEEE Copy of <i>Xiao</i>	17
C.	Library of Congress Copy.....	17
D.	Library of Congress Records	19
E.	Copyright Registration Record	22
F.	University of Wisconsin Copy	23
G.	Records of the University of Wisconsin-Madison Library System ...	24
H.	Usage Record	26
I.	Summary of My Opinion on <i>Xiao</i>	27
III.	Authenticity and Public Availability of <i>Schelén-2</i> (Ex.1007)	28
A.	Authentication.....	28
B.	University of Michigan Library Copy	28
C.	University of Michigan Library Records.....	30
D.	University of California Berkeley Library Copy	33

E.	University of California Berkeley Library Records.....	35
F.	Usage Record.....	38
G.	Summary of My Opinion on <i>Schelén-2</i>	39
IV.	Authenticity and Public Availability of <i>Schelén-1</i> (Ex.1032)	39
A.	Authentication.....	39
B.	Library of Congress Copy.....	40
C.	Library of Congress Records	42
D.	Copyright Registration Record	46
E.	Usage Record.....	47
F.	Summary of My Opinion on <i>Schelén-1</i>	47
V.	Conclusion.....	48

I. Introduction

I, Ingrid Hsieh-Yee, Ph.D., do hereby declare as follows:

1. I have been retained as an independent expert witness on behalf of Cisco Systems, Inc. (“CISCO”) for proceedings before the United States Patent and Trademark Office (“PTO”) regarding the authenticity and public availability of certain prior art references in relation to US Patent No. 7,190,698.

2. I am being compensated for my work in this matter at my customary hourly rate. I am also being reimbursed for any reasonable expenses associated with my work and testimony in this investigation. My compensation is not contingent on the results of my study, the substance of my opinions, or the outcome of this matter.

A. Qualifications and Professional Experience

3. My complete qualifications and professional experience are described in my academic *Curriculum Vitae* (**Appendix A**). The following is a brief summary of my relevant qualifications and professional experience.

4. I was a professor in the Department of Library and Information Science at the Catholic University of America for 32 years and retired at the end of August 2022. I have experience working in an academic library, a medical library, and a legislative library. I hold a Ph.D. in Library and Information Studies from the University of Wisconsin-Madison and a Masters in Library and Information Studies from the University of Wisconsin-Madison.

5. I am an expert on library cataloging and classification and have published two editions on this subject, *Organizing Audiovisual and Electronic Resources for Access: A Cataloging Guide* (2000 and 2006). I taught a variety of courses, including Cataloging and Classification, Advanced Cataloging and Classification, Organization of Internet Resources, Organization of Information, Digital Content Creation and Management, Internet Searches and Web Design, Information Literacy Instruction, Advanced Information Retrieval and Analysis Strategies, and The Information Professions in Society. I am familiar with metadata schema design and implementation. In my teaching I covered the design and implementation of metadata in databases, search engines, digital repositories, digital libraries, and digital archives. I also covered how information organization affects the discovery and access to digital resources on the Internet. My research interests cover cataloging and classification, information organization, metadata, information retrieval, information architecture, digital collections, scholarly communication, user interaction with information systems, and others.

6. I am very familiar with a library cataloging encoding standard known as the “Machine-Readable Cataloging” standard, also known as “MARC,” which became the national standard for sharing bibliographic data in the United States by 1971 and the international standard by 1973. I am also personally familiar with the Online Computer Library Center (OCLC), which is the largest bibliographic

network of the world, with more than 534 million records and thousands of member institutions (many of which are libraries of some type) in more than 100 countries. I am also personally familiar with WorldCat, which is a free web portal for users to explore more than 10,000 libraries worldwide. I provide a brief overview of MARC records, OCLC, and WorldCat below at Section D.

7. I am also personally familiar with many online catalogs, databases, and search engines. In preparing for this declaration, I used the following authoritative information systems to search for records:

- Google Scholar (<https://scholar.google.com>)
- IEEE Xplore Digital Library
(<https://ieeexplore.ieee.org/Xplore/home.jsp>)
- Online catalog of the Library of Congress (<https://catalog.loc.gov/>)
- Online catalog of the Linda Hall Library
(https://catalog.lindahall.org/discovery/search?vid=01LINDAHALL_INS_T:LHL)
- Online catalog of the Stanford University Library
(<https://searchworks.stanford.edu/>)
- Online catalog of the University of California Berkeley Library
(<https://www.lib.berkeley.edu/>)

- Online catalog of the University of Michigan Library
(<https://search.lib.umich.edu>)
- Online catalog of the University of Wisconsin Library System
(<https://search.library.wisc.edu/search/system>)
- Public catalog of the US Copyright Office (<https://cocatalog.loc.gov/>)
- WorldCat (<https://www.worldcat.org>)

B. Scope of This Declaration

8. I have been asked to offer expert opinions on the authenticity and public availability date of the following documents:

- Xiao, X., & Ni, L. M., (1999, March), Internet QoS: A big picture, *IEEE Network*, vol. 13, no. 2, pp. 8-18, available at <https://cs.wmich.edu/~alfuqaha/Fall07/cs6030/lectures/QoS-Paper.pdf>, **Ex.1005** (“Xiao”);
- Schelén, O., & Pink, S., (1998, January), Resource sharing in advance reservation agents, *Journal of High Speed Networks*, vol. 7, nos. 3-4, pp. 213-228, **Ex.1007** (“Schelén-2”);
- Schelén, O., & Pink, S., (1997, November), An agent-based architecture for advance reservations, in *Proceedings of 22nd Annual Conference on Local Computer Networks*, pp. 451-459, published by IEEE, **Ex.1032** (“Schelén-1”).

C. Evidence Considered in Forming My Opinions

9. In the preparation of this declaration, I have reviewed the documents referenced below and any other documents I reference herein, and each of these is a type of material that experts in my field would reasonably rely upon when forming their opinions:

- (1) The document referenced above in Section I.B;
- (2) Library of Congress copy of *Xiao*, obtained on October 7, 2022, **Appendix 1005A**;
- (3) Bibliographic and MARC records for *IEEE Network* that contains *Xiao*, available at <https://lccn.loc.gov/87649810> from the online catalog of the Library of Congress, accessed and obtained on September 30, 2022, **Appendix 1005B**;
- (4) Copyright registration record for vol. 13, issue 2 (March/April 1999) of *IEEE Network* that contains *Xiao*, available at <https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?v1=17&ti=1,17&Search%5FArg=IEEE%20Network&Search%5FCode=TALL&CNT=25&PID=sLdxh8jqgWLtuFlbhh-tWH6ufeD&SEQ=20220930214354&SID=1>, accessed and obtained on September 30, 2022, **Appendix 1005C**;
- (5) University of Wisconsin copy of *Xiao*, obtained on October 7, 2022, **Appendix 1005D**;
- (6) Bibliographic and MARC records for *IEEE Network* that contains

Xiao, available at <https://search.library.wisc.edu/serial/999583751502121> from the online catalog of the University of Wisconsin-Madison Libraries, accessed and obtained on September 30, 2022, **Appendix 1005E**;

- (7) Selected early citations to *Xiao*, obtained from Google Scholar, **Appendix 1005F**;
- (8) Xiao, X., & Ni, L. M., (1999, March), Internet QoS: A big picture, *IEEE Network*, vol. 13, no. 2, pp. 8-18, published by IEEE, **Appendix 1005G**;
- (9) University of Michigan Library copy of *Schelén-2*, obtained on December 22, 2022, **Appendix 1007A**;
- (10) Bibliographic and MARC records for *Journal of High Speed Networks* that contains *Schelén-2*, available at <https://search.lib.umich.edu/catalog/record/990030687990106381> from the online catalog of the University of Michigan Library, accessed and obtained on December 19, 2022, **Appendix 1007B**;
- (11) University of California Berkeley Library copy of *Schelén-2*, obtained on December 21, 2022, **Appendix 1007C**;
- (12) Bibliographic and MARC records for *Journal of High Speed Networks* that contains *Schelén-2*, available at https://search.library.berkeley.edu/permalink/01UCS_BER/1thfj9n/alma991061649799706532 from the online catalog of the University of California Berkeley Library, accessed and obtained on December 21, **Appendix 1007D**;

- (13) Library of Congress copy of *Schelén-1*, obtained on November 1, 2022, **Appendix 1032A**;
- (14) Bibliographic and MARC records for *Proceedings, Conference on Local Computer Networks* that contains *Schelén-1*, available at <https://lccn.loc.gov/98657356> from the online catalog of the Library of Congress, accessed and obtained on October 27, 20202, **Appendix 1032B**;
- (15) Copyright registration record for proceedings of the 22nd *Conference on Local Computer Networks* that contains *Schelén-1*, available at <https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?v1=8&ti=1,8&Search%5FArg=Conference%20on%20local%20computer%20networks&Search%5FCode=TALL&CNT=25&PID=2jCmxoCvf5Dv4yD91n9krxeS50X3Y&SEQ=20221103221511&SID=1> from the public catalog of the U.S. Copyright Office, accessed and obtained on November 3, 2022, **Appendix 1032C**;
- (16) Selected early citations to *Schelén-1*, obtained from Google Scholar, **Appendix 1032D**.

These records are identified and discussed in this Declaration. Experts in the field would reasonably rely on the data described herein to form their opinions.

D. Overview of MARC Records, OCLC, and WorldCat

10. MARC is the primary communications protocol for the transfer and storage of bibliographic metadata in libraries. Experts in my field reasonably rely

upon MARC records when forming their opinions. A MARC record consists of 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 follows. **Appendix B** is a true and correct copy of Parts VII to X of “Understanding MARC Bibliographic: Machine-Readable Cataloging” (<http://www.loc.gov/marc/umb/um07to10.html>) from the Library of Congress that explains commonly used MARC fields. For example, the personal author of a work is recorded in Field 100, the title is recorded in Field 245, publisher information is recorded in Field 260, the physical characteristics of a publication are recorded in Field 300, and topical subjects are recorded in the 650 fields.

11. OCLC was founded in 1967 to promote and support library cooperation. According to the “Third Article, Amended Articles of Incorporation of OCLC Online Computer Library Center, Inc.,” 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” (source: <https://www.oclc.org/content/dam/oclc/membership/articles-of-incorporation.pdf>).

12. OCLC members can contribute original cataloging records in MARC to the system or derive cataloging records from existing records, an activity referred to as “copy cataloging.” When an OCLC participating institution acquires a work, it can create an original MARC record for the work in OCLC’s Connexion system (a system for catalogers to create and share MARC records), and the system will automatically generate a code for the date of record creation in the yymmdd format in Field 008, and the creating library’s OCLC symbol is recorded in subfield “a” of the 040 field. Once the MARC record is in Connexion, it becomes available to other OCLC members for adoption to their local online catalogs (*i.e.*, copy cataloging).

13. After a MARC record is created in Connexion, it also becomes searchable and viewable on WorldCat, which is a free web portal for users to explore more than 10,000 libraries worldwide. The record in WorldCat, however, is not presented in MARC fields. Instead, the data elements are labeled to help users interpret the record. Thus, the information stored in MARC records in Connexion is available to the interested public through the user-friendly WorldCat web portal.

14. Library online catalogs gained acceptance in the early 1980s and many libraries migrated their systems to the World Wide Web in the mid-1990s. Library online catalogs are based on MARC records that represent their collections in order

to help the public understand what materials are publicly accessible in those libraries. Most libraries with online catalogs have made their catalogs freely available on the Web. These online catalogs offer user-friendly search interfaces. Strong user interest in keyword searches and the popularity of Google have led to the “googlization” of library search systems. As a result, many library catalogs now provide a single search box for users to conduct keyword searches, with additional support for searches by author, title, subject terms, and other data elements such as ISBN (International Standard Book Number). Library catalogs these days also offer features for users to narrow their search results by language, year, format, and other elements. Many libraries display MARC records on their online catalogs with labels for the data elements to help the public interpret MARC records. Many libraries also offer the option to display MARC records in MARC fields.

15. Libraries create MARC records for works they acquire, including books, serials, motion pictures, and publications in other formats. MARC records are representations of materials held by libraries and are critical for the discovery and access to those materials. Monograph cataloging is fairly common in libraries, and most libraries make a newly cataloged monograph available to the public soon after the cataloging work is completed, usually within a week. Libraries can create original cataloging records or use an existing record in OCLC to create a copy cataloging record. As soon as the cataloging record is completed, it is added to the

library's online catalog for users. If the record is an original record, it is also entered into OCLC WorldCat. If it is a copy cataloging record, the library's holding symbol is attached to the existing original record in OCLC WorldCat to facilitate searching and interlibrary loan.

16. The cataloging of serials and the serial check-in process are discussed here to show how libraries usually provide access to newly received serial issues. According to the glossary of the *RDA: Resource Description and Access* cataloging standard, a serial is "a mode of issuance of a manifestation issued in successive parts, usually bearing numbering, that has no predetermined conclusion. A serial includes a periodical, monographic series, newspaper, etc." Because the publisher of a serial makes new issues of the serial available successively, a customary cataloging practice is to create one bibliographic record for the serial, and the serial record (encoded in MARC) typically provides information on the beginning date and frequency of the serial, not the dates of individual issues. In other words, libraries typically do not create MARC records for individual issues of a serial. Instead, they rely on a serial check-in system to track the receipt of new issues. A common check-in practice is to date stamp a new issue when it arrives. This practice has become automated since the late 1990s, and libraries now vary in how they share the receipt date of a new serial issue with the public. Some libraries use a date stamp,

some affix a label to indicate the receipt date, some pencil in the receipt date, and some do not provide the information to the public.

17. The serial check-in process usually takes less than an hour, and one of the steps involves placing a date stamp on the new issue to document the date the issue is checked in. After that, the holdings information of the serial is updated in the library's catalog so that users know which issues are available for request or access. After serial check-in is completed, the new issue is placed on the shelf with the previous issues of the serial. Libraries with a public periodical room typically place new issues in the periodical room for easy user access. Because information presented in serials often reflects latest discovery, a general practice of libraries is to make new issues of serials available for user access soon after they are checked in, usually within a week.

II. Authenticity and Public Availability of *Xiao* (Ex.1005)

A. Authentication

18. **Ex.1005** is a true and correct copy of "Internet QoS: A big picture," (*"Xiao"*), by Xiao and Ni, *IEEE Network*, vol. 13, no. 2 (March 1999), pp. 8-18, that I obtained from <https://cs.wmich.edu/~alfuqaha/Fall07/cs6030/lectures/QoS-Paper.pdf>. When I began my research on this work I searched Google Scholar and found a full text of this article from the University of Michigan. This copy is included as **Ex.1005** in this Declaration.

19. Page 1 of **Ex.1005** (internal page 8) shows the title, authors, and abstract of *Xiao*, and the bottom of this page shows that *Xiao* has a “1999” copyright date, with IEEE as the copyright holder, is part of a publication that has the ISSN (International Standard Serial Number) of “0890-8044” and appears in the “March/April 1999” issue of *IEEE Network*. **Ex.1005** shows *Xiao* has a total of 11 pages (from internal pages 8 to 18), including two tables, five figures, and 55 references and two more articles for additional reading.

20. **Ex.1005** is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. As discussed below, I also obtained copies of *Xiao* from the Library of Congress (**Appendix 1005A**), the University of Wisconsin (**Appendix 1005D**), and the publisher IEEE (**Appendix 1005G**). I examined these copies closely and concluded that **Ex.1005** is identical to the two library copies and the publisher copy. The only differences are that the library copies include front matters, the Library of Congress is color, and some pages in the publisher copy have artifacts from what appears to be the copying process, but the content of the four copies of *Xiao* are the same. I therefore see no reason to question the authenticity of **Ex.1005**.

B. IEEE Copy of Xiao

21. My research on Google Scholar also found that *Xiao* was indexed by the IEEE Explore Digital Library, so I searched the digital library and found a metadata record for *Xiao* at <https://ieeexplore.ieee.org/abstract/document/768484>. I followed the PDF link in this IEEE metadata record to purchase this article from the publisher. The IEEE copy is presented as **Appendix 1005G** in this Declaration.

22. Page 1 of **Appendix 1005G** (internal page 8) shows the title, authors, and abstract of *Xiao*, and the bottom of this page shows that *Xiao* has a “1999” copyright date, with IEEE as the copyright holder, is part of a publication that has the ISSN (International Standard Serial Number) of “0890-8044” and appears in the “March/April 1999” issue of *IEEE Network*. **Appendix 1005G** shows *Xiao* has a total of 11 pages (from internal pages 8 to 18), including two tables, five figures, and 55 references and two more articles for additional reading.

C. Library of Congress Copy

23. To determine the availability of *Xiao* in libraries, I searched WorldCat by “IEEE Network” to identify libraries holding this journal and retrieved a record at <https://worldcat.org/title/14393593> that showed the Library of Congress was one of the holding libraries. I searched their online catalog for records, and the search results confirmed the holdings information. I then obtained a scanned copy of *Xiao* from the Library of Congress through Wisconsin TechSearch (WTS), a documentary

delivery service based at the University of Wisconsin-Madison. The Library of Congress copy is presented as **Appendix 1005A** in this Declaration.

24. **Appendix 1005A** is a true and correct copy of *Xiao* obtained from the Library of Congress. Page 1 is the cover that shows “IEEE Network” is the journal title, the issue is “Vol. 13 No. 2” with a “March/April 1999” issue date, and the journal is also available online at www.comsoc.org. The cover shows this issue is a special issue on “Transmission and Distribution of Digital Video: Field Trials and Prototype Implementations” and the bottom of this page shows IEEE is the publisher and this journal is a publication of “the IEEE Communications Society in cooperation with the IEEE Computer Society and the Internet Society.” Page 3 is the table of contents that shows *Xiao* begins on page 8. This page also carries a date stamp of “Copyright Office APR 8 1999 Library of Congress” that indicates the Library of Congress received this issue of *IEEE Network* that contains *Xiao* on April 8, 1999. As discussed above, a common library practice is to put newly received journal issues through the serial check-in process, then process them quickly, usually within a week of serial check-in, to make the newly received issues publicly accessible. It is therefore my opinion that vol. 13, no. 5 of *IEEE Network* that contains *Xiao* would have been accessible at the Library of Congress within a week after receipt, meaning this issue would have been publicly accessible no later than April 15, 1999.

25. Page 4 is the publisher page that shows the ISSN of *IEEE Network* is “0890-8044,” the journal “is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc.” and this issue has a “1999” copyright date, with the Institute of Electrical and Electronics Engineers, Inc. as the copyright holder. Pages 6 to 16 are the *Xiao* article. I have closely compared the Library of Congress copy (**Appendix 1005A**) with the IEEE copy of *Xiao* (**Ex.1005**) and concluded that they have the same content. The only difference is the IEEE copy is black and white, while the Library of Congress copy shows color tables and figures.

26. The Library of Congress copy of *Xiao* (**Appendix 1005A**) has the same content as **Ex.1005** and the publisher’s copy (**Appendix 1005G**). It is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. Moreover, *Xiao* (**Appendix 1005A**) was found in the custody of a national library, a place where, if authentic, it would likely be found. Accordingly, I see no reason to question the authenticity of the Library of Congress copy of *Xiao*.

D. Library of Congress Records

27. **Appendix 1005B** is a true and correct copy of the bibliographic and MARC records for *IEEE Network* that contains *Xiao*. I personally located, identified, and obtained these records from the online catalog of the Library of Congress. These

are the type of records experts in my field would reasonably rely upon when forming their opinions.

28. Pages 1 to 4 of **Appendix 1005B** are the bibliographic record for *IEEE Network* that shows the main title, publisher, and other details of this publication. The “Item Availability” area on page 3 shows the call number for this journal is “TK5105.5.I324 Set 1,” users can request access to journal issues at the Jefferson or Adams Building Reading Rooms, and the library holds many older issues, including “v.3-v.17 (1989-2003)” that include the vol. 13, no. 2 (1999) issue containing *Xiao*.

29. Pages 5 to 7 are the MARC record for *IEEE Network*. The first six digits of Field 008 show the record was created on “861014” (*i.e.*, October 14, 1986), and the “c19879999” code following the record creation date indicates this journal is a continuing resource that began publication in 1987 and has no end date of publication. Subfields “a” and “c” of Field 040 show “NSD” (the National Serials Data Program of the U.S. ISSN Center) created the original record, and subfield “d” shows “DLC” modified the record later. According to the *Directory of OCLC Members* (<https://www.oclc.org/en/contacts/libraries.html>), “DLC” is the OCLC library symbol for the Library of Congress. Data in Fields 008 and 040 inform my opinion that the National Serials Data Program created the original serial record for this journal on October 14, 1986, and the Library of Congress later used the original

record to create their copy cataloging record.

30. Field 022 in the MARC record (**Appendix 1005B**) shows the ISSN of this journal is “0890-8044” and Field 245 shows the journal title. Three variant titles are provided in Fields 246 to help users discover this journal. Field 260 shows the Institute of Electrical and Electronics Engineers of New York, NY began publishing this journal in 1987, Field 362 shows the journal began with “Vol. 1, no. 1” in January 1987, and Field 310 shows publication frequency has been “Bimonthly” since 1988. Fields 710 show the Institute of Electrical and Electronics Engineers and the IEEE Communications Society are included as additional access points to help users discover this journal. The ISSN, title, publisher, and publication frequency match the information presented in the Library of Congress copy of *Xiao* (**Appendix 1005A**), and confirm this serial record represents *IEEE Network*.

31. Subjects of this journal are represented in Field 050 by a Library of Congress Classification (LCC) number of “TK5105.5” that represents general works about “Computer Networks,” in Field 082 by a Dewey Decimal Classification (DDC) number of “004.6” that represents “Interfacing and communications,” and in Field 650 by one Library of Congress subject heading, with “Computer networks” as the main heading, followed by “Periodicals” as a form heading encoded in subfield “v” (\$v) to indicate the topic is treated in a journal.

32. The MARC record (**Appendix 1005B**) shows that *IEEE Network* is a long-running journal and the MARC record has made it searchable at the Library of Congress at least by April 1999, if not earlier. Interested users would have been able to conduct subject searches to discover this journal. They would also have been able to discover and access the journal by the ISSN, journal title, variant titles, IEEE and IEEE Communications Society.

E. Copyright Registration Record

33. **Appendix 1005C** is a true and correct copy of the copyright registration record for six issues of volumes 13 of *IEEE Network*. I obtained this record from the public catalog of the United States Copyright Office in the Library of Congress by first accessing the <http://www.copyright.gov/> website, selecting the hyperlink under “Search Copyright Records,” and performing a journal title search to retrieve the record. I personally located, identified and obtained this record. This is the type of record experts in my field would reasonably rely upon when forming their opinions.

34. The “Type of Work” field identifies this publication as “Serial”. The title, serial publication year, serial key title, imprint, ISSN, frequency and copyright claimant match the information presented in the Library of Congress copy of *Xiao* (**Appendix 1005A**). The “Issues Registered” field shows “v. 13, no. 2, Mar-Apr99.; Pub. 1999-03-29; Reg. 1999-05-03; TX0004881643,” meaning v. 13, no. 2 that

contains *Xiao* was published on March 29, 1999, and registered with the US Copyright Office on May 3, 1999, with a registration number of TX0004881643.

F. University of Wisconsin Copy

35. My research on WorldCat also identified the University of Wisconsin-Madison as one of the holding libraries of *IEEE Network* that contains *Xiao*. After I verified the holdings information in that library's online catalog, I obtained a scanned copy of *Xiao* from that library via Wisconsin TechSearch (WTS). This copy is presented as **Appendix 1005D** in this Declaration.

36. **Appendix 1005D** includes a color cover that identifies this issue as vol. 13, no. 2 of *IEEE Network* that has an issue date of March/April 1999, a table of content, a copyright page and the *Xiao* article. Page 2 shows a date stamp of "Kurt F. Wendt Library College of Engineering APR 08 1999 UW-Madison, WI 53706" that indicates the engineering library of the University of Wisconsin-Madison received this issue on April 8, 1999. As discussed above, a common library practice is to put newly received journal issues through the serial check-in process, then process them quickly, usually within a week of serial check-in, to make the newly issues publicly accessible. It is therefore my opinion that vol. 13, no. 5 of *IEEE Network* that contains *Xiao* would have been accessible at the engineering library of the University of Wisconsin-Madison within a week after serial check-in was completed, meaning this issue would have been publicly accessible no later than

April 15, 1999.

37. I have closely compared **Appendix 1005D** with **Ex.1005**, the Library of Congress copy (**Appendix 1005A**) and the IEEE copy (**Appendix 1005G**) and concluded that these four copies have the same content. The only differences are that the University of Wisconsin copy (**Appendix 1005D**) includes the front matter (a color copy of the cover, the table of contents and the copyright page) and the images in this copy are clearer than some of the images in the publisher copy.

38. The University of Wisconsin copy of *Xiao* (**Appendix 1005D**) has the same content as the publisher's copy (**Ex.1005**). It is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. Moreover, **Appendix 1005D** was found in the custody of a university library, a place where, if authentic, it would likely be found. Accordingly, I see no reason to question the authenticity of the University of Wisconsin copy of *Xiao* (**Appendix 1005D**).

G. Records of the University of Wisconsin-Madison Library System

39. **Appendix 1005E** is a true and correct copy of the bibliographic and MARC records for *IEEE Network* that contains *Xiao*. I personally located, identified, and obtained these records from the online catalog of the University of Wisconsin-

Madison General Library System. These are the type of records experts in my field would reasonably rely upon when forming their opinions.

40. Pages 1 and 2 of **Appendix 1005E** are the bibliographic record with holdings information. Page 1 shows the journal is available online and also available for user request. The Verona Shelving Facility area shows this location has “v.13:no.1(1994:01/02)-v.14:no.6(2000:11/12),” meaning it holds the print copy of vol. 13, no. 2 that contains *Xiao*. The journal title, contributors, publishers, publication date and ISSN match the information presented in the University of Wisconsin copy of *Xiao* (**Appendix 1005D**) and the Library of Congress copy of *Xiao* (**Appendix 1005A**), confirming this record represents *IEEE Network*.

41. Pages 3 and 4 of **Appendix 1005E** are the MARC record for the print version of *IEEE Network*. Fields 008 and Field 040 inform me that, like the Library of Congress record (**Appendix 1005B**), this University of Wisconsin record is based on the original MARC record created by the National Serials Data Program. I have closely compared this MARC record with the Library of Congress MARC record for *IEEE Network* and found them to have the same data for all the critical fields, including the Library of Congress control number in Field 010, the National Bibliographic Agency Control Number in Field 016, ISSN in Field 022, source of acquisition information in Field 037, the record authentication code in Field 042,

Library of Congress classification number in Field 050, Dewey Decimal Classification number in Field 082, title in Field 245, three variant titles in Fields 246, publisher information in Field 260, publication frequency in Field 310, beginning volume information in Field 362, the issuing body note in Field 550, a Library of Congress subject heading in Field 650, and two Fields 710 for the two contributing organizations.

42. This MARC record shows that *IEEE Network* is a long-running journal, and the MARC record has made it searchable in the online catalog of the University of Wisconsin General Library System no later than April 8, 1999. Interested users would have been able to conduct subject searches to discover this journal. They would also have been able to discover and access the journal by the ISSN, journal title, variant titles, IEEE and IEEE Communications Society.

H. Usage Record

43. Actual usage of a publication is reflected by the papers that make reference to it. The citation history on Google Scholar shows *Xiao* has been cited at least 1,459 times, and 102 citations appeared in 1999 and 2000. **Appendix 1005F** presents four citations published in 1999 and 2000 to demonstrate early usage. In addition, page 2 of **Appendix 1005F** is a screen capture of Google Scholar showing 1,459 citations have cited *Xiao*, and page 3 is a screen capture showing 102 citations to *Xiao* were made from 1999 to 2000. *Xiao* was published in March 1999, and the

earliest citation appeared in May of 1999, further supporting my opinion that *Xiao* was publicly accessible no later than April 15, 1999.

I. Summary of My Opinion on *Xiao*

44. I have compared the University of Michigan copy (**Ex.1005**), the Library of Congress copy (**Appendix 1005A**), the University of Wisconsin copy (**Appendix 1005D**), and the publisher copy (**Appendix 1005G**), and found them to have the same content. The University of Michigan copy of *Xiao* (**Ex.1005**) shows it was published in vol. 13, no. 2 (March/April 1999) of *IEEE Network*, the copyright registration record (**Appendix 1005C**) shows this issue was published on March 29, 1999, and the Library of Congress copy (**Appendix 1005A**) and the University of Wisconsin copy (**Appendix 1005D**) of this issue were received on April 8, 1999, at those libraries. Based on these data and my knowledge of common library cataloging and processing practices it is my opinion that *IEEE Network* would have been searchable in the online catalogs of the Library of Congress and the University of Wisconsin General Library System at least by April 8, 1999, if not earlier. Citation history shows that the earliest citation appeared in May of 1999.

45. Therefore, it is my opinion that *Xiao* (**Ex.1005**) was publicly available to interested persons exercising reasonable diligence by more than one year before the earliest claimed priority date of the '698 patent, but in no event later than April 13, 2000.

III. Authenticity and Public Availability of *Schelén-2* (Ex.1007)

A. Authentication

46. **Ex.1007** is a true and correct copy of “Resource sharing in advance reservation agents,” (“*Schelén-2*”), by Olov Schelén and Stephen Pink, *Journal of High Speed Networks*, vol. 7, nos. 3-4 (January 1998), pp. 213-228, that I obtained from the publisher IOS Press. I began my research on Google Scholar by the title of *Schelén-2* and found a record that led me to the publisher’s webpage for this article at <https://content.iospress.com/articles/journal-of-high-speed-networks/jhs148>. I followed the purchase link in this webpage to obtain this article. This publisher copy is presented as **Ex.1007** in this Declaration.

47. **Ex.1007** page 1 (internal page 213) shows the title, authors, abstract and seven keywords (“QoS, advance reservations, resource sharing, agents, admission control, simulation, virtual leased lines”) for this article. The bottom of this page shows this article appears in a work with an ISSN (International Standard Serial Number) of “0926-6801” that has a “1998” copyright date, and the IOS Press is the copyright holder. **Ex.1007** shows *Schelén-2* has a total of 16 pages, including 15 figures, two equations, and 21 references.

B. University of Michigan Library Copy

48. To determine the availability of *Schelén-2* in libraries, I searched WorldCat by the journal title “high speed networks” for records and retrieved a

record at <https://worldcat.org/title/973879779> that showed the University of Michigan Library and the University of California Berkeley Library are two of the holding libraries. I searched the online catalog of the University of Michigan Library and their records (<https://search.lib.umich.edu/catalog/record/99187257762906381>) confirmed the holdings information. I then obtained a scanned copy of *Schelén-2* from this library through Wisconsin TechSearch (WTS), a documentary delivery service based at the University of Wisconsin-Madison. The University of Michigan Library copy is presented as **Appendix 1007A** in this Declaration.

49. **Appendix 1007A** is a true and correct copy of *Schelén-2* obtained from the University of Michigan Library. Page 1 is the issue cover that identifies the issue as “Volume 7, Numbers 3,4, 1998” of *Journal of High Speed Networks* whose ISSN is “0926-6801.” The cover also shows this issue is on “Multimedia Networking.” Page 2 is the copyright page that shows this issue has a “1998” copyright date and the IOS Press holds the copyright. The “Aims and scope” area shows the journal serves as “an archive for papers describing original results of lasting significance in both the theory and practice of high speed networks” and is published quarterly. Page 3 (internal page 187) is the guest-editorial and carries a date stamp of “University of Michigan Mar 18 1999 Media Union Library” that indicates the library received this issue on March 18, 1999. Pages 4 to 19 (internal pages 213 to 228) are the *Schelén-2* article that shows this article has a total of 16 pages, including

15 figures, two equations and 21 references.

50. I have closely examined the University of Michigan Library copy. I also compared it with the publisher's copy (**Ex.1007**) and found them to have the same content. The University of Michigan Library copy (**Appendix 1007A**) is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. Moreover, this copy was found in the custody of a university library, a place where, if authentic, it would likely be found. Accordingly, I see no reason to question the authenticity of the University of Michigan Library copy of *Schelén-2* (**Appendix 1007A**).

C. University of Michigan Library Records

51. **Appendix 1007B** is a true and correct copy of the bibliographic and MARC records for *Journal of High Speed Networks* that contains *Schelén-2*. I personally located, identified, and obtained these records from the online catalog of the University of Michigan Library. These are the type of records experts in my field would reasonably rely upon when forming their opinions.

52. Pages 1 to 3 of **Appendix 1007B** are the bibliographic record for *Journal of High Speed Networks*. The top of the record shows this publication is a

“Serial” and a “Journal” and the IOS Press has published this journal since 1992 with vol. 1, no. 1. It also shows the journal is published quarterly and some issues are called special issues. Page 2 shows the ISSN of the print journal is “0926-6801” and the “Available at” field shows the journal is available online and in print; the print volumes of “1-18: 1992-2012” are shelved offsite with the call number of “TK 5101 .A1 J68”; and volume 7 that contains *Shelén-2* is “on shelf.”

53. Pages 3 to 7 are the MARC record for this journal. The first six digits of Field 008 show the record was created on “960717” (*i.e.*, July 17, 1996). Based on my experience with records from the University of Michigan Library, I understand this record date in Field 008 is the date when this record was created at this library. The “c19929999” code following the record creation date indicates this publication is a continuing resource that began publication in 1992 and has no end date of publication. Subfields “a” and “c” of Field 040 show “NED” created the original record, and subfield “d” shows “EYM” modified the record later for copy cataloging. According to the *Directory of OCLC Members* (<https://www.oclc.org/en/contacts/libraries.html>), “NED” is the OCLC library symbol for the library of the Northeastern University and “EYM” is the OCLC library symbol for the University of Michigan Library. Data in Fields 008 and 040 inform my opinion that the library of the Northeastern University cataloged this journal first by creating a MARC serial record, and the University of Michigan

Library used the original MARC record to create a copy cataloging record for their copy on July 17, 1996, making *Journal of High Speed Networks* discoverable in their online catalog from that date on.

54. Field 022 in the MARC record (**Appendix 1007B**) shows the ISSN of this journal is “0926-6801.” Field 245 records the journal title, and Fields 246 show two variant titles—“High speed networks” and “JHSN”—are included to help users discover this journal. Field 580 shows this journal is also available as an electronic journal. Fields 710 show IOS Press and OCLC FirstSearch Electronic Collections Online as two sources for access to the electronic version of this journal, and Fields 864 show two links to the electronic journal, and access requirements and issues available from each service provider.

55. Subjects of this journal are represented in Field 050 by a Library of Congress Classification (LCC) number of “TK5105” that represents general works on “Data transmission systems” and in Fields 650 by two Library of Congress subject headings, with “Telecommunication systems” and “Computer networks” as the main headings, each followed by “Periodicals” as a form subdivision encoded in subfield v (“|v”) to indicate the main topics are treated in a journal. These subject representations mean that users can conduct subject searches to discover this journal.

56. The date stamp on the University of Michigan Library copy of *Schelén-*

2 (**Appendix 1007A**) has the appearance of a typical library date stamp that shows the date of receipt. The “MAR 18 1999” stamp indicates the University of Michigan Library received vol. 7, nos. 3,4 (1998) of *Journal of High Speed Networks* containing *Schelén-2* on March 18, 1999. As discussed earlier, a common library practice is to put newly received journal issues through the serial check-in process, then process them quickly, usually within a week of serial check-in, to make the newly issues publicly accessible. It is therefore my opinion that this special issue on Multimedia Networking that contains *Schelén-2* would have become accessible at the University of Michigan Library within a week after receipt, meaning the copy would have been publicly accessible no later than March 25, 1999.

57. The MARC record (**Appendix 1007B**) shows that *Journal of High Speed Networks* is a long-running journal and the MARC record has made it discoverable at the University of Michigan Library since July 17, 1996, when the copy cataloging record was created. Interested users would have been able to conduct subject searches to discover this journal. They would also have been able to discover and access the journal by the title, variant titles, and the ISSN.

D. University of California Berkeley Library Copy

58. My research on WorldCat (<https://worldcat.org/title/973879779>) also identified the University of California Berkeley Library as a holding library of the journal containing *Schelén-2*. After I checked their online catalog to confirm the

holdings information, I obtained a scanned copy of this article from this library through Wisconsin TechSearch (WTS), a document delivery service based at the University of Wisconsin-Madison. The University of California Berkeley Library copy is presented as **Appendix 1007C** in this Declaration.

59. **Appendix 1007C** is a true and correct copy of *Schelén-2* obtained from the University of California Berkeley Library. This copy includes the issue cover of vol. 7, nos. 3-4 (January 1998) of *Journal of High Speed Networks*, the copyright page, table of contents, and the *Schelén-2* article. The cover and copyright page are identical to the corresponding pages in the University of Michigan Library copy (**Appendix 1007A**). The cover shows a call number label of “TK5105.5 J687 v. 7 no. 3-4 ENGI Display” and a date stamp of “The Library – UC Berkeley, Received on: 04-12-99, Journal of high speed networks.” The table of contents page shows this journal is abstracted or indexed by 13 abstracting and indexing services, which enable interested users to discover individual articles published in this journal. It also shows *Schelén-2* appears from pages 213 to 228 in volume 7, numbers 3,4 (1998) of this journal. Pages 4 to 19 of **Appendix 1007C** are the *Schelén-2* article. I have closely examined this copy and compared it with the publisher’s copy (**Ex.1007**) and the University of Michigan Library copy (**Appendix 1007A**), and concluded that these three copies contain the same *Schelén-2* article.

60. The University of California Berkeley Library copy (**Appendix 1007C**) is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. Moreover, this copy was found in the custody of a university library, a place where, if authentic, it would likely be found. Accordingly, I see no reason to question the authenticity of the University of California Berkeley Library copy of *Schelén-2* (**Appendix 1007C**).

E. University of California Berkeley Library Records

61. **Appendix 1007D** is a true and correct copy of the bibliographic and MARC records for *Journal of High Speed Networks* that contains *Schelén-2*. I personally located, identified, and obtained these records from the online catalog of the University of California Berkeley Library. These are the type of records experts in my field would reasonably rely upon when forming their opinions.

62. Pages 1 to 3 of **Appendix 1007D** are the bibliographic record for *Journal of High Speed Networks*. Page 1 of the record shows this publication is a peer reviewed “Journal” that is available in print at the “Northern Regional Library Facility NRLF (UCB) (TK5105.5 J687)” and online at four service platforms. Page 2 shows the library holds print issues “v.4(1995)-18:3(2011/12)” and volume 7 (1998) that contains *Schelén-2* is located at the Northern Regional Library Facility

NRLF (UCB) TK5105.5 J687 and loanable. Page 3 shows details of the record, including title, other titles, format, creation date, publisher, frequency, general notes, additional form available, subjects, related titles, identifier and source. The bottom of this page shows the URL of this bibliographic record.

63. Page 4 of **Appendix 1007D** is the MARC record. The first six digits of Field 008 show the record was created on “930129” (*i.e.*, January 29, 1993), and the “c19929999” code following the record creation date indicates this publication is a continuing resource that began publication in 1992 and has no end date of publication. Subfields “a” and “c” of Field 040 show “NED” created the original record. According to the *Directory of OCLC Members* (<https://www.oclc.org/en/contacts/libraries.html>), “NED” is the OCLC library symbol for the library of the Northeastern University. The OCLC library symbol of the University of California Berkeley Library (“CUY”) is not included in a subfield “d” in Field 040, meaning the library used the NED record for copy cataloging but did not modify the original record. Data in Fields 008 and 040 inform my opinion that the library of the Northeastern University cataloged this journal and created a MARC serial record on January 29, 1993, and the University of California Berkeley Library used the original MARC record to create a copy cataloging record later.

64. Because the University of Michigan Library MARC record (**Appendix**

1007B) and the University of California Berkeley Library MARC record (**Appendix 1007D**) are based on the same original record created by the Northeastern University Library, the two copy cataloging records have the same essential data, including the ISSN in Field 022, the OCLC number of “27344595” that represents the Northeastern University Library MARC record in OCLC Connexion, the journal title in Field 245, variant titles in Fields 246, publisher information in Field 260, physical description in Field 300, publication frequency in Field 310, initial publication in Field 362, a note in Field 500 on the source of the title information, a note in Field 515 indicating some issues are called special issues, and the same Library of Congress subject headings in Fields 650. A difference between these two university MARC records is that the University of Michigan Library uses “TK5105” as the call number, while the University of California Berkeley Library uses “TKTK5101.A1” as their call number for this journal. Another difference is that each university includes electronic journal access information that is specific to their campus.

65. The “Received on: 04-12-99” date stamp on the University of California Berkeley Library copy of *Schelén-2* (**Appendix 1007C**) has the appearance of a typical library date stamp that shows the library received vol. 7, nos. 3-4 (1998) of *Journal of High Speed Networks* containing *Schelén-2* on April 12, 1999. As discussed earlier, a common library practice is to put newly received

journal issues through the serial check-in process, then process them quickly, usually within a week of serial check-in, to make the newly issues publicly accessible. It is therefore my opinion that this special issue on Multimedia Networking that contains *Schelén-2* would have become accessible at the University of California Berkeley Library within a week after receipt, meaning the copy would have been publicly accessible no later than April 19, 1999.

66. The MARC record (**Appendix 1007D**) shows that *Journal of High Speed Networks* is a long-running journal and the MARC record has made it discoverable at the University of California Berkeley Library by at least April 1999, if not earlier. Interested users would have been able to conduct subject searches to discover this journal. They would also have been able to discover and access the journal by the title, variant titles, and the ISSN.

F. Usage Record

67. Actual usage of a publication is reflected by the papers that make reference to it. The citation history on Google Scholar shows *Schelén-2* has been cited 65 times, and a unique citation and a self-citation appeared in 1999. The unique citation is a dissertation completed in 1999: Chu, H. H. (1999). CPU service classes: A soft real-time framework for multimedia applications. University of Illinois at Urbana-Champaign.

G. Summary of My Opinion on *Schelén-2*

68. The date stamps on the University of Michigan Library copy (**Appendix 1007A**) and the University of California Berkeley Library copy (**Appendix 1007C**) of *Schelén-2* inform my opinion that copies of the journal issue containing this article were received by these libraries in March and April of 1999, respectively, and would have become publicly accessible about a week after the copies were received. Google Scholar's citation history shows a dissertation citing *Schelén-2* was completed in 1999.

69. Therefore, it is my opinion that *Schelén-2* (**Ex.1007**) was publicly available to interested persons exercising reasonable diligence by more than one year before the earliest claimed priority date of the '698 patent, but in no event later than April 13, 2000.

IV. Authenticity and Public Availability of *Schelén-1* (Ex.1032)

A. Authentication

70. **Ex.1032** is a true and correct copy of An agent-based architecture for advance reservations, ("*Schelén-1*"), by Olov Schelén and Stephen Pink, in *Proceedings of 22nd Annual Conference on Local Computer Networks*, pp. 451-459 that I obtained from the publisher IEEE. I began my research on the IEEE Xplore Digital Library and searched by the title of *Schelén-1* for record. The retrieved metadata record at <https://ieeexplore.ieee.org/abstract/document/631014> showed

this article was indexed by IEEE. I followed the PDF link in this IEEE record to purchase from the publisher. The IEEE copy is presented as **Ex.1032** in this declaration.

71. **Ex.1032** is a true and correct copy of *Schelén-1* that I obtained from the IEEE Xplore Digital Library, the digital library created and maintained by the Institute of Electrical and Electronics Engineers. Page 1 (internal page 451) shows the title, authors, three keywords (“Advance Reservations, Reservation Agents, Quality of Service), and abstract, and the bottom of this page shows this article has a “1997” copyright date, with IEEE as the copyright holder, and it appears in a publication that has an ISSN (International Standard Serial Number) of “0742-1303.” **Ex.1032** shows *Schelén-1* has a total of nine pages, including one figure and 31 references.

B. Library of Congress Copy

72. To determine the availability of *Schelén-1* in libraries I searched WorldCat by “Conference on local computer networks” for records and retrieved a record at <https://worldcat.org/title/17022071> that showed the Library of Congress was one of the holding libraries. I searched their online catalog for records, and the search results confirmed the holdings information. I then obtained a scanned copy of *Schelén-1* from the Library of Congress through Wisconsin TechSearch (WTS), a documentary delivery service based at the University of Wisconsin-Madison. The

Library of Congress copy is presented as **Appendix 1032A** in this declaration.

73. **Appendix 1032A** is a true and correct copy of *Schelén-1* obtained from the Library of Congress. Page 1 shows the spine of a bound volume showing a title of “Conference on Local Computer Networks Proceedings” and a call number of “TK 5105.5 .C66a 22nd Nov 2-5 1997.” Page 2 is the front cover that shows “Proceedings, 22nd Conference on Local Computer Networks, Minneapolis, Minnesota, November 2-5, 1997, LCN ‘97” and the IEEE Computer Society and the IEEE Computer Society Technical Committee on Computer Communications as conference sponsors. Page 3 is the back cover that shows this work was published by the IEEE Computer Society of Los Alamitos, CA, and it has four identifiers, including “PR08141” as the IEEE Computer Society Press Order Number, “97TB100179” as the IEEE Catalog Number, “0-8186-8141-1” as the ISBN (International Standard Book Number), and “0742-1303” as the ISSN (International Standard Serial Number). Page 4 is the title page that shows the same information as the cover. It also carries a date stamp of “Library of Congress DEC 02 1997 Copyright Office” and a barcode of a number “TX4-577-503” (also as “TX0004577503”) that has the appearance of a copyright registration number. Page 5 is the copyright page that shows this work has a “1997” copyright date, with The Institute of Electrical and Electronics Engineers, Inc. as the copyright holder. This page shows the same four identifiers listed on the back cover, and two more ISBNs.

Pages 6 to 10 are the table of contents and page 9 (internal page viii) shows *Schelén-I* begins on page 451 and is included in “Session XIII: Quality of Service II.” Pages 11 to 19 are the full text of this article. **Appendix 1032A** shows *Schelén-I* has nine page, including one figure and 31 references.

74. I have compared the Library of Congress copy of *Schelén-I* (**Appendix 1032A**) with the IEEE copy (**Ex.1032**) and found them to be identical. The Library of Congress copy is in a condition that creates no suspicion about its authenticity. Specifically, this copy is not missing any intermediate pages, the text on each page appears to flow seamlessly from one page to the next, and there are no visible alterations to the document. Moreover, this copy was found in the custody of a national library, a place where, if authentic, it would likely be found. Accordingly, I see no reason to question the authenticity of the Library of Congress copy of *Schelén-I* (**Appendix 1032A**).

C. Library of Congress Records

75. **Appendix 1032B** is a true and correct copy of the bibliographic and MARC records for *Proceedings, Conference on Local Computer Networks* that contains *Schelén-I*. I personally located, identified, and obtained these records from the online catalog of the Library of Congress. These are the type of records experts in my field would reasonably rely upon when forming their opinions.

76. Pages 1 to 4 of **Appendix 1032B** are the bibliographic record for *Proceedings, Conference on Local Computer Networks*. The top of the record shows this publication is a “Periodical or Newspaper” and the bibliographic record shows the meeting name, main title, publisher and starting date of publication, and other details. The “Item Availability” area on page 3 shows the call number for this serial is “TK5105.5 .C66a Set 1,” the status of the serial is “Not Charged” (meaning available for user request), and users can request access to journal issues at the Jefferson or Adams Building Reading Rooms. The “Older receipts” field on page 4 shows the Library of Congress holds “11-24 (1986-1999)” (meaning the library has the *22nd Annual Conference on Local Computer Networks* that contains *Schelén-I*).

77. Pages 5 to 7 are the MARC record for this serial. The first six digits of Field 008 show the record was created on “871124” (*i.e.*, November 24, 1987), and the “c19869999” code following the record creation date indicates this publication is a continuing resource that began publication in 1986 and has no end date of publication. Subfields “a” and “c” of Field 040 show “MYG” created the original record, and subfield “d” shows “DLC” modified the record later. According to the *Directory of OCLC Members* (<https://www.oclc.org/en/contacts/libraries.html>), “MYG” is the OCLC library symbol for the library of the Massachusetts Institute of Technology and “DLC” is the OCLC library symbol for the Library of Congress. Data in Fields 008 and 040 inform my opinion that the library of the Massachusetts

Institute of Technology cataloged this serial on November 24, 1987, and the Library of Congress used their record to create a copy cataloging record later.

78. Field 022 in the MARC record (**Appendix 1032B**) shows the ISSN of this serial is “2831-7742” and the ISSN of “0742-1303” in subfield “y” is an ISSN printed in a serial but assigned by the International Serials Data System to another serial. Field 111 shows the name of the conference is treated as the author of this serial and Field 245 shows the main title is “Proceedings” and the conference is responsible for this work. Fields 246 show that six variant titles are included to help users discover this serial. Field 260 shows the IEEE Computer Society Press of Washington, DC has published this serial since 1986. Field 362 shows the serial began with the “11th” conference in 1986, and Field 310 shows the frequency of this serial is “Annual.” A 500 field shows the publisher’s location has been “Los Alamitos, Calif.” since 1996, and Field 530 shows the serial is also “issued online.” Field 710 shows the Technical Committee on Computer Communications of the IEEE Computer Society is included as an additional access point to help users discover this serial. The ISSN, title, publisher, publication frequency and sponsor match the information presented in the Library of Congress copy of *Schelén-1* (**Appendix 1032A**), and confirm this serial record represents *Proceedings of the Conference on Local Computer Networks*.

79. Subjects of this serial are represented in Field 050 by a Library of Congress Classification (LCC) number of “TK5105.5” that represents general works about “Computer Networks,” in Field 082 by a Dewey Decimal Classification (DDC) number of “004.68” that represents “Local-area networks and other specific kinds of networks,” and in Field 650 by a Library of Congress subject heading that has “Local area networks (Computer networks)” as the main heading, followed by “Congresses” as a form heading encoded in subfield v (“|v”) to indicate the topic is treated at conferences. These subject representations mean that users can conduct subject searches to discover this serial.

80. The date stamp on the Library of Congress copy of *Schelén-1* has the appearance of a typical library date stamp that shows the date of receipt. The “DEC 02 1997” stamp indicates the Library of Congress received the proceedings of the *22nd Annual Conference on Local Computer Networks* on December 2, 1997. As discussed earlier, a common library practice is to put newly received journal issues through the serial check-in process, then process them quickly, usually within a week of serial check-in, to make the newly issues publicly accessible. It is therefore my opinion that this conference proceedings that contains *Schelén-1* would have become accessible at the Library of Congress within a week after receipt, meaning the copy of this work would have been publicly accessible no later than December 9, 1997.

81. The MARC record (**Appendix 1032B**) shows that *Proceedings Conference on Local Computer Networks* is a long-running serial and the MARC record has made it discoverable at the Library of Congress at least by December 1997, if not earlier. Interested users would have been able to conduct subject searches to discover this serial. They would also have been able to discover and access the serial by the ISSN, conference name, variant titles, and the sponsor.

D. Copyright Registration Record

82. **Appendix 1032C** is a true and correct copy of the copyright registration record for no. 22 of the IEEE Conference on Local Computer Networks that contains *Schelén-1*. I located this record from the public catalog of the United States Copyright Office in the Library of Congress by first accessing the <http://www.copyright.gov/> website, following the hyperlink “Learn More” under “Search Copyright Records,” and selecting “Official Public Catalog (1978-Present)” to search by the conference name to retrieve the record. I personally located, identified and obtained this record. This is the type of record experts in my field would reasonably rely upon when forming their opinions.

83. The “Type of Work” field identifies this publication as “Serial”. The title, serial publication year, serial key title, imprint, ISSN, frequency and copyright claimant match the information presented in the Library of Congress copy of *Schelén-1* (**Appendix 1032A**). The “Issues Registered” field shows “no. 22, 1997”

and “Pub. 1997-11-02; Reg. 1997-12-02; TX0004577503,” meaning the 22nd conference proceedings was published on November 2, 1997, and registered with the US Copyright Office on December 2, 1997, with a registration number of TX0004577503, which is the same as the label on the title page of the Library of Congress copy of *Schelén-1* (**Appendix 1032A**).

E. Usage Record

84. Actual usage of a publication is reflected by the papers that make reference to it. The citation history on Google Scholar shows *Schelén-1* has been cited 60 times, and seven citations appeared in 1998 and 1999. **Appendix 1032D** presents three citations to demonstrate early usage. In addition, page 2 of **Appendix 1032D** is a screen capture of Google Scholar showing 60 citations to *Schelén-1*, and page 3 shows a screen capture of seven citations from 1998 to 1999 that cited this article.

F. Summary of My Opinion on *Schelén-1*

85. The IEEE copy of *Schelén-1* (**Ex.1032**) shows this article was published in the proceedings of the *22nd Annual Conference on Local Computer Networks* in November 1997. The date stamp on the Library of Congress copy (**Appendix 1032A**) shows this volume was received by the library on December 2, 1997, the MARC record (**Appendix 1032B**) shows the serial has been discoverable at the online catalog of the library, and the copyright registration record (**Appendix**

1032C) shows this volume was published on November 2, 1997. Based on these data and my knowledge of common library cataloging and processing practices it is my opinion that the volume of proceedings of the *22nd Annual Conference on Local Computer Networks* that contains *Schelén-1* would have been discoverable in the online catalogs of the Library of Congress by at least December 1997, if not earlier. The copy of *Schelén-1* would have been publicly accessible in the library no later than December 9, 1997. Citation history (**Appendix 1032D**) shows that the earliest citations appeared in 1998.

86. Therefore, it is my opinion that *Schelén-1* (**Ex.1032**) was publicly available to interested persons exercising reasonable diligence by more than one year, but in no event later than April 13, 2000, the earliest claimed priority date of the '698 patent.

V. Conclusion

87. I hereby declare that all statements made herein on 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 made under penalty of perjury.

Date: 2-10-2023

Executed: Ingrid Hsieh-Yee
Ingrid Hsieh-Yee, Ph.D.

APPENDIX A

Ingrid Hsieh-Yee
Professor Emerita
Dept. of Library and Information Science
Catholic University of America
Washington, D.C. 20064
E-mail: hsiehyee@cua.edu

Education

Ph.D. 1990. Library and Information Studies, University of Wisconsin-Madison
Minors: Sociology and Psychology

M.A. 1984. Library and Information Studies, University of Wisconsin-Madison.

M.A. 1983. Comparative Literature, University of Wisconsin-Madison.

B.A. 1980. Foreign Languages and Literature, National Taiwan University.

Work Experience

Professor, School/Dept. of Library and Information Science, Catholic University of America,
1990-2022 (Assistant Professor, 1990-1996; Associate Professor, 1997-2004; Professor
2004-2022; Professor Emerita 2022)

Co-Chair, Dept. of Library and Information Science, Catholic University of America, June 2015-
August 2016.

Acting Dean, School of Library and Information Science, Catholic University of America,
January 2010-June 2012.

Cataloger, Dept. of Legislative Reference Library, Annapolis, Maryland, 1989-1990.

Lecturer, School of Library and Information Studies, University of Wisconsin-Madison, 1988.

Teaching Assistant, School of Library and Information Studies, University of Wisconsin-
Madison, 1986-1988.

Cataloger, Health Sciences Library, University of Wisconsin-Madison, 1984-1986.

Areas of Teaching and Research Interests

Information Organization and Access; Metadata; Cataloging & Classification; Information
Architecture; Information Retrieval; Digital Collections; Social Media; Scholarly

Communication; Information Behavior; Health Informatics; Human Computer Interaction; Usability Studies

Grants & Honors

Cultural Heritage Information Management Project. IMLS grant. Amount: \$498,741. Period: Aug. 2012 to July 2015. Co-PI with Dr. Youngok Choi.

D.C. Health Information Technology (HIT4): Building Capacity & Providing Access in Our Nation's Capital. Dept. of Labor H2B Training Grant. Grant amount: \$4,175,500. Grant period: Nov. 2011 to Dec. 2015. Partner with the Metropolitan School of Professional Studies of the Catholic University of America, Children's National Medical Center, D.C. Department of Employment Services, Holy Cross Hospital, Howard University, Center for Urban Progress, Providence Hospital, and Sibley Memorial Hospital.

Capital Health Careers Project. Department of Labor Healthcare Sector and Other High Growth and Emerging Industries Grant. Grant amount: \$4,953,999. Grant period: March 2010 – February 2013. Awarded to a group of healthcare organizations and educational institutions in Washington, D.C. Providence Health Foundation of Providence Hospital (Lead institution). Part of the grant supported the development of a Master's degree program in Information Technology with a concentration in Health Information Technology offered by the School of Library and Information Science.

The Washington D.C. School Librarians Project. IMLS grant. Grant amount: \$412,660. Grant period: Aug. 2007 – June 2011. The School partnered with the District of Columbia Public Schools (DCPS) and the District of Columbia Library Association to educate and mentor school media specialists for the DCPS system. PI, Jan. 2010 to June 2011.

SIG Member of the Year, American Society for Information Science and Technology (2009).

Most Outstanding Paper of *OCLC Systems & Services* (2001).

ALISE Research Grant (2001).

Most Outstanding Paper of *OCLC Systems & Services* (2000).

Research Grant from ERIC (1999-2000).

Best Research Paper Award; Association for Library and Information Science Education (1998).

Research Grants, Catholic University of America. 1991, 1992, 1993, 1996, 1998, 1999, 2004, 2005, 2006, 2007, 2013-14.

Cooperative Faculty Research Grant, Consortium of Universities in the Washington Metropolitan Area (1993-1994).

Cooperative Research Grant, Council on Library Resources (1993-1994).

Journal of the American Society for Information Science Best Paper Award (1993).

ASIS/ISI Information Science Doctoral Dissertation Scholarship (1989).

HEA Title IIB Fellowship (Dept. of Education) (1989)

Chinese-American Librarians Association Scholarship (1987).

Beta Phi Mu (1985).

Vilas Fellowship, University of Wisconsin-Madison. 1984

Publications

Hsieh-Yee, I. (2021). Can We Trust Social Media? *Internet Reference Services Quarterly*, 1-15.
Published online: 28 July, 2021. 10.1080/10875301.2021.1947433.

Bailey, T., & Hsieh-Yee, I. (2020). Combating the Sharing of False Information: History, Framework & Literacy Strategies, *Internet Reference Services Quarterly*, 24 (1-2), 9-30.

Choi, Y., and Hsieh-Yee, I. (2010). Finding Images in an OPAC: Analysis of User Queries, Subject Headings, and Description Notes. *Canadian Journal of Information and Library Science*, 34(3): 271 – 295.

Hsieh-Yee, I. (2008). Educating Cataloging Professionals in a Changing Information Environment. *Journal of Education for Library and Information Science*, 46(2): 93-106.

Vellucci, S. L., Hsieh-Yee, I., and Moen, W.E. (2007). The Metadata Education and Research Information Commons (MERIC): A Collaborative Teaching and Research Initiative. *Education for Information*, 25(3&4): 169-178.

NISO Framework for Guidelines for Building Good Digital Collections. 3rd ed. Baltimore, MD: National Information Standards Organization, 2007. Also available online: <http://www.niso.org/framework/framework3.pdf> (NISO Working Group members: Priscilla Caplan (chair), Grace Agnew, Murtha Baca, Tony Gill, Carl Fleischhauer, Ingrid Hsieh-Yee, Jill Koelling, and Christie Stephenson.)

Choi, Y., Hsieh-Yee, I., and Kules, B. (2007). Retrieval Effectiveness of TOC and LCSH. *Proceedings of the Joint Conference on Digital Libraries*, pp. 233-234.

Vellucci, S., and Hsieh-Yee, I. (2007). They Didn't Teach Me That in Library School! Building

a Digital Teaching Commons to Enhance Metadata Teaching, Learning and Research. *Proceedings of the National Conference of the Association of College and Research Libraries, Baltimore, MD*, pp. 26-31.

Mitchell, Vanessa, and Ingrid Hsieh-Yee. (2007). Converting Ulrich's Subject Headings to FAST Headings: A Feasibility Study. *Cataloging & Classification Quarterly*, 45(1): 59-85.

Hsieh-Yee, I., Tang, R., and Zhang, S. (2007). User Perceptions of a Federated Search System. *IEEE Technical Committee on Digital Libraries Bulletin*, Summer 3(2) (URL = <http://www.ieee-tcdl.org>).

Tang, R., Hsieh-Yee, I., and Zhang, S. (2007). User Perceptions of MetaLib Combined Search: An Investigation of How Users Make Sense of Federated Searching." *Internet Reference Services Quarterly*, 12(12): 211-236.

Hsieh-Yee, I., Tang, R., and Zhang, S. (2006). User Perceptions of a Federated Search System. *Proceedings of the Joint Conference on Digital Libraries, June 11-15, 2006, Chapel Hill*, p. 338.

Hsieh-Yee, I. (2006). *Organizing Audiovisual and Electronic Resources for Access: A Cataloging Guide*. 2nd ed. Westport, Conn.: Libraries Unlimited.

NISO A Framework of Guidance for Building Good Digital Collections. 2nd ed. Bethesda, MD: National Information Standards Organization, 2004. Framework Advisory Group: Grace Agnew, Liz Bishoff, Priscilla Caplan (Chair), Rebecca Gunther and Ingrid Hsieh-Yee.

Hsieh-Yee, I. (2004). Cataloging and Metadata Education in North American LIS Programs. *Library Resources & Technical Services*, 48(1): 59-68.

Hsieh-Yee, I. (2004). Cataloging and Metadata Education. In Gary E. Gorman (Ed.), *International Yearbook of Library and Information Management 2003: Metadata Applications and Management*, (pp.204-234). London: Facet Publishing.

Yee, P. L., Hsieh-Yee, I., Pierce, G.R., Grome, R., and Schantz, L. (2004). Self-Evaluative Intrusive Thoughts Impede Successful Searching on the Internet. *Computers in Human Behavior*, 20(1): 85-101.

Hsieh-Yee, I. (2003). Cataloging and Metadata Education: A Proposal for Preparing Cataloging Professionals of the 21st Century. A report submitted to the ALCTS-Education Task Force in response to Action Item 5.1 of the *Bibliographic control of Web Resources: A Library of Congress Action Plan*. Approved by the Association for Library Collections and Technical Services. Web version available since April 2003 at <http://lcweb.loc.gov/catdir/bibcontrol/CatalogingandMetadataEducation.pdf>.

- Hsieh-Yee, I. (2002). Cataloging and Metadata Education: Asserting a Central Role in Information Organization. *Cataloging & Classification Quarterly* 34(½): 203-222.
- Hsieh-Yee, I., and Smith, M. (2001). The CORC Experience: Survey of Founding Libraries, Part I. *OCLC Systems & Services*, 17: 133-140. (Received "The Most Outstanding Paper of OCLC Systems & Services in 2001" award.)
- Hsieh-Yee, I., and Smith, M. (2001). The CORC Experience: Survey of Founding Libraries, Part II, Automated Tools and Usage. *OCLC Systems & Services*, 17: 166-177. (Received "The Most Outstanding Paper of OCLC Systems & Services in 2001" award.)
- Hsieh-Yee, I. (2001). ERIC User Services: Changes and Evaluation for the Future. *Government Information Quarterly*, 18: 31-42.
- Hsieh-Yee, Ingrid. (2001). Research on Web Search Behavior. *Library and Information Science Research*, 23: 167-185.
- Logan, E., and Hsieh-Yee, I. (2001). Library and Information Science Education in the Nineties. *Annual Review of Information Science and Technology*, 35: 425-477.
- Hsieh-Yee, I. (Ed.) (2001). *Library and Information Science Research*, 23 (2). A special issue in honor of the retirement of Douglas L. Zweizig.
- Hsieh-Yee, I. (2000). *ERIC User Services: Evaluation in a Decentralized Environment*. Washington, D.C.: Dept. of Education.
- Hsieh-Yee, Ingrid. (2000). *Organizing Audiovisual and Electronic Resources for Access: A Cataloging Guide*. Littleton, CO: Libraries Unlimited.
- Hsieh-Yee, I. (2000). Organizing Internet Resources: Teaching Cataloging Standards and Beyond. *OCLC Systems & Services*, 16: 130-143. (Received "The Most Outstanding Paper of OCLC Systems & Services in 2000" award.)
- Hsieh-Yee, I. (1998). The Retrieval Power of Selected Search Engines: How Well Do They Address General Reference Questions and Subject Questions? *Reference Librarian*, 60: 27-47.
- Hsieh-Yee, I. (1998). Search Tactics of Web Users in Searching for Texts, Graphics, Known Items and Subjects: A Search Simulation Study. *Reference Librarian*, 60: 61-85. (Received the 1997 Best ALISE Research Paper Award.)
- Hsieh-Yee, I. (1997). Access to OCLC and Internet Resources: LIS Educators' Views and Teaching Practices. *RQ*, 36: 569-86.

- Hsieh-Yee, I. (1997). Teaching Online and CD-ROM Resources: LIS Educators' Views and Practices. *Journal of Education for Library and Information Science*, 38: 14-34.
- Hsieh-Yee, I. (1996). The Cataloging Practices of Special Libraries and Their Relationship with OCLC. *Special Libraries*, 87: 10-20.
- Hsieh-Yee, I. (1996). Modifying Cataloging Practice and OCLC Infrastructure for Effective Organization of Internet Resources. In *Proceedings of the OCLC Internet Cataloging Colloquium*. [Online]. Available: <http://www.oclc.org/oclc/man/colloq/hsieh.htm>
- Hsieh-Yee, I. (1996). Student Use of Online Catalogs and Other Information Channels. *College & Research Libraries*, 57: 161-175.
- Hsieh-Yee, I. (1995). Ten entries in James S. C. Hu (Ed.), *Encyclopedia of Library & Information Science*, 913, 1028-29, 1036, 1037, 1145-46, 1514, 1575, 1763-64, 2216-27, 2378-79. Taipei, Taiwan: Sino-American Publishing. (Topics include "Advanced Technology/Libraries," "Information Ethics," "Instruction on Cataloging and Classification," "Instruction on Reference Services.")
- Hsieh-Yee, I. (1993). Effects of Search Experience and Subject Knowledge on Online Search Behavior: Measuring the Search Tactics of Novice and Experienced Searchers. *Journal of the American Society for Information Science*, 44: 161-174. (Received the 1993 Best JASIS Paper Award.)

Works in Progress

- Hsieh-Yee, I. Actionable lessons for debunking false information: A comparative analysis of literacy lessons in communication and library and information science
- Hsieh-Yee, I. Representing diverse voices and communities in cataloging and indexing: The evolution of controlled vocabularies
- Hsieh-Yee, I. Information Organization and Retrieval in Practice (a monograph scheduled for completion and publication in 2023)

Presentations

- Bailey, T. & Hsieh-Yee, I. (February 2020). The Phenomena of Sharing Misinformation and the Need for Information Literacy. Presented at the 2020 Bridging the Spectrum Symposium, Washington, D.C., February 14, 2020.

- Hsieh-Yee, I. and Fragan-Fly, J. (May 2018) *Trends, Design & Strategies for Digital Scholarship Services*. Presented at the 2018 Maryland/Delaware Library Association Conference, Cambridge, MD.
- Hsieh-Yee, I. (February, 2018) *Research Data Management: What It Takes to Succeed*. Presented at the 10th Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee, I. (February, 2017) *Research Data Management: New Competencies and Opportunities for Information Professionals*. Presented at the 9th Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee, I. and Lawton, P. (February, 2017) *Enhancing Catholic Portal Searches with User Terms and LCSH*. Presented at the 9th Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee, I. (2016, October) *Visualizing Data for Information*. Presented at the 2016 Virginia Library Association Conference, Hot Springs, VA.
- Hsieh-Yee, I. (2016, August) *Religious Materials Toolbox for Archivists: Solutions to Problems Facing the Profession*. Presented at Archives * Records 2016, Atlanta, GA.
- Hsieh-Yee, I. and Lawton, P. (2016, March) *Enhancing Retrieval of Catholic Materials with LCSH Knowledge Structure*. Presented at the 2016 Catholic Library Association Conference, San Diego, CA.
- Fagan-Fry, J. and Hsieh-Yee, I. (2016, February) *Approaches to Digital Scholarship at Top Universities around the World: Scholarly Publishing in the Digital Age*. Presented at the 8th Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee and Fagan-Fry, J. (2016, January) *Innovative Services for Digital Scholarship at Top 100 Research Libraries of the World*. Poster presented at the 2016 Annual Conference of the Association for Library and Information Science Education, Boston, Mass.
- Hsieh-Yee, I. and Lawton, P. (2015, June). *Crowdsourcing terms for CRRA portal themes*. Poster presented at the third CRRA symposium and annual meeting, Bringing the created toward the Creator: Liturgical art and design since Vatican II. Catholic Theological Union, Chicago, Illinois.
- Hsieh-Yee, I. and Lawton, P. (2015, February). *Crowdsourcing terms for thematic exploration in the Catholic Portal*. Poster presented at the 7th Annual Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee, I., James, R., and Fagan-Fry, J. (2015, February). *Support for digital scholarship at top university libraries of the world*. Poster presented at the 7th Annual Bridging the Spectrum Symposium, Washington, D.C.

- Hsieh-Yee, I., Zhang, S., Lin, K., and Cherry, S. (2015, February). *Thus said the end users: Summon experience and support for research workflows*. Poster presented at the 7th Annual Bridging the Spectrum Symposium, Washington, D.C.
- Yontz, E., Hsieh-Yee, I., & Houston, S. (2015, February). *Healthy Heroes Summer Reading Club: Developing healthy youth at public libraries*. 11th Annual Jean Mills Health Symposium, Greenville, North Carolina.
- Yontz, E., Hsieh-Yee, I., and Houston, S. (2015, January). *Healthy youth and libraries: A pilot study*. Association for Library & Information Science Education (ALISE) Annual Conference, Chicago, Illinois.
- Hsieh-Yee, I. (2014, May). *Linking CRRA resources to portal themes via authority files*. Presented at the Catholic Research Resources Alliance 2014 Membership Meeting, Marquette, WI.
- Hsieh-Yee, I. (2014, April). *Enhancing subject access to CRRA resources*. Presented at the 2014 Catholic Library Association Conference, Pittsburgh, PA.
- Hsieh-Yee, I. (2014, January). *Health Information Technology Program: Educational entrepreneurship in action*. Presented at the 2014 annual Conference of the Association for Library and Information Science Education, Philadelphia, PA.
- Hsieh-Yee, I., Zhang, S., Lin, K., and Cherry, S. (2014, January). *Discovering information through Summon: An analysis of user search strategies and search success*. Paper presented at the 6th Bridging the Spectrum Symposium, Washington, D.C.
- Hsieh-Yee, I. (2012, December). *National Digital Stewardship Alliance and SLIS at CUA: An Educational Partnership*. Paper presented at Best Practices Exchange: Acquiring, Preserving, and Providing Access to Government Information in the Digital Era, Annapolis, MD.
- Choi, Y. and Hsieh-Yee, I. (2010, January). *Finding Images in an OPAC: Analysis of User Queries, Subject Headings, and Description Notes*. Paper presented at 2nd Annual Bridging the Spectrum Symposium, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I. and Coogan, J. (2010, January). *Google Scholar vs. Academic Search Premier: What Libraries and Searchers Need to Know*. Paper presented at 2nd Annual Bridging the Spectrum Symposium, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I. (2009, November). *Information Science Education: An LIS School's Perspective*. Paper presented at Annual Meeting of the American Society for Information Science and Technology, Vancouver, British Columbia, Canada.

- Hsieh-Yee, I., Menard, E., Ya-Ning Chen, A., Shu-Jiun Chen, S., Kalfatovic, M. R., Wisser, K. M. (2009, November). *Information Organization in Libraries, Archives and Museums: Converging Practices and Collaboration Opportunities*. Presented at Annual Meeting of the American Society for Information Science and Technology, Vancouver, British Columbia, Canada. (Organizer and moderator of this panel.)
- Hsieh-Yee, I. and Coogan, J. (2009, July). *Catching up to Google Scholar: The Retrieval Power of Academic Search Premier and Google Scholar*. Poster presented at American Library Association Conference, Chicago, Illinois.
- Hsieh-Yee, I., with the CUA Scholarly Communications Project Team. (2009, January). *Digital Scholarship@CUA: Developing an Institutional Repository for CUA*. Poster presented at 1st Annual Bridging the Spectrum Symposium, Catholic University of America, Washington, D.C.
- Wise, M., Cylke, K., and Hsieh-Yee, I. (2009, January). *Digital Talking Books: Meeting the Needs of the Blind and the Handicapped*. Paper presented at the Bridging the Spectrum Symposium, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I. (2009, January). *User Expectations of MERIC*. Presented at the Information Organization Competencies for the 21st Century Discussion Session of the 2009 Conference of the Association for Library and Information Science Education, Denver, Colorado.
- Choi, Y., and Hsieh-Yee, I. (2008, November). *Subject Access for Images in an OPAC*. Annual Meeting of the American Society for Information Science and Technology, Columbus, Ohio. (Also co-organized a panel on Retrieving and Using Visual Resources: Challenges and Opportunities for Research and Education.)
- Hsieh-Yee, I. (2008, June). *Educating Cataloging Professionals in a Changing Information Environment*. National Taiwan University, Taipei, Taiwan.
- Vellucci, S. L., Moen, W.E., Hsieh-Yee, I., Marson, B., and Wisser, K. (2008, January) *Building a Metadata Education and Research Community through MERIC (Metadata Education and Research Information Commons): Demo and Stakeholder Input*. A panel presented at the 2008 Conference of the Association for Library and Information Science Education, Philadelphia, Pennsylvania.
- Hsieh-Yee, I., Choi, Y. and Kules, B. (2007, October). *Searching for Books and Images in OPAC: Effects of LCSH, TOC and Subject Domains*. A poster presented at the American Society for Information Science and Technology Annual Meeting, Milwaukee, Wisconsin.
- Hsieh-Yee, I. and Coogan, J. (2007, August) *Google Scholar vs. Academic Search Premier: A Comparative Analysis*. Presented to the Faculty and Staff of the University of the District

of Columbia.

- Hsieh-Yee, I. and Coogan, J. (2007, June). *Google Scholar vs. Academic Search Premier: A Comparative Analysis*. Presented to the Washington Research Library Consortium Community, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I., Choi, Y., and Kules, B.. (2007, June). *What Users Need for Subject Access: Table of Contents or Subject Headings?* A poster presented at the 2007 American Library Association Annual Conference, Washington, D.C., June 2007.
- Choi, Y., Hsieh-Yee, I., and Kules, B. (2007, June). *Retrieval Effectiveness of TOC and LCSH*. A paper presented at the Joint Conference on Digital Libraries 2007, Vancouver, Canada.
- Vellucci, S. L., Hsieh-Yee, I., and Moen, W.E. (2007, May). *If We Build It, Will They Come? Building a Community of Practice for Metadata Stakeholders*. A poster presented at the Rutgers University Research Day, Bridgeton, New Jersey.
- Hsieh-Yee, I. (2007, May). *Federated Searching: User Experience & Perceptions*. International Conference on Information Organization & Retrieval, National Taiwan University, Taipei, Taiwan.
- Hsieh-Yee, I. (2007, May). *Search Performance of Google Scholar and Academic Search Premier*. International Conference on Information Organization & Retrieval, National Taiwan University, Taipei, Taiwan.
- Hsieh-Yee, I. (2007, May) *MERIC: Building a Digital Commons for Metadata Education & Research*. International Conference on Information Organization & Retrieval, National Taiwan University, Taipei, Taiwan.
- Hsieh-Yee, I., and Coogan, J. (2007, March/April). *A Comparative Analysis of Google Scholar and Academic Search Premier*. Poster presented at the Association of College & Research Libraries 13th National Conference, Baltimore, Maryland.
- Vellucci, S. L. and Hsieh-Yee, I. (2007, March/April) *They Didn't Teach Me That in Library School! Building a Digital Teaching Commons to Enhance Metadata Teaching, Learning and Research*. On-site presentation and Webcast by Elluminate. A contributed paper presented at the Association of College & Research Libraries 13th National Conference, Baltimore, Maryland. The acceptance rate for contributed paper was 20%. This paper was one of 10 conference papers chosen for live webcast during the conference.
- Moen, W., Hsieh-Yee, I. and Vellucci, S.L. (2007, January) *A DSpace Foundation for a Teaching & Research Commons: The Metadata Education and Research Information Commons*. A poster session presented at the Open Repositories Conference 2007, San Antonio, Texas.

- Tang, R., Hsieh-Yee, I., and Zhang, S. (2006, November) *User Perception of MetaLib Combined Search*. Paper presented at the Annual Meeting of the American Society for Information Science and Technology, Austin, Texas, Nov. 2006.
- Hsieh-Yee, I. (2006, November). *Federated Searching: User Perceptions, System Design, and Library Instructions*. Paper presented at the Annual Meeting of the American Society for Information Science and Technology, Austin, Texas. (Panel organizer, moderator, presenter).
- Hsieh-Yee, I. (2006, November). *Building a Digital Teaching Commons to Enhance Teaching and Learning: The MERIC Experience and Challenges*. Paper presented at the Annual Meeting of the American Society for Information Science and Technology, Austin, Texas. (Panel organizer, moderator, presenter)
- Hsieh-Yee, I. (2006, September). *Search Performance of Google Scholar and Academic Search Premier*. Paper presented at the ERIC Publishers Meeting, Washington, D.C.
- Hsieh-Yee, I., Zhang, S., and Rong Tang, R. (2006, June). *User Perceptions of a Federated Search System*. Poster presented at Joint Conference on Digital Libraries, Chapel Hill, North Carolina.
- Hsieh-Yee, I. and Zhang, S. (2006, June). *Preparing Users for Federated Search: Implications of a MetaLib User Perceptions Study*. Paper presented at the 2006 Ex Libris User Groups of North America Conference, Knoxville, Tennessee.
- Hsieh-Yee, I. (2006, January). *MERIC Organizations and Navigation*. Paper presented at the 2006 ALISE Annual Conference, San Antonio, Texas.
- Hsieh-Yee, I. (2006, January). *Metadata and Cataloging Education: Recommended Competencies*. Paper presented at the 2006 ALISE Annual Conference, San Antonio, Texas.
- Hsieh-Yee, I. (2005, November). *Digital Library Evaluation: Progress & Next Steps*. Presentation at the Annual Meeting of the American Society for Information Science & Technology, Charlotte, North Carolina.
- Hsieh-Yee, I. (2005, August). *Providing Access to Digital Content: Issues for DL Managers*. Presentation at MDK12 Digital Library Steering Committee Meeting, Columbia, Maryland.
- Hsieh-Yee, I. (2005, April). *Enhancing Teaching and Learning: The Role of School Library Media Specialists*. Presentation at Meeting of the Baltimore County Public School System School Media Specialists, Baltimore, Maryland.

- Hsieh-Yee, I. (2005, January). *Subject Access and Users: Insights & Inspirations from Marcia J. Bates*. Paper presented at the Historical Perspectives SIG, 2005 Conference of the Association for Library and Information Science Education, Boston, Massachusetts.
- Hsieh-Yee, I. (2005, January). *Electronic Resource Management: Practice, Employer Expectations, & CE Interests*. Paper presented at Technical Services Education SIG, 2005 Conference of the Association for Library and Information Science Education, Boston, Massachusetts.
- Hsieh-Yee, I. (2004, October). *Library Professionals for the Digital Age: Competencies & Preparation*. Paper presented at Bibliographic Access Management Team meeting, Library of Congress, Washington, D.C.
- Hsieh-Yee, I. (2004, January). *Cataloging and metadata expertise for the digital era*. Presented at Preparing 21st Century Cataloging and Metadata Professionals: A Workshop for Educators and Trainers, San Diego and sponsored by ALCTS, ALISE, LC, and OCLC.
- Hsieh-Yee, I. (2004, January). *Educating catalogers for the digital era*. Paper presented at the Technical Services SIG, 2004 Conference of the Association for Library and Information Science Education, San Diego.
- Hsieh-Yee, I. (2003, July). *Cataloging Education for the 21st Century*. A presentation at the Library of Congress, Washington, D.C.
- Hsieh-Yee, I. (2002, January) *Metadata Education and Research Priorities: A Delphi Study of Metadata Experts*. Presentation at the 2002 Conference of the Association for Library and Information Science Education, New Orleans.
- Hsieh-Yee, I. (2001, November). *A Delphi Study of Metadata: Preliminary Findings*. Poster session at the 2001 Annual Meeting of the American Society for Information Science & Technology, Washington, D.C.
- Hsieh-Yee, I. (2001, June). *Resources on Asian American Children: Analysis of Retrieval by Search Engines and WorldCat*. Presentation at the National Conference on Asian Pacific American Librarians, San Francisco.
- Hsieh-Yee, I. (2001, January). *Delphi Study on Metadata: Project Design*. Presentation at Research Awards Session, Association for Library & Information Science Education, Washington, D.C.
- Hsieh-Yee, I. (2000, May). *Web Search Behavior Research: Progress and Implications*. Presentation at the Symposium on Evaluating Library and Information Science Research, University of Wisconsin-Madison, Madison, Wisconsin.

- Hsieh-Yee, I. (2000, March). *ERIC User Services: Evaluation in a Decentralized Environment*. Presentation at the National ERIC Joint Directors/Technical Meeting, Arlington, Virginia.
- Hsieh-Yee, I. (2000, January). *Enhancing Learning with Web Technology*. Presentation at Faculty Conversations, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I. (2000, January). *From Surrogates to Objects: CUA's Approaches to Organizing Electronic Resources*. Paper presentation at the Annual Conference of the Association for Library and Information Science Education, San Antonio, Texas.
- Yee, P., and Hsieh-Yee, I. (1997, November). *Individual Differences in Search Behavior on the WWW*. A poster session presented at the 38th Annual Meeting of the Psychonomic Society, Philadelphia, Pennsylvania.
- Hsieh-Yee, I. (1997, April). *Research + Marketing + Preparation = Job!* Presented at the "Workshop on Resume and Interview Techniques," Special Libraries Association, Student Chapter, Catholic University of America, Washington, D.C.
- Hsieh-Yee, I. (1997, February). *Creating CyberCatalogers: Education and Training*. Presentation at ALA's Midwinter Meeting, Washington, D.C.
- Hsieh-Yee, I. (1997, February). *Search Tactics of Web Users in Searching for Texts, Graphics, Known Items and Subjects: A Search Simulation Study*. Presented at the Conference of the Association for Library and Information Science Education, Washington, D.C.
- Hsieh-Yee, Ingrid. "Beginning Your Special Library/Information Center Career." Presented at SLA's "Career Day," Jan. 11, 1997, Catholic University of America.
- Hsieh-Yee, I. (1996, September). *The Roles of Library and Information Scientists in Managing Electronic Information*. Presentation at Hamilton College, Clinton, New York.
- Hsieh-Yee, I. (1996, May). *The Future of Cataloging as a Profession*. Presented at "The Cataloging Forum, Library of Congress, Washington, D.C.
- Hsieh-Yee, I. (1994, October). *The Impact of the Internet on OPACs*. Presented at the Third Workshop on User Interfaces for OPACs, Library of Congress, Washington, D.C.

Reports

- Hsieh-Yee, I., with Knowledge Management Competencies and Performance Action Group of the Federal Knowledge Management Initiative. "From Knowledge Management Competencies to Improved Organizational Performance." April 9, 2009.

- Hsieh-Yee, I., with Knowledge Practices Action Group of the Federal Knowledge Management Initiative. "KM Practice in Government Agencies: Findings and Recommendations." April 9, 2009.
- Hsieh-Yee, I. "Delphi Study on Metadata." 2001. Three quarterly reports submitted to the Association for Library and Information Science Education.
- Hsieh-Yee, I. "College Students' Information Channels: Patterns of Use and Possible Factors in Channel Selection." 1995. Submitted to the Catholic University of America.
- Hsieh-Yee, I. "The Information-Seeking Patterns of Scholars and Their Use of an Online Information System." 1994. Submitted to the Council on Library Resources.

Book Reviews

- Review of *The Measurement and Evaluation of Library Services*, by Sharon L. Baker and F. Wilfrid Lancaster. *Information Processing and Management* 30 (1994): 450-52.
- Review of *Subject Access to Films and Videos*, by Sheila S. Intner and William E. Studwell; and *Cataloging Unpublished Nonprint Materials*, by Verna Urbanski with Bao Chu Chang and Bernard L. Karon. *Information Processing and Management* 30 (1994): 449-50.
- Review of *Automated Information Retrieval in Libraries: A Management Handbook*, by Vicki Anders. *Journal of Library and Information Science* 19 (1993): 98-100.
- Review of *Full Text Databases*, by Carol Tenopir and Jung Soon Ro. *Information Processing and Management* 28 (1992): 667-68.
- Review of *Descriptive Cataloging for the AACR2R And USMARC: A How-to-Do It Workbook*, by Larry Millsap and Terry Ellen Ferl. *Information Processing and Management* 28 (1992): 809-11.
- Review of *MARC Manual: Understanding and Using MARC Records*, by Deborah J. Byrne. *Information Processing and Management* 28 (1992): 537-38.

Service

Professional Associations and Societies

- Library of Congress. RDA Training Program for the Profession. Co-authored with Tim Carlton. 2013-2014.
- 2014 Digital Preservation Outreach & Education Survey. Contributed to the design of the survey, 2014.
- National Digital Stewardship Alliance. Outreach Committee. 2011-2014.
- National Digital Stewardship Residency Program. Advisory Group, 2012-2013.
- FEDLINK Health Information Technology Advisory Council, 2011-2015.

- 2012 Joint Conference on Digital Libraries. Program Planning Committee, Pre-Conference Proposals Review Committee, 2012
- Catholic Research Resources Alliance. Five-Year Strategic Plan Task Force, 2011-2012
- Institute of Museum and Library Services. Grant reviewer. 2004, 2005, 2010.
- Association for Library and Information Science Education.
 - * ALISE Bodan Wynar Research Paper Award Committee, 2015, 2016, 2017
 - * ALISE Eugene Garfield Dissertation Award Competition, Jury, 2013, 2014
 - * ALISE Research Grant Competition Committee. Chair, 2012
 - * Pratt-Severn Faculty Innovation Award. Chair, 2009, 2010
 - * ALISE Doctoral Poster Jury, 2012
 - * “Information Organization Competencies for the 21st Century” Discussion session leader. 2009 Conference of the Association for Library and Information Science Education.
 - * Assisted Technical Services SIG Convener in organizing a program, ““Building a Metadata Education and Research Community through MERIC (Metadata Education and Research Information Commons): Demo and Stakeholder Input” for the 2008 ALISE conference.
 - * Association for Library Collections and Technical Services/Association for Library and Information Science Education (ALCTS/ALISE) Metadata Education and Research Information Center (MERIC) Advisory Board, Co-Chair (with Sherry Vellucci), 2005-2007. Chair, 2008-2009 (leading the effort to build MERIC, a repository and collaborative space for metadata educators, practitioners, and researchers)
 - * Technical Services SIG, Convener, 2004-2005. Organized a program on “Electronic Resources Management: Current Practices, Employer Expectations, and Teaching Strategies” for the 2005 conference in Boston, Massachusetts.
 - * Technical Services SIG, Convener, 2003-2004. Organized a program on “Organizing Information with Metadata: Desired Competencies and Teaching Innovations” for the 2004 conference.
 - * Technical Services SIG, Convener, 1999-2000. Organized a program on "Teaching the Organization of Electronic Resources" for the 2000 conference.
 - * Curriculum SIG, Co-convener (with Sibyl Moses), 1996-97. Organized a program on “Government Information Policy” for the 1997 conference.
- American Society for Information Science & Technology.
 - * Reviewer, Conference program panel submissions and poster submissions, 2005, 2006, 2007, 2009, 2011, 2012, 2013, 2014, 2015, 2016, 2017
 - * Nomination Committee, 2009-2011
 - * Information Science Education Special Interest Group. American Society for Information Science and Technology. Chair-Elect, 2007-2008. Chair 2008-2009.
 - * Committee on Information Science Education. 1999-2006.
 - * Committee on Information Science Education. Organizing Committee for an orientation program for students at ASIS annual meetings, 1999-2001
 - * Committee on Information Science Education. Sub-committee on Student Welfare (focusing on issues related to master's education), 1998-2001

- * SIG ED. Organizing Committee for the "Seminar on Research and Career Development" for junior researchers. 1995-96 (chair), 1997-2001
- * ISI Doctoral Dissertation Proposal Scholarship Jury, 1997; 2001, 2002
- * Pratt-Severn Best Student Research Paper Award Jury. Chair. 1997
- * 1998 Midyear Meeting (referee of contributed papers), 1997
- * Organizer and moderator of the ASIS Doctoral Forum and the Doctoral Research Seminar 1994-1995
- * SIG Human Computer Interaction. Chair-Elect, Chair, 1993-1995
- * Doctoral Forum Award Jury, 1995
- * Best Student Paper Award Jury, 1995
- American Library Association.
 - * Committee on Accreditation, External Review Panelist, 2009- (site visiting team 2013-2014; site visiting team 2016-2017)
 - * Association for Library Collections and Technical Services Task Force on Competencies and Education for a Career in Cataloging, member, 2008-2009
 - * Facilitator for "What They Don't Teach in Library School: Competencies, Education and Employer Expectations for a Career in Cataloging," an Association for Library Collections and Technical Services Preconference, June 22, 2007 in Washington, D.C. Also a local liaison for bringing this program to the Catholic University of America.
 - * Facilitator for a discussion on "Effect of Electronic Resources on Technical Services" at ALA's Midwinter Meeting held in Feb. 1997 in Washington, D.C.
 - * International Relations Committee, Subcommittee Task Force for IFLA and China, 1994-1997
- Virginia Association of School Librarians. Scholarships and Awards Committee. 2010-2012
- Federal Knowledge Management Initiative, Knowledge Management Practices Action Group. Member. 2009 (leading the effort to build a knowledge management repository)
- Federal Knowledge Management Initiative, Knowledge Management Competencies & Learning Action Group. Member. 2009 (developing an action plan for helping government knowledge workers and government agencies to develop knowledge management competencies)
- National Center for Education Statistics. Technical Review Panel. 2008.
- External evaluator for a case of promotion to full professorship. University of Tennessee. 2008.
- National Information Standards Organization (NISO). Advisory Board, Revision of "IMLS Framework of Guidance for Building Good Digital Collections," 2004, 2007.
- Library of Congress, Bibliographic Control of Web Resources: A Library of Congress Action Plan. Principal Investigator of Action Item 5.1, focusing on cataloging and metadata education for students and new librarians, 2002-2003. (worked with the Association for Library Collections and Technical Services, Education Task Force)

- Chinese American Librarians Association
 - * Chinese American Librarians Association Outstanding Library Leadership Award in Memory of Dr. Margaret Chang Fung, Award Committee, 2016-2017
 - * Achievement Award Jury, 2000-2001
 - * CALA Goal 2000 Task Force, 1997
 - * Scholarship Committee, 1995, 1996-1997 (chair)
 - * Board of Directors, 1994-1997
 - * Publication Committee, 1993-1995
 - * International Relations Committee, 1993-1996
- SailorSM Assessment Advisory Group (An impact study of Sailor, Maryland's Public Information Network), 1995
- Editorial boards
 - Journal of Library & Information Science. Editorial Board, 2012-
 - Chinese American Librarians Association, *Occasional Papers Series*. Editorial Board, 2009-2016.
 - Library Quarterly*. Editorial Board, 2003-2008
 - Bulletin of the Medical Library Association*, 1994-97
 - Newsletter editor for the Chinese American Librarians Association, 1989-92
- Referee for the following journals
 - Information Processing and Management*
 - Journal of Digital Information*
 - Journal of Education for Library and Information Science*
 - Journal of Information Science*
 - Journal of Library & Information Science*
 - Journal of Library Metadata*
 - Journal of the American Society for Information Science & Technology*
 - Library and Information Science Research*
 - Library Quarterly*
- Expert reviewer, “Digital Library” course, Evaluation module, University of North Carolina, Chapel Hill, 2007-2008.
- Expert reviewer, “Information Organization” course, University of Michigan, Ann Arbor. 2007.

Catholic University of America

- Academic Senate Committee on Appointments and Promotions, 2005-2008
- Academic Senate Committee on Committees and Rules, 2009-2012, 2019-2022
- Academic Senate representative (for School of Arts & Sciences), 2017-2020

- Academic Senate representative (for School of Library & Information Science) 2003-2012
- Academic Senate Library Committee, Interim Chair (2007), Member, 2008-2012
- Academic Senate Committee on Computing, 1995-2003
- President's Administrative Council, 2010-2012
- Deans' Council, 2010-2012
- Academic Leadership Group, 2010-2012
- Graduate Board, 2010-2012
- School of Arts & Sciences, Committee on Appointments and Promotions, 2015-2019
- School of Arts & Sciences, Academic Council, 2015-2016.
- School of Arts & Sciences, Ordinary Professor Group, 2013-
- Doctoral Dissertation Defense Committee, Chair, Dept. of Psychology, 2016, 2017, 2018, 2019, 2020, 2021, 2022
- Doctoral Dissertation Defense Committee, Chair, Dept. of Education, 2014, 2015, 2017, 2018, 2019
- Doctoral Dissertation Defense Committee, Chair, National Catholic School of Social Services, 2019, 2020, 2021
- Doctoral Dissertation Defense Committee, Chair, School of Nursing, 2006, 2008
- CUA Scholarly Communication Project Team, Member (2007), Chair, 2008-2009
- Dean Search Committee, 1992-1994, 1998-1999, 2002-2003, 2006-2007
- Fulbright Review Panel, 2006
- CUA Service Learning Advisory Board, 2001-2002
- CUA Faculty Conversations on Enhancing Teaching and Learning through Technology, Planning Group, 1999-2001
- CUA Initiative on Technology and Teaching, 1998-2001

Dept. of Library and Information Science

- Admissions Committee, 2007-2009, Chair 2010-2012, Member 2013-2015, Member 2018-2022
- Appointments and Promotions Committee, 1991-2022
- Student advisement, 1990-2022
- 4+1 Bachelor to Master program, Program Advisor, 2014-2022
- LIS Minor, Advisor, 2017-2022
- Curriculum Committee, 2020-21 (Member), 2021-2022 (Chair)
- Scholarship & Awards Committee, 2016-2018, 2020-21
- Comprehensive examination editor, 2020-21
- Symposium, Colloquium, Lecture Series Committee, May 2018-May 2020 (Chair)
- Symposium and Colloquium Committee, fall 2016-May 2018
- Bridging the Spectrum Symposium session moderator, 2018, 2019, 2020, 2021

- Organization of Information Core Course Outcomes Assessment, 2019, 2020
- The Information Professions in Society Core Course Outcomes Assessment, 2019, 2020
- Community Services Librarianship, Course of Study Review and Revision, 2020
- Organization of Information, Course of Study Review and Revision, 2019
- Organization of Information, Competency and Course Offerings Review, 2018-2019
- Accreditation presentation, Chair, June 2015-August 2016
- Interim Co-Chair, June 2015-August 2016.
- Blended/OWL Learning Committee, spring 2016-2018
- Technology Committee, fall 2016-2017
- Comprehensive examination editor, 2016-2017, reader (every year since 1990)
- LIS Advisory Board, 2015-2016 (chair); fall 2016- May 2018 (member)
- Committee on Planning and Assessment, 2015-2016 (chair)
- Senior Faculty Committee, 2014-2016.
- Accreditation Steering Committee, 2014-2016 (Chair, 2015-2016)
- Accreditation Students Standard Committee, co-chair, 2014-2016
- Accreditation Mission, Goals, and Objectives Standard Committee, co-chair, 2014-2016
- Accreditation Curriculum Standard, member 2014-2-16
- Accreditation Administration and Finance Standard, member 2014-2016
- Cultural Heritage Information Management Project (IMLS-funded), Co-PI, 2012-2015
- Cultural Heritage Information Management Forum (scheduled for June 2015), Co-Organizer, 2013-2015
- Health Information Technology Interim Review Committee, 2015 (chair)
- Health Sciences Librarianship Advisory Group, 2015- (chair)
- Comprehensive examination editors, 2013-2014, 2016-2017
- National Digital Stewardship Alliance liaison, 2011-2014
- Advisory Board, Chair 2010-2012
- Academic Honesty Committee, Chair, 2008-2012
- Blended Learning Committee, 2010-2012
- Colloquium Committee, 2010-2012
- Comprehensive Examination Administration, 2010-2012
- Cultural Heritage Information Management Advisory Committee, 2010-2012 (chair), 2013-
- Curriculum Committee, 1991-2003, 2007-2009, Chair 2010-2012, member 2013-
- Curriculum Subcommittee on Comprehensive Examination, Chair 2009-2012
- Health Information Technology Advisory Board, Chair 2010-2012. Member 2013-
- Health Sciences Advisory Committee, 2009, Chair 2010-2012. Member 2013-
- HIT Expert Forum, Chair 2012. Member 2013-2015
- Health Information Technology Student Group Advisor, 2011-2012
- State Council for Higher Education of Virginia, SLIS Representative, 2010-2012
- Symposium Planning Committee, 2010-2012
- Website Management Team, Chair, 2010-2012
- Urban School Librarianship Project (IMLS-Funded), PI, 2007-2011 (chair, 2010-11)

- Failing Grades Committee, 1995-1997 (chair), 2000-2001 (chair), 2004-2005 (chair), 2007 (chair)-2011
- Faculty Search Committee, 1994-1998, 2002-2004, 2006 (chair), Fall 2007-2009, Chair fall 2009-2012
- Recruitment Committee, Chair 2010-2012
- Strategic Planning Committee, Chair 2010-2012
- Technology Committee, 2010-2012
- Accreditation Advisory Committee, 2007-2009
- Accreditation Coordinating Committee, 2007-2009
- Accreditation Steering Committee, 2007-2009
- SLIS Advisory Group, 2007-2009
- Accreditation Curriculum Standard Committee, Co-chair, 2007-2009
- Accreditation Faculty Standard Committee, Co-chair, 2007-2009
- LSC 551 Information Organization Review Team, Co-chair, 2008-2009, 2015-2016.
- Curriculum Subcommittee on Portfolios, 2009
- LSC 555 Information Systems in Libraries and Information Centers Review Team, contributor, 2008-2009
- Redesign of LSC 730 Use and Users of Libraries and Information. 2009-
- Development of a metadata institute that was taught as LSC 715 Organization of Internet Resources in 2008. The institute is being revised and will be offered in 2010 under a new course title.
- Development of lesson plans, assignments, and evaluation rubrics for LSC 606, Cataloging and Classification, for the School's NCATE accreditation. 2008
- Howard and Mathilde Rovelstad Scholarship Committee, Chair, 2004-2007
- Assistant Dean Search Committee, Chair, Fall 2007
- Liaison to the Association for Library Collections and Technical Services to bring its preconference program, Cataloging Education and Employer Expectations, to CUA during the 2007 American Library Association Annual Meeting in Washington, D.C.
- Organizer of the colloquium presentation and reception for Tamar Sadeh of Ex Libris on PRIMO June 2007
- Practicum review and design (work with potential supervisors, such as the American Indian Museum internship description revision) 2006-
- Comprehensive examinations (edits, proctoring, and grading), 1990-
- SLIS Web site redesign: Comments and suggestions. Fall 2007
- Conducted surveys of current students and alumni in preparation for the 2005 re-accreditation, 2004-2005
- Technology Committee, 1992-1999 (chair, 1996-1998), 2002-2003 (member)
- Colloquia Committee 1997-1999, 2002-2003.
- Advisor of the CUA Student Chapter of the American Society for Information Science and Technology, 2002-2003
- Visiting Professor Search Committee, 1999, 2000, 2001
- Leader, Participation in the CORC experiment, 1999-2000
- Advisor of the Special Libraries Association Student Chapter, 1993-1999; the group was recognized for outstanding leadership by SLA in 1999.

- COA planning Committee, Task Force on Electronic Presentation of SLIS Reports (team leader) 1997-1998
- COA Planning Committee, Subcommittee on Technology 1996-1998
- NLM practicum coordinator, 1997-1998
- Computer Literacy Workshops: Assisted with the development and evaluation of the workshops, 1996-1998
- Leader, Participation in the InterCat project, 1995-1997

APPENDIX B

[Library of Congress](#) >> [MARC](#) >> [Understanding MARC](#)

MARC 21 Reference Materials

[Part VII: A Summary of Commonly Used MARC 21 Fields](#)

[Part VIII: A List of Other Fields Often Seen in MARC Records](#)

[Part IX: The Leader](#)

[Part X: Field 008 for Books](#)

Part VII:

A Summary of Commonly Used MARC 21 Fields

This is a summary of the MARC 21 tags used most frequently by libraries in entering their own bibliographic records. For full listings of all MARC 21 tags, indicators, and subfield codes, see *MARC 21 Format for Bibliographic Data*.

In the explanations on these pages:

Tags -- The tags (3-digit numbers) are followed by the names of the fields they represent. In this summary, and in the *MARC 21 Format for Bibliographic Data*, if a tag can appear more than once in one bibliographic record, it is labeled repeatable (R). If it can only be used once, it is labeled non-repeatable (NR). For example, a catalog record can have several subjects, so the tags for subject added entries (6XX) are labeled repeatable (R).

Indicators -- The use of indicators is explained in fields where they are used. Indicators are one-digit numbers. Beginning with the 010 field, in every field -- following the tag -- are two character positions, one for Indicator 1 and one for Indicator 2. The indicators are not actually defined in all fields, however. And it is possible that a 2nd indicator will be used, while the 1st indicator remains undefined (or vice versa). When an indicator is undefined, the character position will be represented by the character # (for blank space).

Subfield codes -- All the data in each field (beginning with the 010 field) is divided into subfields, each of which is preceded by a delimiter-subfield code combination. The most common subfield codes used with each tag are shown. Each subfield code is preceded by the character \$, signifying a delimiter. The name of the subfield follows the code.

In general, every field **MUST** have a subfield 'a' (**\$a**). One exception that is often seen is in Field 020 (ISBN), when the ISBN information (subfield **\$a**) is unavailable but the price (subfield **\$c**) is known. Some subfields are repeatable. In this summary, repeatability is noted for only the more common repeatable subfields.

Examples: Examples follow the explanation for each field. For clarity, one space has

been placed between the tag and the first indicator, one space has been placed between the second indicator and the first delimiter- subfield code, and one space has been inserted between the delimiter-subfield code and the subfield data.

010 Library of Congress Control Number -- (LCCN)
(NR, or Not Repeatable)

Indicators undefined.

Subfield used most often:

\$a -- Library of Congress control number

Example: 010 ## \$a ###86000988#

020 International Standard Book Number -- (ISBN)
(R, or Repeatable)

Indicators undefined.

Subfields used most often:

\$a -- International Standard Book Number

\$c -- Terms of availability (often a price)

\$z -- Cancelled/invalid ISBN (R)

Example: 020 ## \$a 0877547637

040 Cataloging source -- (NR)

Indicators undefined.

Subfields used most often:

\$a -- Original cataloging agency

\$c -- Transcribing agency

\$d -- Modifying agency (R)

Example: 040 ## \$a DLC
 \$c DLC
 \$d gwhs

100 Main entry -- Personal name -- (primary author)
(NR; there can be only one main entry)

Indicator 1: Type of personal name entry element

0 -- Forename

1 -- Surname (this is the most common form)

3 -- Family name

Indicator 2 undefined.

Indicator 2 became obsolete in 1990. Older records may display 0 or 1

*Subfields used most often:***\$a** -- Personal name**\$b** -- Numeration**\$c** -- Titles and other words associated with a name (R)**\$q** -- Fuller form of name**\$d** -- Dates associated with a name (generally, year of birth)

Example: 100 1# \$a Gregory, Ruth W.
 \$q (Ruth Wilhelme),
 \$d 1910-

130 Main entry -- Uniform title -- (NR)*Indicator 1: Nonfiling characters*

0-9 -- Number of nonfiling characters present (for initial articles, including spaces)

Indicator 2 undefined.

Indicator 2 became obsolete in 1990. (See 100 above.)

*Subfields used most often:***\$a** -- Uniform title**\$p** -- Name of part/section of a work (R)**\$l** -- Language of a work**\$s** -- Version**\$f** -- Date of a work

Example: 130 0# \$a Bible.
 \$p O.T.
 \$p Psalms.

240 Uniform title (NR)*Indicator 1: Uniform title printed or displayed*

0 -- Not printed or displayed

1 -- Printed or displayed (most common)

Indicator 2: Nonfiling characters

0-9 -- Number of nonfiling characters present (for initial articles, including spaces)

Subfields used most often:

\$a -- Uniform title

\$l -- Language of a work

\$f -- Date of a work

Example: 240 10 \$a Ile mystérieuse.
 \$l English.
 \$f 1978

245 Title Statement (NR)

Indicator 1: Title added entry

(Should the title be indexed as a title added entry?)

0 -- No title added entry

(indicates a title main entry; i.e. no author is given)

1 -- Title added entry

(the proper indicator when an author given in 1XX; the most common situation)

Indicator 2: Nonfiling characters

0-9 -- Number of nonfiling characters present, including spaces; usually set at zero, except when the title begins with an article; e.g., for *The robe*, the second indicator would be set to 4. The letters *T*, *h*, *e*, and the space following them are then ignored in alphabetizing titles. The record will be automatically filed under "*r*" -- for *Robe*.

Subfields used most often:

\$a -- Title proper

\$h -- Medium (often used for non-book media)

\$p -- Name of part/section of a work (R)

\$b -- Reminder of title (subtitles, etc.)

\$c -- Remainder of title page transcription/Statement of responsibility

Example: 245 14 \$a The DNA story :
 \$b a documentary history of gene
 cloning /
 \$c James D. Watson, John Tooze.

246 Varying form of title (R)

Indicator 1: Note/title added entry controller

1 -- Note, title added entry

3 -- No note, title added entry

Indicator 2: Type of title

-- No information provided

0 -- Portion of title

1 -- Parallel title

4 -- Cover title

8 -- Spine title

Subfield used most often:

\$a -- Title proper

Example: 246 3# \$a Four corners power review

250 Edition statement (NR)

Indicators undefined.

Subfield used most often:

\$a -- Edition statement

Example: 250 ## \$a 6th ed.

260 Publication, distribution, etc. (Imprint) (R)

Indicator 1: Sequence of publishing statements

-- No information provided

Indicator 2: Undefined

Subfields used most often:

\$a -- Place of publication, distribution, etc. (R)

\$b -- Name of publisher, distributor, etc. (R)

\$c -- Date of publication, distribution, etc. (R)

Example: 260 ## \$a New York :
 \$b Chelsea House,
 \$c 1986.

300 Physical description (R)

Indicators undefined.

Subfields used most often:

- \$a** -- Extent (number of pages) (R)
- \$b** -- Other physical details (usually illustration information)
- \$c** -- Dimensions (cm.) (R)
- \$e** -- Accompanying material (for example, "teacher's guide" or "manual")

Example: 300 ## \$a 139 p. :
 \$b ill. ;
 \$c 24 cm.

440 Series statement / Added entry--Title

This field was made obsolete in 2008 to simplify the series statement. See 490 and 830.

490 Series statement (No added entry is traced from field) (R)

Indicator 1: Specifies whether series is traced (whether an 8XX tag is also present)

- 0 -- Series not traced
- 1 -- Series traced (8XX is in record)

Indicator 2 undefined.

Subfield used most often:

- \$a** -- Series statement (R)
- \$v** -- Volume number (R)

Example: 490 1# \$a Colonial American craftsmen

500 General note (R)

Indicators undefined.

Subfield used most often:

- \$a** -- General note (Used when no specialized note field has been defined for the information. Examples: Notes regarding the index; the source of the title; variations in title; descriptions of the nature, form, or scope of the item.)

Example: 500 ## \$a Includes index.

504 Bibliography, etc. note (R)

\$a -- Bibliography, etc. note

505 Formatted contents note (R)

\$a -- Formatted contents note

520 Summary, etc. note (R)

\$b -- Expansion of summary note

Ex. 1039
CISCO SYSTEMS, INC. / Page 78 of 245

600 Subject added entry -- Personal name (R)*Indicator 1:* Type of personal name entry element

- 0 -- Forename
- 1 -- Surname (this is the most common form)
- 3 -- Family name

Indicator 2: Subject heading system/thesaurus (identifies the specific list or file which was used)

- 0 -- Library of Congress Subject Headings
- 1 -- LC subject headings for children's literature
- 2 -- Medical Subject Headings
- 3 -- National Agricultural Library subject authority file
- 4 -- Source not specified
- 5 -- Canadian Subject Headings
- 6 -- Répertoire de vedettes-matière
- 7 -- Source specified in subfield \$2

(Note regarding Sears subject headings: The MARC 21 format does not provide an assigned indicator for Sears subject headings. Therefore, an indicator of 7 is used, and the MARC defined code "sears" is placed in subfield \$2.)

Subfields used most often:

- \$a** -- Personal name (surname and forename)
- \$b** -- Numeration
- \$c** -- Titles and other words associated with a name (R)
- \$q** -- Fuller form of name
- \$d** -- Dates associated with a name (generally, year of birth)
- \$t** -- Title of a work
- \$v** -- Form subdivision (R)
- \$x** -- General subdivision (R)
- \$y** -- Chronological subdivision (R)
- \$z** -- Geographic subdivision (R)
- \$2** -- Source of heading or term (used with 2nd indicator of 7)

Example: 600 10 \$a Shakespeare, William,
 \$d 1564-1616
 \$x Comedies
 \$x Stage history.

Example: 600 10 \$a Shakespeare, William,
 \$d 1564-1616
 \$x Knowledge

\$z Rome \$v Congresses.

Notice that subfields \$v, \$x, and \$z in the 600 field are repeatable. Subfields \$v, \$x, \$y, and \$z do not have to be in alphabetical order. They will be in the order prescribed by the instructions given by the subject heading system.

610 Subject added entry -- Corporate name (R)

Indicator 1: Type of corporate name entry element

- 0 -- Inverted name (not used with AACR2)
- 1 -- Jurisdiction name
- 2 -- Name in direct order

Indicator 2: Subject heading system/thesaurus.

See indicator 2 under 600

Subfields used most often:

- \$a** -- Corporate name or jurisdiction name as entry element
- \$b** -- Subordinate unit (R)
- \$v** -- Form subdivision (R)
- \$x** -- General subdivision (R)
- \$y** -- Chronological subdivision (R)
- \$z** -- Geographic subdivision (R)
- \$2** -- Source of heading or term (used with 2nd indicator of 7)

<i>Example:</i> 610 10 \$a United States. \$b Army Air Forces \$v Biography.
--

650 Subject added entry -- Topical term (Most subject headings fit here.) (R)

Indicator 1: Level of subject

- # -- No information provided

Indicator 2: Subject heading system/thesaurus

(identifies the specific list or file which was used)

- 0 -- Library of Congress Subject Headings
- 1 -- LC subject headings for children's literature
- 2 -- Medical Subject Headings
- 3 -- National Agricultural Library subject authority file
- 4 -- Source not specified
- 5 -- Canadian Subject Headings
- 6 -- Répertoire de vedettes-matière
- 7 -- Source specified in subfield \$2

Note regarding Sears subject headings: The MARC 21 format does not provide an assigned indicator for Sears subject headings. Therefore, an indicator of 7 is used, and the MARC defined code "sears" is placed in subfield \$2.)

Subfields used most often:

- \$a -- Topical term
- \$v -- Form subdivision (R)
- \$x -- General subdivision (R)
- \$y -- Chronological subdivision (R)
- \$z -- Geographic subdivision (R)
- \$2 -- Source of heading or term used with 2nd indicator of 7)

Example: 650 #0 \$a Theater
 \$z United States
 \$v Biography
 \$v Dictionaries.

Notice that subfields \$v, \$x, and \$z in the 650 field are repeatable. Subfields \$v, \$x, \$y, and \$z do not have to be in alphabetical order. They will be in the order prescribed by the instructions given by the subject heading system.

651 Subject added entry -- Geographic name (R)

Indicator 1: undefined.

Indicator 2: Subject heading system/thesaurus.

See indicator 2 under 600

Subfields used most often:

- \$a -- Geographic name
- \$v -- Form subdivision (R)
- \$x -- General subdivision (R)
- \$y -- Chronological subdivision (R)
- \$z -- Geographic subdivision (R)
- \$2 -- Source of heading or term (used with 2nd indicator of 7)

Example: 651 #0 \$a United States
 \$x History
 \$v Chronology.

Notice that subfields \$v, \$x, and \$z in the 651 field are repeatable. Subfields \$v, \$x, \$y, and \$z do not have to be in alphabetical order. They will be in the order prescribed by the instructions given by the subject heading system.

700 Added entry -- Personal name (R)

Indicator 1: Type of personal name entry element

- 0 -- Forename
- 1 -- Surname (this is the most common form)
- 3 -- Family name

Indicator 2: Type of added entry

- # -- No information provided (most common; co-authors, editors, etc.)
- 2 -- Analytical entry (The values for Indicator 2 changed in 1994 with Format Integration, and older records may display additional values. An analytical entry involves an author/title of an item contained in a work.)

Subfields used most often:

- \$a** -- Personal name
- \$b** -- Numeration
- \$c** -- Titles and other words associated with a name (R)
- \$q** -- Fuller form of name
- \$d** -- Dates associated with a name (generally, year of birth)
- \$e** -- Relator term (such as ill.) (R)
- \$4** -- Relator code (R)

Example: 700 1# \$a Baldrige, Letitia.

710 Added entry -- Corporate name (R)

Indicator 1: Type of corporate name entry element

- 0 -- Inverted name (not used with AACR2)
- 1 -- Jurisdiction name
- 2 -- Name in direct order

Indicator 2: Type of added entry.

- See Indicator 2 under 700
- # -- No information provided
- 2 -- Analytical entry

Subfields used most often:

- \$a** -- Corporate name or jurisdiction name as entry element
- \$b** -- Subordinate unit (R)

Example: 710 2# \$a Sunburst Communications (Firm)

740 Added entry -- Uncontrolled related/analytical title (R)

Indicator 1: Nonfiling characters

0-9 -- Number of nonfiling characters present (for initial articles, including spaces)

Indicator 2: Type of added entry. See Indicator 2 under 700

-- No information provided

2 -- Analytical entry

(This field was redefined in 1994 with Format Integration. Prior to 1994, the field was also used for variant titles, such as a different wording on a spine title. In records created since Format Integration, those variant titles appear in a 246 field.)

Subfield used most often:

\$a -- Title

Example: 740 02 \$a Uncle Vanya.

800 Series added entry -- Personal name (R)

Indicator 1: Type of personal name entry element

0 -- Forename

1 -- Surname

3 -- Family name

Indicator 2 undefined.

Subfields used most often:

\$a -- Personal name

\$b -- Numeration

\$c -- Titles and other words associated with a name (R)

\$q -- Fuller form of name

\$d -- Dates associated with a name (generally, year of birth)

\$t -- Title of a work (the series)

\$v -- Volume number

Example: 800 1# \$a Fisher, Leonard Everett.
 \$t Colonial American craftsmen.

830 Series added entry -- Uniform title (R)

Indicator 1 undefined.

Indicator 2: Nonfiling characters

0-9 -- Number of nonfiling characters present (for initial articles, including spaces)

Subfield used most often:

\$a -- Uniform title

\$v -- Volume number

Example: 830 #0 \$a Railroads of America (Macmillan)

[[Back to Top of Page](#)]

Part VIII:

A List of Other Fields Often Seen in MARC Records

001	Control number
003	Control number identifier
005	Date and time of latest transaction
006	Fixed-length data elements -- additional material characteristics
007	Physical description fixed field
008	Fixed length data elements (See Part X)
022	International Standard Serial Number (ISSN)
037	Source of acquisition
041	Language code
043	Geographic area code
050	Library of Congress call number
060	National Library of Medicine call number
082	Dewey Decimal classification number (the one recommended by the Library of Congress; locally-assigned call numbers may appear elsewhere)
110	Main entry -- Corporate name (less frequent under AACR2 rules)
256	Computer file characteristics
263	Projected publication date (indicates a CIP -- Cataloging in Publication -- record)
306	Playing time
508	Creation/production credits note
510	Citation/references note (review sources)
511	Participant or performer note
521	Target audience note (first indicator: 0 = reading grade level, 1 = interest age level, 2 = interest grade level, 3 = special audience characteristics, 4 = motivation interest level)
530	Additional physical form available note
538	System details note
586	Awards note
656	Index term -- Occupation
730	Added entry -- Uniform title
852	Location

- 856 Electronic location and access
- 9XX Reserved for local use. (They are used by vendors, systems, or individual libraries to exchange additional data)

[\[Back to Top of Page \]](#)

Part IX:

The Leader

There are 24 positions in the Leader, numbered from 00 to 23. For fuller explanation, see the *MARC 21 Format for Bibliographic Data*.

- 00-04 Record length (calculated by the computer for each record)
- 05 Record status
 - a = increase in encoding level
 - c = corrected or revised
 - d = deleted
 - n = new
 - p = increase in encoding from prepublication (previous CIP)
- 06 Type of record
 - a = language material
 - c = printed music
 - d = manuscript music
 - e = cartographic material
 - f = manuscript cartographic material
 - g = projected medium
 - i = nonmusical sound recording
 - j = musical sound recording
 - k = 2-dimensional nonprojectable graphic
 - m = computer file
 - o = kit
 - p = mixed materials
 - r = 3-dimensional artifact or naturally occurring object
 - t = manuscript language material
- 07 Bibliographic level
 - a = monographic component part
 - b = serial component part
 - c = collection
 - d = subunit
 - i = integrating resource
 - m = monograph/item
 - s = serial
- 08 Type of control
 - # = no specified type
 - a = archival

- 09 **Character coding scheme**
 # = MARC-8
 a = UCS/Unicode
- 10 **Indicator count** (always "2")
- 11 **Subfield code count** (always "2")
- 12-16 **Base address of data** (calculated by the computer for each record)
- 17 **Encoding level**
 # = full level
 1 = full level, material not examined
 2 = less-than-full level, material not examined
 3 = abbreviated level
 4 = core level
 5 = partial (preliminary) level
 7 = minimal level
 8 = prepublication level (CIP)
 u = unknown
 z = not applicable
- 18 **Descriptive cataloging form**
 # = non-ISBD
 a = AACR2
 i = ISBD
 u = unknown
- 19 **Multipart resource record level**
 # = Not specified or not applicable
 a = Set
 b = Part with independent title
 c = Part with dependent title
- 20 **Length of the length-of-field portion** (always "4")
- 21 **Length of the starting-character-position portion** (always "5")
- 22 **Length of the implementation-defined portion** (always "0")
- 23 **Undefined** (always "0")

[\[Back to Top of Page \]](#)

Part X:

Field 008 for Books

Field 008 is used for Fixed Length Data Elements ("Fixed Field Codes"). There are 40 character positions in field 008, numbered from 00-39. Undefined positions must contain either a blank (#) or a fill character (|). Positions 00-17 and 35-39 are defined the same way for all media.

The information shown here for positions 18-34 applies only to books. For explanation of all the positions below and for positions 18-34 for other media, see the *MARC 21 Format*

for Bibliographic Data.

Note that field 008 has no indicators or subfield codes.

- 00-05 Date entered on file (YYMMDD),
 where Y=year, M=month, and D=day
- 06 Type of date/publication status:
 b = no dates given; B.C. date involved
 e = detailed date
 s = single known date/probable date
 m = multiple dates
 r = reprint/reissue date (Date 1) and original date (Date 2)
 n = dates unknown
 q = questionable date
 t = publication date and copyright date
 | = no attempt to code
- 07-10 Date 1/beginning date of publication
- 11-14 Date 2/ending date of publication

Date fields contain the year(s) of publication. The type of date(s) in these elements are specified in fixed field element 06: Type of date/publication status. (For further details, see the field 008 description in the *MARC 21 Format for Bibliographic Data*.)

- 15-17 Place of publication, production, or execution
 For example:
 pk# = Pakistan
 cau = California (US)

(For a full list of codes used in these positions, see the [MARC Code List for Countries](#).)

- 18-21 Illustrations (up to 4 codes):
 # = no illustrations
 a = illustrations
 b = maps
 c = portraits
 d = charts
 e = plans
 f = plates
 g = music
 h = facsimiles
 i = coats of arms
 j = genealogical tables
 k = forms
 l = samples
 m = phonodisc, phonowire, etc.
 o = photographs
 p = illuminations
 | = no attempt to code

- 22 **Target audience:**
= unknown or not specified
a = preschool
b = primary
c = pre-adolescent
d = adolescent
e = adult
f = specialized
g = general
j = juvenile
| = no attempt to code
- 23 **Form of item:**
= none of the following
a = microfilm
b = microfiche
c = microopaque
d = large print
f = braille
r = regular print reproduction
s = electronic
| = no attempt to code
- 24-27 **Nature of contents (up to 4):**
= no specified nature of contents
a = abstracts/summaries
b = bibliographies (is one or contains one)
c = catalogs
d = dictionaries
e = encyclopedias
f = handbooks
g = legal articles
i = indexes
j = patent document
k = discographies
l = legislation
m = theses
n = surveys of literature
o = reviews
p = programmed texts
q = filmographies
r = directories
s = statistics
t = technical reports
u = standards/specifications
v = legal cases and notes
w = law reports and digests
z = treaties
| = no attempt to code

- 28 **Government publication:**
 # = not a government publication
 i = international intergovernmental
 f = federal/national
 a = autonomous or semi-autonomous component
 s = state, provincial, territorial, dependent, etc.
 m = multistate
 c = multilocal
 l = local
 z = other type of government publication
 o = government publication -- level undetermined
 u = unknown if item is government publication
 | = no attempt to code
- 29 **Conference publication:**
 0 = not a conference publication
 1 = conference publication
 | = no attempt to code
- 30 **Festschrift:**
 0 = not a festschrift
 1 = festschrift
 | = no attempt to code
- 31 **Index:**
 0 = no index
 1 = index present
 | = no attempt to code
- 32 **Undefined (since 1990)** (Earlier records may contain the values 0 or 1)
 # = Undefined
 | = no attempt to code
- 33 **Literary form:**
 0 = not fiction (not further specified)
 1 = fiction (not further specified)
 c = comic strips
 d = dramas
 e = essays
 f = novels
 h = humor, satires, etc.
 i = letters
 j = short stories
 m = mixed forms
 p = poetry
 s = speeches
 u = unknown
 | = no attempt to code
- 34 **Biography:**
 # = no biographical material
 a = autobiography
 b = individual biography

- c = collective biography
- d = contains biographical information
- | = no attempt to code

35-37 Language:

A three-letter code. For example: eng fre ger spa rus ita

(For a full list of codes used in these positions, see the [MARC Code List for Languages](#).)

38 Modified record:

- # = not modified
- x = missing characters (because of characters unavailable in MARC character set)
- s = shortened
- d = "dashed-on" information omitted
- r = completely romanized/printed cards in script
- o = completely romanized/printed cards romanized
- | = no attempt to code

39 Cataloging source:

- # = national bibliographic agency
- c = cooperative cataloging program
- d = other sources
- u = unknown
- | = no attempt to code

[[Back to Top of Page](#)]

[[Back to Table of Contents](#)] -- [[Continue to Part 11](#)]



Library of Congress

[Library of Congress Help Desk](#) (10/27/2009)

APPENDIX 1005A


5ER

IEEE NETWORK[®]

March/April 1999 Vol. 13 No. 2

THE MAGAZINE OF GLOBAL INFORMATION EXCHANGE

www.comsoc.org



*Transmission and Distribution of
Digital Video:
Field Trials and Prototype
Implementations*




A Publication of the IEEE Communications Society
in cooperation with the IEEE Computer Society and the Internet Society

Ex. 1039

CISCO SYSTEMS, INC. / Page 92 of 245

EXTREME!

Extreme Broadband Testing From the Edge to the Core



The new AX/4000 Version II modules from Adtech are smashing broadband testing barriers you never knew you were up against. Like being able to test ATM and Frame Relay with the same module. Or running ATM and Frame relay interworking tests inside the same chassis. Or adding future hardware upgrades that support even more broadband technologies by simply restarting the AX/4000 system.

And that's only the beginning. Version II modules pack plenty of performance too. Like automatic identification of up to 2,040 substreams based on VPI/VCI or DLCI. Not to mention continuous full-rate analysis for all 2,040 substreams with full QoS and performance measurements and individual policing.

And at the same time they are taking measurements of incoming streams, Version II modules also generate their own 155.52 Mbps traffic with eight fully programmable traffic sources per port.

Of course we could go on and on, about feature after feature, such as our remote multi-user access, or GPS synchronized point-to-point network tests, and bi-directional analysis and decoding, but that might be just a little too

EXTREME!

Adtech, Inc. 3465 Waiālae Avenue, Suite 200 Honolulu, HI 96816
800-348-0080, 808-734-3300 EMAIL: contact@adtech-inc.com

www.comsoc.org/productinnovations#2
OR Circle #2 on reader service card

How EXTREME can a test system be? Imagine any of these...

A test system so scaleable that it can test everything from a single ATM chip to a huge world-wide network with dozens of switches and combined ATM and Frame Relay technologies.

A totally modular mix-and-match test system with 12 different test modules including 622 Mbps ATM, 23 interchangeable interfaces, and two chassis types to choose from.

A multi-port test system that can easily maintain 500,000 measurements at up to 622.08 Mbps continuously, at full-rate, and in real time.

A single test system that multiple users can access simultaneously to perform their own independent tests.

A "virtual" test system consisting of several Version II and other test modules, each located in a different city, but operating together as a single unit under the remote control of one person and one application.

**AX/4000
VERSION**

If you can imagine a test system like any or all of the above for testing your own network and switches, give us a call.

*If you can't, call us anyway.
We'll help you get the picture.*

800-348-0080

www.adtech-inc.com



ADTECH™

A4150A-03

IEEE NETWORK[®]

THE MAGAZINE OF GLOBAL INFORMATION EXCHANGE

MARCH/APRIL 1999 VOL. 13 NO. 2

F E A T U R E S

8

Internet QoS: A Big Picture



A framework for the emerging Internet QoS is presented, including RSVP, differentiated services, MPLS, and constraint-based routing.

Xipeng Xiao and Lionel M. Ni

20

Admission Control for Statistical QoS: Theory and Practice



A broad set of proposed admission control algorithms, key components of future multi-service networks, is described.

Edward W. Knightly and Ness B. Shroff

30

The Integration of Real-Time I/O and Network Support in the Stony Brook Video Server



How a real-time I/O subsystem interacts with the RETHER subsystem, and the performance cost of this integration, are described.

Tzi-cker Chiueh, Chitra Venkatramani, and Michael Vernick

37

AMUSE: Advanced Broadband Services Trials for Residential Users



The AMUSE project within the ACTS program is presented, and different technologies under experimentation are explored.

Marco Di Concetto, Giovanna Pavarani, Corrado Rosa, Fabio Rossi, Siegfried Paul, and Pierluigi Di Martino

46

Video Multicast over the Internet



Techniques that aim to improve the fairness among receivers by sending video in multiple (layered or replicated) streams are considered.

Xue Li, Mostafa H. Ammar, and Sanjoy Paul

61

Multimedia Service Interworking over Heterogeneous Networking Environments



The authors discuss technical issues related to delivery of multimedia services over heterogeneous networking environments.

Sehyeong Cho and Youngmee Shin

70

A CORBA-Based Quality of Service Management Framework for Distributed Multimedia Services and Applications



QoS management system for managing and controlling the QoS of a distributed multimedia system called MAESTRO is described.

James Won-Ki Hong, Jong-Seo Kim, and Jae-Kyu Park



THIS ISSUE

includes the second installment of a two-part Special Issue on Transmission and Distribution of Digital Video.

Cover graphics:
Image Bank

Icons explained



The mortarboard indicates a tutorial article.



The light bulb indicates an article that describes a unique idea or a new approach.



The meshing gears indicate an article that describes an implementation.



The graph indicates an article that contains measurements from a real system.



The sigma indicates an article that contains theory or analysis.



The clock indicates a timely article.

DEPARTMENTS

Editor's Note	2
IEEE 802 Perspectives	4
Guest Editorial	6

Director of Magazines
Hussein Mouftah, Queen's University, Canada
Editor-in-Chief
Jörg Liebeherr, University of Virginia, USA
Senior Technical Editors
Salah Aidarous, NEC America, USA
Thomas M. Chen, Southern Methodist U., USA
Biswanath Mukherjee, Univ. of California-Davis, USA

IEEE Network Technical Editorial Board
Anindo Banerjee, USC/ISI, USA
Jon Bennett, Harvard University, USA
Ernst Biersack, Eurecom, France
Chatschik Bisdikian, IBM Research, USA
Jim Carlo, Texas Instruments, USA
Anna Charny, MIT, USA
Jon Crowcroft, University College London, UK
John Daigle, University of Mississippi, USA
Julio Escobar, SENACYT, Panama
Sally Floyd, Lawrence Berkeley National Laboratory, USA
David Greaves, University of Cambridge, UK
Paolo Giacomazzi, Politecnico di Milano, Italy
Alden Jackson, BBN, USA
Nikhil Jain, Qualcomm, USA
Tim King, British Telecom, USA
Edward Knightly, Rice University, USA
Yi-Bing (Jason) Lin, National Chiao Tung University, Taiwan
Frank Magee, AT&T Research, USA
Allison Mankin, USC/ISI
Ioannis Nikolaidis, University of Alberta, Canada
Peter O'Reilly, GTE Laboratories, USA
Guru Parulkar, Washington University, St. Louis, USA
Krzysztof Pawlikowski, University of Canterbury, New Zealand
Stephen Pink, SICS, Sweden
K. K. Ramakrishnan, AT&T Research, USA
Kazem Sohraby, Bell Labs, USA
Martha Steenstrup, BBN, USA
C.-K. Toh, Georgia Tech, USA
Joe Touch, USC/ISI, USA
Vittorio Trecordi, CEFRIEL, Italy
Bulent Yener, Bell Labs, USA
Martina Zitterbart, TU Braunschweig, Germany

Feature Editors
Jim Carlo, "IEEE 802 Perspectives," Texas Instruments, USA
Allison Mankin, "IETF Perspectives," USC/ISI, USA
John D. Spragins, "New Books and Multimedia," Clemson Univ., USA

IEEE Production Staff
Joseph Milizzo, Publications Manager
Elizabeth Wilber, Managing Editor
Catherine Kemelmacher, Production Editor
Eric Levine, Advertising Sales Manager
Joanne O'Rourke, Staff Assistant
Susan Lange, Layout Editor
Jennifer Porcello, Editorial Assistant

1999 IEEE Communications Society Officers
Thomas J. Plevyak, *President*
Stephen B. Weinstein, *Past President*
Roberto de Marca, *President-Elect*
Douglas N. Zuckerman, *VP-Technical Activities*
Curtis A. Siller, *VP-Membership Services*
Ron Horn, *VP-Membership Development*
Tetsuya Miki, *VP-Society Relations*
Ross C. Anderson, *Treasurer*
John M. Howell, *Secretary*

Board of Governors
The elected officers above plus Members-at-Large:
Class of 1999

Jeffrey H. Derby — Ronald W. Horn
Harold Sobol — Desmond P. Taylor
Class of 2000

Gerhard Fettweis — Paul Hartmann
Michael Kincaid — William R. Robinson
Class of 2001

Laura Cerchio, Leonard Cimini
Roberta Cohen, William Tranter

1999 IEEE Officers
Kenneth R. Laker, *President*
Bruce A. Eisenstein, *President-Elect*
Maurice Papo, *Secretary*
David A. Conner, *Treasurer*
Joseph Bordogna, *Past President*
Daniel J. Senese, *Executive Director*
David G. Daut, *Director, Division III*



EDITOR'S NOTE

Eliminate Barriers of Time and Space



Jörg Liebeherr

What are the most important research challenges in network research?" It is difficult to think of a question which may heat up a conversation among a group of *IEEE Network* readers more than this one. Indeed, when a group of panelists was asked to respond to this question at a recent workshop, the range of answers covered a wide spectrum.

One of the panelists, however, gave a response which I thought did not exist — a response which is as simple as it is agreeable. The panelist, Hilarie Orman from DAPRA, stated that there is only one grand challenge in network research: to eliminate all barriers of time and space.

I see several implications of Dr. Orman's response, all of which I find encouraging. For one, trying to eliminate the barriers of time and space is a worthy goal and a great contribution to society at large. Furthermore, I find it useful to identify the existence of a common thread which holds together the very diverse and large networking community. Finally, the formulation of the grand challenge implies that a ranking of importance of networking areas is not only futile, but also counter-productive. In other words, comparing the relative importance of "universal wireless access" and "Terabit networks" distracts from the fact that both serve the same greater goal.

This issue contains the second part of a Special Issue on "Transmission and Distribution of Digital Video" by guest editors Bhupim Khasnabish and Anindo Banerjee. The articles in this part focus on testbeds and trials of digital video networks. The guest editors provide an introduction to the articles in their Guest Editorial.

In addition to the guest-edited articles, this issue also contains two feature articles. Both articles shed light on the state-of-the-art of QoS in packet networks. The first article, "Internet QoS: The Big Picture" by Xipeng Xiao and Lionel M. Ni discusses how constrained-based routing, differentiated services, and label switching can constitute the building blocks of a QoS framework for the Internet. The second article, "Admission Control for Statistical QoS: Theory and Practice," reviews the state of the art of our understanding of the multiplexing gain in packet switching networks. I hope you enjoy this issue of *IEEE Network*. I appreciate any feedback sent to me via e-mail to jorg@cs.virginia.edu. Please visit our website at <http://www.comsoc.org/socstr/techcom/ntrwk> for up-to-date information on calls-for-papers of upcoming guest-edited issues.

IEEE NETWORK ISSN 0890-8044 is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc. Headquarters address: IEEE, 3 Park Avenue, 17th Floor, New York, NY 10016-5997, USA; tel: +1-212-705-8900; e-mail: ieee.network@ieee.org. Responsibility for the contents rests upon authors of signed articles and not the IEEE or its members. Unless otherwise specified, the IEEE neither endorses nor sanctions any positions or actions espoused in *IEEE Network*.

ANNUAL SUBSCRIPTION: \$22 in addition to IEEE Communications Society or any other IEEE Society member dues. Non-member prices: \$115. Single copy prices: Member, \$10; Non-member, \$20.

EDITORIAL CORRESPONDENCE: Address to: Jörg Liebeherr, Editor-in-Chief, *IEEE Network*, IEEE Communications Society, 305 East 47th Street, New York, NY 10017-2394, USA; e-mail: jorg@cs.virginia.edu

COPYRIGHT AND REPRINT PERMISSIONS: Abstracting is permitted with credit to the source. Libraries are permitted to

photocopy beyond the limits of U.S. Copyright law for private use of patrons: those articles that carry a code on the bottom of the first page provided the per copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, USA. For other copying, reprint, or republication permission, write to Director, Publishing Services, at IEEE Headquarters. All rights reserved. Copyright © 1999 by the Institute of Electrical and Electronics Engineers, Inc.

POSTMASTER: Send address changes to *IEEE Network*, IEEE, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA. Printed in USA. Periodical-class postage paid at New York, NY and at additional mailing offices.

SUBSCRIPTIONS, orders, address changes should be sent to IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA. Tel. +1-908-981-0060. GST Reg# 125634188.

ADVERTISING: Advertising is accepted at the discretion of the publisher. Address correspondence to *IEEE Network*, 305 East 47th Street, New York, NY 10017-2394, USA:

The Revolution Continues...

in Perfecting Network Performance and **Now Monitoring!**



CACI revolutionized network planning in 1984 when we introduced the first network simulation tool for network designers and IT Professionals.

COMNET III continues to be the industry standard for high fidelity network performance prediction utilizing hybrid discrete event/analytical technology. **Now COMNET III has two companion products and three new add-on modules.**

NEW COMNET III Add-on Modules

Satellite and Mobile Module

Predict throughput and performance for wireless to wireline network integration.

Circuit-Switched Traffic Module

Predict the impact and throughput of circuit-switched traffic over your networks. **Model voice over IP.**

Application Profiler

Predict the impact and performance for new applications on the existing network infrastructure **before you deploy.**

Now Available!

CACI's Next Generation Technology



Patent pending decomposition flow technology provides the next generation tool for day to day network performance management. **Don't be fooled by amateur alternatives.**

- Easy to set up, learn, and use
- Lightning-fast and extremely accurate
- Scalable to any size network
- **Seamless integration with network management**



Next Generation performance monitoring software.

- Continuous-24 hour, 7days a week, 365 days a year-monitoring
- The big picture view of how users utilize the network and how applications affect it
- Doesn't bog down network traffic
- **Free Software Agents!**

**Contact us NOW for
New Break-Through Pricing
at 619.824.5200 or
on-line at www.caciasl.com**





Internet QoS: A Big Picture

Xipeng Xiao and Lionel M. Ni, Michigan State University

Abstract

In this article we present a framework for the emerging Internet quality of service (QoS). All the important components of this framework — integrated services, RSVP, differentiated services, multiprotocol label switching (MPLS), and constraint-based routing — are covered. We describe what integrated services and differentiated services are, how they can be implemented, and the problems they have. We then describe why MPLS and constraint-based routing have been introduced into this framework, how they differ from and relate to each other, and where they fit into the differentiated services architecture. Two likely service architectures are presented, and the end-to-end service deliveries in these two architectures are illustrated. We also compare ATM networks to router networks with differentiated services and MPLS. Putting all these together, we give the readers a grasp of the big picture of the emerging Internet QoS.

Today's Internet only provides *best-effort service*. Traffic is processed as quickly as possible, but there is no guarantee as to timeliness or actual delivery. With the rapid transformation of the Internet into a commercial infrastructure, demands for service quality have rapidly developed [1-4].

It is becoming apparent that several service classes will likely be demanded. One service class will provide predictable Internet services for companies that do business on the Web. Such companies will be willing to pay a certain price to make their services reliable and give their users a fast feel of their Web sites. This service class may contain a single service, or it may contain *gold service*, *silver service*, and *bronze service*, with decreasing quality. Another service class will provide low-delay and low-jitter services to applications such as Internet telephony and videoconferencing. Companies will be willing to pay a premium price to run a high-quality videoconference to save travel time and cost. Finally, best-effort service will remain for those customers who only need connectivity.

Whether mechanisms are even needed to provide quality of service (QoS) is a hotly debated issue. One opinion is that fibers and wavelength-division multiplexing (WDM) will make bandwidth so abundant and cheap that QoS will be automatically delivered. The other opinion is that no matter how much bandwidth the networks can provide, new applications will be invented to consume them; therefore, mechanisms will still be needed to provide QoS. This argument is beyond the scope of this article [5]. Here we simply note that even if bandwidth will eventually become abundant and cheap, it is not going to happen soon. For now, some simple mechanisms are definitely needed in order to provide QoS on the Internet. Our view is supported by the fact that all the major router/switch vendors now provide some QoS mechanisms in their high-end products [6-11].

The Internet Engineering Task Force (IETF) has proposed

many service models and mechanisms to meet the demand for QoS. Notably among them are the integrated services/Resource Reservation Protocol (RSVP) model [4, 12], the differentiated services (DS) model [13, 14], multiprotocol label switching (MPLS) [15], traffic engineering [16], and constraint-based routing [17].

The integrated services model is characterized by resource reservation. For real-time applications, before data are transmitted, the applications must first set up paths and reserve resources. RSVP is a signaling protocol for setting up paths and reserving resources. In differentiated services, packets are marked differently to create several packet classes. Packets in different classes receive different services. MPLS is a forwarding scheme. Packets are assigned labels at the ingress of an MPLS-capable domain. Subsequent classification, forwarding, and services for the packets are based on the labels. Traffic engineering is the process of arranging how traffic flows through the network. Constraint-based routing is to find routes that are subject to some constraints, such as bandwidth or delay requirement.

Although there are many papers on each of integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing, to the best of the authors' knowledge, they are never discussed together in a single paper. As a result, it is difficult for readers to understand the relationships among them and to grasp the big picture of the QoS framework.

In this article we give an introduction to integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing. We describe how they differ from, relate to, and work with each other to deliver QoS on the Internet. Through this, we intend to present to readers a clear overview of Internet QoS.

The organization of the rest of the article is as follows. In the next two sections, we describe integrated services, RSVP,

Flow	A stream of packets with the same source IP address, source port number, destination IP address, destination port number, and protocol ID.
Service level agreement (SLA)	A service contract between a customer and a service provider that specifies the forwarding service a customer should receive. A customer may be a user organization or another provider domain (upstream domain).
Traffic profile	A description of the properties of a traffic stream, such as rate and burst size.
Differentiated services (DS) field	The field in which the differentiated services class is encoded. It is the Type of Service (TOS) octet in the IPv4 header or the traffic class octet in the IPv6 header.
Per-hop behavior (PHB)	The externally observable behavior of a packet at a DS-compliant router.
Mechanism	A specific algorithm or operation (e.g., queuing discipline) that is implemented in a router to realize a set of one or more per-hop behaviors.
Admission control	The decision process of whether to accept a request for resources (link bandwidth plus buffer space).
Classification	The process of sorting packets based on the content of packet headers according to defined rules.
Behavior aggregate (BA) classification	The process of sorting packets based only on the contents of the DS field.
Multifield (MF) classification	The process of classifying packets based on the content of multiple fields such as source address, destination address, TOS byte, protocol ID, source port number, and destination port number.
Marking	The process of setting the DS field in a packet.
Policing	The process of handling out-of-profile traffic (e.g., discarding excess packets).
Shaping	The process of delaying packets within a traffic stream to cause it to conform to some defined traffic profile.
Scheduling	The process of deciding which packet to send first in a system of multiple queues.
Queue management	Controlling the length of packet queues by dropping packets when necessary or appropriate.
Traffic trunk	An aggregation of flows with the same service class that can be put into an MPLS label-switched path.

■ Table 1. Frequently used terminologies.

and differentiated services, their characteristics, mechanisms, and problems. A likely differentiated services architecture and the complete process for delivering end-to-end services in this architecture are also presented. MPLS and a service architecture based on MPLS are then described. We then describe traffic engineering and constraint-based routing. ATM networks and router networks are compared. Finally, we summarize the article in the last section.

The frequently used terminologies in this article are defined in Table 1.


Integrated Services and RSVP

The integrated services model [4] proposes two service classes in addition to best-effort service. They are:

- *Guaranteed service* [18] for applications requiring fixed delay bound
- *Controlled-load service* [19] for applications requiring reliable and enhanced best-effort service

The philosophy of this model is that "there is an inescapable requirement for routers to be able to reserve resources in order to provide special QoS for specific user packet streams, or flows. This in turn requires flow-specific state in the routers" [4].

RSVP was invented as a signaling protocol for applications to reserve resources [12]. The signaling process is illustrated in Fig. 1. The sender sends a PATH message to the receiver specifying the characteristics of the traffic. Every intermediate



■ Figure 1. RSVP signaling.

router along the path forwards the PATH message to the next hop determined by the routing protocol. Upon receiving a PATH message, the receiver responds with a RESV message to request resources for the flow. Every intermediate router along the path can reject or accept the request of the RESV message. If the request is rejected, the router will send an error message to the receiver, and the signaling process will terminate. If the request is accepted, link bandwidth and buffer space are allocated for the flow, and the related flow state information will be installed in the router.

Recently, RSVP has been modified and extended in several ways to reserve resources for aggregation of flows, to set up explicit routes (ERs) with QoS requirements, and to do some other signaling tasks [20–22]. Whether or not this is appropriate is a hotly debated issue in the IETF and is beyond the scope of this article.

Integrated services is implemented by four components: the

signaling protocol (e.g., RSVP), the admission control routine, the classifier, and the packet scheduler. Applications requiring guaranteed or controlled-load service must set up the paths and reserve resources before transmitting their data. The admission control routines will decide whether a request for resources can be granted. When a router receives a packet, the classifier will perform a *multifield* (MF) classification and put the packet in a specific queue based on the classification result. The packet scheduler will then schedule the packet accordingly to meet its QoS requirements.

The integrated services/RSVP architecture is influenced by the work of Ferrari *et al.* [23, 24]. It represents a fundamental change to the current Internet architecture, which is founded on the concept that all flow-related state information should be in the end systems [25].

The problems with the integrated services architecture are:

- The amount of state information increases proportionally with the number of flows. This places a huge storage and processing overhead on the routers. Therefore, this architecture does not scale well in the Internet core.
- The requirement on routers is high. All routers must have RSVP, admission control, MF classification, and packet scheduling.
- Ubiquitous deployment is required for guaranteed service. Incremental deployment of controlled-load service is possible by deploying controlled-load service and RSVP functionality at the bottleneck nodes of a domain and tunneling the RSVP messages over other parts of the domain.

Differentiated Services

Because of the difficulty in implementing and deploying integrated services and RSVP, differentiated services (DS) are introduced.

An Introduction to Differentiated Services

IPv4 header contains a Type of Service (TOS) byte. Its meaning was previously defined in [26]. Applications can set three bits in the TOS byte to indicate the need for low-delay, high-throughput, or low-loss-rate service. However, choices are limited. Differentiated services defines the layout of the TOS byte (*DS field*) and a base set of packet forwarding treatments (per-hop behaviors, or PHBs) [27]. By marking the DS fields of packets differently and handling packets based on their DS fields, several differentiated service classes can be created. Therefore, differentiated services is essentially a relative-priority scheme.

In order for a customer to receive differentiated services from its Internet service provider (ISP), it must have a service level agreement (SLA) with its ISP. An SLA basically specifies the service classes supported and the amount of traffic allowed in each class. An SLA can be static or dynamic. *Static SLAs* are negotiated on a regular (e.g., monthly or yearly) basis. Customers with *dynamic SLAs* must use a signaling protocol (e.g., RSVP) to request services on demand.

Customers can mark DS fields of individual packets to indicate the desired service or have them marked by the leaf router based on MF classification.

At the ingress of the ISP networks, packets are classified, policed, and possibly shaped. The classification, policing, and shaping rules used at the ingress routers are derived from the SLAs. The amount of buffering space needed for these operations is also derived from the SLAs. When a packet enters one domain from another domain, its DS field may be remarked as determined by the SLA between the two domains.

Using the classification, policing, shaping, and scheduling mechanisms, many services can be provided, such as:

- *Premium service* for applications requiring low-delay and low-jitter service
- *Assured service* for applications requiring better reliability than best-effort service
- *Olympic service*, which provides three tiers of services: *gold*, *silver*, and *bronze*, with decreasing quality [28, 29]

Note that differentiated services only defines DS fields and PHBs. It is ISPs' responsibility to decide which services to provide.

Differentiated services is significantly different from integrated services. First, there are only a limited number of service classes indicated by the DS field. Since service is allocated in the granularity of a class, the amount of state information is proportional to the number of classes rather than the number of flows. Differentiated services is therefore more scalable. Second, sophisticated classification, marking, policing, and shaping operations are only needed at the boundary of the networks. ISP core routers need only to have *behavior aggregate* (BA) classification. Therefore, it is easier to implement and deploy differentiated services.

There is another reason the second feature is desirable for ISPs. ISP networks usually consist of boundary routers connected to customers and core routers/switches interconnecting the boundary routers. Core routers must forward packets very quickly, and therefore must be simple. Boundary routers need not forward packets very quickly because customer links are relatively slow. Therefore, they can spend more time on sophisticated classification, policing and shaping [3]. Boundary routers at the *network access points* (NAPs) are exceptions. They must forward packets very quickly and do sophisticated classification, policing, and shaping. Therefore, they must be well equipped.

In the differentiated services model, incremental deployment is possible for assured service. DS-incapable routers simply ignore the DS fields of the packets and give the assured service packets best-effort service. Since assured service packets are less likely to be dropped by DS-capable routers, the overall performance of assured service traffic will be better than that of best-effort traffic.

An End-to-End Service Architecture

In this section a service architecture for differentiated services is presented. This architecture provides assured service and premium service in addition to best-effort service. It is mainly based on the architecture proposed in [28]. Other possible service architectures also exist [30].

Assured Service — Assured service is intended for customers that need reliable services from their service providers, even in times of network congestion. Customers will have SLAs with their ISPs. The SLAs will specify the amount of bandwidth allocated for the customers. Customers are responsible for deciding how their applications share that amount of bandwidth. One possible service allocation process is described later. SLAs for assured service are usually static, meaning that customers can start data transmission whenever they want without signaling their ISPs.

Assured service can be implemented as follows. First, classification and policing are done at the ingress routers of the ISP networks. If the assured service traffic does not exceed the bit rate specified by the SLA, they are considered *in profile*; otherwise, the excess packets are considered *out of profile*. Second, all packets, *in* and *out*, are put into an *assured queue* (AQ) to avoid out of order delivery. Third, the queue is managed by a queue management scheme called *random early detection* (RED) with *In* and *Out* — *RIO* [31].

RED [32] is a queue management scheme that drops packets randomly. This will trigger the TCP flow control mecha-

nisms at different end hosts to reduce send rates at different time. By doing so, RED can prevent the queue at the routers from overflowing, and therefore avoid the tail-drop behavior (dropping all subsequent packets when a queue overflows). Tail-drop triggers multiple TCP flows to decrease and later increase their rates simultaneously. It causes network utilization to oscillate and can hurt performance significantly. RED has proven to be useful and has been widely deployed.

RIO is a more advanced RED scheme. It basically maintains two RED algorithms, one for in packets and one for out packets. There are two thresholds for each queue. When the queue size is below the first threshold, no packets are dropped. When the queue size is between the two thresholds, only out packets are randomly dropped. When the queue size exceeds the second threshold, indicating possible network congestion, both in and out packets are randomly dropped, but out packets are dropped more aggressively. In addition to breaking the TCP flow-control synchronization, RIO prevents, to some extent, greedy flows from hurting the performance of other flows by dropping the out packets more aggressively.

Because in packets have low loss even in cases of congestion, the customers will perceive a predictable service from the network if they keep traffic conformant. When there is no congestion, out packets will also be delivered. The networks are thus better utilized.


Best-effort traffic can be treated differently, or identically, from assured service out traffic. In this article we assume that they are treated identically. Therefore, conceptually, we can consider that there is an *A-bit* in the DS field. The A-bits of assured service in packets are set to 1, while the A-bits of assured service out packets and best-effort packets are reset to 0.

Premium Service — Premium service provides low-delay and low-jitter service for customers that generate fixed peak bit rate traffic. Each customer will have an SLA with its ISP. The SLA specifies a desired peak bit-rate for a specific flow or an aggregation of flows. The customer is responsible for not exceeding the peak rate; otherwise, excess traffic will be dropped. The ISP guarantees that the contracted bandwidth will be available when traffic is sent. Premium service is suitable for Internet telephony, videoconferencing, or for creating virtual lease lines for virtual private networks (VPNs) [33].

Because premium service is more expensive than assured service, it is desirable for ISPs to support both static and dynamic SLAs. Dynamic SLAs allow customers to request premium service on demand without subscribing to it. Admission control is needed for dynamic SLAs.

Premium service can be implemented as follows. At the customer side, some entity will decide which application flow can use premium service. The leaf routers directly connected to the senders will do MF classifications and shape the traffic. Conceptually, we can consider that there is a *P-bit* in the DS field. If the P-bit of a packet is set, this packet belongs to the premium class; otherwise, the packet belongs to the assured service or best-effort class. After the shaping, the P-bits of all packets are set for the flow that is allowed to use premium service. The exit routers of the customer domain may need to reshape the traffic to make sure that the traffic does not exceed the peak rate specified by the SLA. At the provider side, the ingress routers will police the traffic. Excess traffic is dropped. All packets with the P-bit set enter a *premium queue* (PQ). Packets in the PQ will be sent before packets in the AQ.

First, by admission control the amount of premium traffic can be limited to a small percentage, say 10 percent, of the bandwidth of input links. Second, excess packets are dropped at the ingress routers of the networks. Nonconformant flows cannot impact the performance of conformant flows. Third,



■ Figure 2. Uneven distribution of premium traffic in an ISP. The shaded area is the core of the ISP.

premium packets are forwarded before packets of other classes; they can potentially use 100 percent of the bandwidth of the output links. Since most links are full-duplex, the bandwidth of the input links equals the bandwidth of the output links. Therefore, if premium traffic is distributed evenly among the links, these three factors should guarantee that the service rate of the PQ is much higher than the arrival rate. Therefore, arriving premium packets should find the PQ empty or very short most of the time. The delay or jitter experienced by premium packets should be very low. However, premium service provides no quantified guarantee on the delay or jitter bound.

However, uneven distribution of premium traffic may cause a problem for premium service. In ISP networks, aggregation of traffic from the boundary routers to a core router (e.g., CR1 in Fig. 2) is inevitable; but this is not a problem because the output link is much faster than the input links. However, aggregation of premium traffic in the core at CR4 may invalidate the assumption that the arrival rate of premium traffic is far below the service rate. Differentiated services alone cannot solve this problem. Traffic engineering/constraint-based routing must be used to avoid congestion caused by such uneven traffic distribution.

By limiting the total amount of bandwidth requested by premium traffic, the network administrators can guarantee that premium traffic will not starve the assured and best-effort traffic. Another scheme is to use Weighted Fair Queuing (WFQ) [34] between the PQ and the AQ.

Service Allocation in Customer Domains — Given an SLA, a customer domain should decide how its hosts share the services specified by the SLA. This process is called *service allocation*.

There are basically two choices:

- Each host makes its own decision as to which service to use.
- A resource controller called the *bandwidth broker* (BB) [28] makes decision for all hosts.

A BB can be a host, a router, or a software process on an exit router. It is configured with the organizational policies and manages the resources of a domain. A domain may also have backup BBs. Since all hosts must cooperate to share a limited amount of resources specified by the SLA, it is technically better to have a BB to allocate resources.

At the initial deployment stage, hosts need no DS mechanism. They simply send their packets unmarked. The exit routers marked them before sending them out to the ISPs. The packets are treated as best-effort traffic inside the customer domain. In later deployment stages, hosts may have some signaling or marking mechanisms. Before a host starts sending packets, it may decide the service class for the packets by itself, or it may consult a BB for a service class. The host may mark the packets by itself or send the packets unmarked. If the host

sends the packets unmarked, the BB must use some protocols, such as RSVP or Lightweight Directory Access Protocol (LDAP) [35], to set the classification, marking, and shaping rules at the leaf router directly connected to the sender so that the leaf router knows how to mark the sender's packets.

If the SLA between a customer and its ISP is dynamic, the BB in the customer domain must also use some signaling protocol to request resources on demand from its ISP. From now on, we assume that RSVP is used as the signaling protocol.

Resource Allocation in ISP Domains — Given the SLAs, ISPs must decide how to configure their boundary routers so that they know how to handle the incoming traffic. This process is called *resource allocation*.

For static SLAs, boundary routers can be manually configured with the classification, policing, and shaping rules. Resources are therefore statically allocated for each customer. Unused resources can be shared by other customers.

For a dynamic SLA, resource allocation is closely related to the signaling process. The BB in the customer domain uses RSVP to request resources from its ISP. At the ISP side, admission control decisions can be made in a distributed manner by the boundary routers, or centrally by a BB. If boundary routers are directly involved in the signaling process, they are configured with the corresponding classification, policing, and shaping rules when they grant a request. If a BB is involved rather than the boundary routers, the BB must configure the boundary routers when it grants a request. Such procedures will be detailed in the next section.

In both cases, the ISP core routers must be shielded from the requests to avoid the scalability problem.

Examples of End-to-End Service Delivery

Example 1: Delivery of Assured Service with a Static SLA — In Fig. 3, host S in corporate network 1 (CN1) wants to use assured service to send data to host D in corporate network 2 (CN2). CN1 has a static SLA with ISP1. The numbers inside the circles are the step numbers in the service delivery process, described below.

- 1 Host S sends a RSVP message to the local BB, CN1-BB, requesting for assured service for its traffic.
- 2 If CN1-BB grants the request, it will configure leaf router LR1 so that LR1 can set the A-bits of the packets of this flow. CN1-BB will then reply to host S; otherwise, an error message is sent to host S.
- 3 Host S sends packets to leaf router LR1.
- 4 If LR1 is configured to mark the traffic, it will set the A-bits of the packets.
- 5 Every router from LR1 (exclusive) to ER1 (inclusive) does a BA classification. Packets with the A-bit set are consid-

ered in, while packets with the A-bit reset are considered out. All packets enter the AQ. RIO is applied on the AQ.

- 6 BR1 polices the traffic. All out traffic remains out. If the in traffic exceeds its bit rate, the excess packets' A-bits will be reset. All packets enter the AQ. RIO is applied on the queue.
- 7 All routers between boundary routers BR1 and BR2 (inclusive) perform BA classifications and apply RIO on their AQs.
- 8 ER2 performs the same operations as BR1.
- 9 The packets are eventually delivered to host D.


Note that:

- If there are multiple ISPs between CN1 and CN2, steps 6 and 7 will be repeated multiple times, once per ISP.
- If CN1 does not have any SLA with ISP1, it can only send traffic as best effort. No matter how the routers in CN1 mark the DS fields of their packets, the A-bits will be reset at BR1.


Example 2: Delivery of Premium Service with a Dynamic SLA — In Fig. 4, host S in CN1 wants to use premium service to send data to host D in CN2. CN1 has a dynamic SLA with ISP1.

Phase 1: Signaling

- 1 Host S sends an RSVP PATH message to the local BB, CN1-BB.
- 2 CN1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to host S. The signaling process ends.
- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to ISP1-BB.
- 4 ISP1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to CN1-BB. Sender S will be notified.
 - If the request is accepted, ISP1-BB sends the PATH message to CN2-BB.
- 5 CN2-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to ISP1-BB. Sender S will be notified.
 - If the request is accepted, CN2-BB will use LDAP or RSVP to set the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to ISP1-BB.
- 6 When ISP1-BB receives the RESV message, it will configure the classification and policing rules on router BR1, and the policing and reshaping rules on router BR2. It will then send the RESV message to CN1-BB.
- 7 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so that if the traffic of the admitted flow is nonconformant, LR1 will shape it. CN1-BB will also set the policing and reshaping



■ Figure 3. The delivery process of assured service with a static SLA.



■ Figure 4. The delivery process of premium service with a dynamic SLA.

rules on router ER1. CN1-BB will then send the RESV message to host S.

8 When host S receives the RESV message, it can start transmitting data.

Note that:

- This signaling process is significantly different from the signaling process in the integrated services/RSVP model. First, it is the sender who requests resources, not the receiver. Second, a request can be rejected when the BB receives the PATH message from the sender. In integrated services/RSVP, a request is rejected only when a router receives the RESV message from the receiver. Third, a BB can aggregate multiple requests and make a single request to the next BB. Fourth, each domain behaves like a single node, represented by the BB. ISP core routers are not involved in this process.
- The state information installed by the BB on the boundary routers is soft state. It must be regularly refreshed, or it will time out.
- If there are multiple ISPs between CN1 and CN 2, repeat steps 4 and 6 once for each ISP.
- If the SLA between CN1 and ISP1 is static, simply skip steps 3–6 in the signaling process.

Phase 2: Data Transmission

9 Host S sends packets to leaf router LR1.

10 Leaf router LR1 performs an MF classification. If the traffic in nonconformant, LR1 will shape it. LR1 will also set the P-bits of the packets. All packets enter the PQ.

11 Each intermediate router between LR1 and ER1 performs a BA classification, puts the packets into the PQ, and sends them out.

12 ER1 performs a BA classification and reshapes the traffic to make sure that the negotiated peak rate is not exceeded. Reshaping is done for the aggregation of all flows heading toward BR1, not for each individual flow.

13 BR1 classifies and polices the premium traffic. Excess premium packets are dropped.

14 Intermediate routers between leaf routers BR1 and BR2 (inclusive) perform BA classifications. BR2 also reshapes the premium traffic.

15 ER2 classifies and polices the premium traffic. Excess premium packets are dropped.

16 The premium packets are delivered to host D.

Requirements on Routers

The requirements on routers to support premium service and assured service are summarized below.

- The leaf routers in customer domains need MF classifications, marking, and shaping.

- The ISP ingress routers need policing and remarking.

- The ISP egress routers optionally need re-shaping.

- All routers need BA classification and two queues with strict priority.

- If dynamic SLAs are supported, each customer domain will need a BB to request service on behalf of the domain and to allocate services inside the domain. Signaling and admission control mechanisms are needed in both customer and ISP domains.

If assured service is to be replaced by olympic service, the AQ must be replaced by three queues: a *gold queue*, a *silver queue*, and a *bronze queue*. WFQ can be used to schedule these queues. The rate parameters of these queues can be manually configured based on experience.

MPLS


The motivation for MPLS is to use a fixed-length label to decide packet handling. MPLS is also a useful tool for traffic engineering [16, 36].

An Introduction to MPLS

MPLS is a forwarding scheme. It evolved from Cisco's *Tag Switching*. In the open systems interconnection (OSI) seven-layer model, it is between layer 2 (L2, link layer) and layer 3 (L3, network layer).

Each MPLS packet has a header. The header contains a 20-bit label, a 3-bit Class of Service (COS) field, a 1-bit label stack indicator, and an 8-bit Time to Live (TTL) field. The MPLS header is encapsulated between the link layer header and the network layer header. An MPLS-capable router, called the *label-switched router* (LSR), examines only the label in forwarding the packet. The network protocol can be IP or others. This is why it is called *multiprotocol* label switching.

MPLS needs a protocol to distribute labels to set up *label-switched paths* (LSPs). Whether a generic *label distribution protocol* (LDP) [37] should be created or RSVP should be extended [22] for this purpose is another hotly debated issue. MPLS labels can also be piggybacked by routing protocols. An LSP is similar to an ATM virtual circuit (VC) and is unidirectional from the sender to the receiver. MPLS LSRs use the protocol to negotiate the semantics of each label, that is, how to handle a packet with a particular label from the peer. LSP setup can be control-driven (i.e., triggered by control traffic such as routing updates) or data-driven (i.e., triggered by the request of a flow or *traffic trunk*). In MPLS, a traffic trunk is an aggregation of flows with the same service class that can be put into an LSP. The LSP between two routers can be the same as the L3 hop-by-hop route, or the sender LSR can specify an *explicit route* (ER) for the LSP. The ability to set up



■ Figure 5. The signaling process of dynamic premium service in an MPLS-based architecture.

ERs is one of the most useful features of MPLS. A forwarding table indexed by labels is constructed as the result of label distribution. Each forwarding table entry specifies how to process packets carrying the indexing label.

Packets are classified and routed at the ingress LSRs of an MPLS-capable domain. MPLS headers are then inserted. When an LSR receives a labeled packet, it will use the label as the index to look up the forwarding table. This is faster than the process of parsing the routing table in search of the longest match done in IP routing [38, 39]. The packet is processed as specified by the forwarding table entry. The incoming label is replaced by the outgoing label, and the packet is switched to the next LSR. This label-switching process is similar to ATM's VCI/VPI processing. Inside an MPLS domain, packet forwarding, classification, and QoS service are determined by the labels and the COS fields. This makes core LSRs simple. Before a packet leaves an MPLS domain, its MPLS label is removed.

MPLS LSPs can be used as tunnels. After LSPs are set up, a packet's path can be completely determined by the label assigned by the ingress LSR. There is no need to enumerate every intermediate router of the tunnel. Compared to other tunneling mechanisms, MPLS is unique in that it can control the complete path of a packet without explicitly specifying the intermediate routers.

In short, MPLS is strategically significant because:

- It provides faster packet classification and forwarding.
- It provides an efficient tunneling mechanism.

These features, particularly the second one, make MPLS useful for traffic engineering [16, 36].

A Service Architecture Based on MPLS

MPLS can be used together with differentiated services to provide QoS. In such an architecture, LSPs are first configured between each ingress-egress pair. For LSP(LSR1 → LSR2) and LSP(LSR2 → LSR1), their intermediate LSRs need not be reciprocal. It is likely that for each ingress-egress pair, a separate LSP is created for each traffic class. In this case, a total number of $C*N*(N-1)/2$ LSPs are needed, where C is the number of traffic classes and N is the number of boundary routers. In order to reduce the number of LSPs, the LSPs from all ingress routers to a single egress router can be merged into a *sink tree*. The total number of sink trees needed is $C*N$. It is also possible to use a single sink tree to transmit packets of different traffic classes, and use the COS bits to differentiate packet classes. In this case, the number of sink trees is reduced to N . In this architecture, as the number of transiting flows increases, the number of flows in each LSP or sink tree also increases. But the number of LSPs or sink trees need not increase. This architecture is therefore scalable.

The operations of the routers are basically the same in this architecture as in the DS-field-based architecture described earlier. There are three differences in the processing of a packet:

- At the ingress of the ISP network, in addition to all the processing described in the DS-field-based architecture, an MPLS header is inserted into the packet.
- Core routers process the packet based on its label and COS field rather than its DS field.
- At the egress, unless interdomain LSPs are configured, the MPLS header is removed.

Note that, with such schemes, the MPLS effect is confined within the ISPs that use MPLS. Whether a particular ISP's architecture is DS-field-based or MPLS-based is transparent to other ISPs. Therefore, the DS-field based and MPLS-based architectures can easily interoperate.

Each customer domain still needs a BB to allocate services, and to request resources on behalf of the customer domain when the SLA is dynamic. But since LSPs are configured within the ISPs, resource requests can easily be hidden from the core routers by tunneling them from the ingress routers to the egress routers. Therefore, BBs may not be needed in MPLS-based ISP networks. Admission control is done in a distributed fashion by the ingress and egress routers.

Without BBs in the ISP networks, the signaling process for dynamic SLAs is slightly different from the one described earlier. It is depicted in Fig. 5 and described below. The data transmission process remains the same except for the three differences noted above.

- 1 Host S sends an RSVP PATH message to its local domain BB CN1-BB.
- 2 CN1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to host S. The signaling process ends.
- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to BR1.
- 4 BR1 decides if there are enough resources to send the traffic to egress router BR2.
 - If no, the request is denied. An error message is sent back to CN1-BB. Sender S will be notified.
 - If yes, ISP1-BB sends the PATH message through an LSP to BR2.
- 5 BR2 sends the PATH message to CN2-BB.
- 6 CN2-BB decides if its domain can support the traffic.
 - If no, the request is denied. An error message is sent back to BR2. Sender S will be notified.
 - If yes, the request is accepted. CN2-BB will use LDAP or RSVP to set the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to BR2.
- 7 BR2 configures the reshaping rules for the traffic. It then sends the RESV message through an LSP to BR1.

8 BR1 configures the classification and policing rules for the traffic. It then sends the RESV message to CN1-BB.
 9 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so if the traffic of the admitted flow is nonconformant, LR1 can shape it. CN1-BB will also set the reshaping rules on router ER1. CN1-BB will then pass the RESV message to host S.
 10 Sender S starts transmitting data.
 If there are multiple ISPs between CN1 and CN2, repeat steps 4, 5, 7, and 8 once per ISP.

Traffic Engineering and Constraint-Based Routing

QoS schemes such as integrated services/RSVP and differentiated services essentially provide differentiated degradation of performance for different traffic when traffic load is heavy. When the load is light, integrated services/RSVP, differentiated services, and best-effort service differ little. Then why not avoid congestion at the first place? This is the motivation for traffic engineering.

Traffic Engineering

Network congestion can be caused by lack of network resources or uneven distribution of traffic. In the first case, all routers and links are overloaded, and the only solution is to provide more resources by upgrading the infrastructure. In the second case, some parts of the network are overloaded while other parts are lightly loaded. Uneven traffic distribution can be caused by the current dynamic routing protocols such as RIP, OSPF, and IS-IS, because they always select the shortest paths to forward packets. As a result, routers and links along the shortest path between two nodes may become congested while routers and links along a longer path are idle. The *equal-cost multipath* (ECMP) option of OSPF [40], and recently of IS-IS [41], is useful in distributing load to several shortest paths; but if there is only one shortest path, ECMP does not help. For simple networks, it may be possible for network administrators to manually configure the cost of the links so that traffic can be evenly distributed. For complex ISP networks, this is almost impossible.

Traffic engineering is the process of arranging how traffic flows through the network so that congestion caused by uneven network utilization can be avoided. *Constraint-based routing* is an important tool for making the traffic engineering process automatic.

Avoiding congestion and providing graceful degradation of performance in congestion are complementary. Traffic engineering therefore complements differentiated services.

Constraint-Based Routing

In a sentence, constraint-based routing is used to compute routes that are subject to multiple constraints.

Constraint-based routing evolves from *QoS routing*. Given the QoS request of a flow or an aggregation of flows, QoS routing returns the route that is most likely to be able to meet the QoS requirements. Constraint-based routing extends QoS routing by considering other constraints of the network such as policy. The goals of constraint-based routing are:

- To select routes that can meet certain QoS requirements
- To increase utilization of the network

While determining a route, constraint-based routing considers not only network topology, but also requirements of the flow, resource availability of the links, and possibly other policies specified by the network administrators. Therefore, constraint-based routing may find a longer but lightly loaded path

better than the heavily loaded shortest path. Network traffic is thus distributed more evenly.

In order to do constraint-based routing, routers need to distribute new link state information and to compute routes based on such information.

Distribution of Link State Information — A router needs topology and resource availability information in order to compute QoS routes. Here, resource availability information means link available bandwidth [42]. Buffer space is assumed to be sufficient and is not explicitly considered [42].

One approach to distributing bandwidth information is to extend the link state advertisements of protocols such as OSPF and IS-IS [42, 43]. Because link residual bandwidth is frequently changing, a trade-off must be made between the need for accurate information and the need to avoid frequent flooding of link state advertisements.

To reduce the frequency of link state advertisements, one possible way is to distribute them only when there are topology or significant bandwidth changes (e.g., more than 50 percent or more than 10 Mb/s) [44]. A hold-down timer should always be used to limit the frequency of such advertisements. A recommended timer value is 30 seconds [45]. An approach to limit the flooding scope of such advertisements is described in [46].

Route Computation — The routing table computation algorithms in constraint-based routing and the complexity of such algorithms depend on the metrics chosen for the routes.

Common route metrics in constraint-based routing are monetary cost, hop count, bandwidth, reliability, delay, and jitter. Routing algorithms select routes that optimize one or more of these metrics.

Metrics can be divided into three classes. Let $d(i, j)$ be a metric for link (i, j) . For any path $P = (i, j, k, \dots, l, m)$, metric d is:

- *Additive* if $d(P) = d(i, j) + d(j, k) + \dots + d(l, m)$
- *Multiplicative* if $d(P) = d(i, j) * d(j, k) * \dots * d(l, m)$
- *Concave* if $d(P) = \min\{d(i, j), d(j, k), \dots, d(l, m)\}$

According to this definition, metrics *delay*, *jitter*, *cost*, and *hop count* are additive, *reliability* ($1 - \text{loss rate}$) is multiplicative, and *bandwidth* is concave.

A well-known theorem in constraint-based routing is that computing optimal routes subject to constraints of two or more additive and/or multiplicative metrics is *NP-complete* [47]. That is, algorithms that use two or more of delay, jitter, hop count, and loss probability as metrics, and optimize them simultaneously are NP-complete. The computationally feasible combinations of metrics are bandwidth and one of those metrics.

However, the proof of NP-completeness in [47] is based on the assumptions that all the metrics are independent, and the delay and jitter of a link are known a priori. Although such assumptions may be true in circuit-switched networks, metrics bandwidth, delay, and jitter are not independent in packet networks. As a result, polynomial algorithms for computing routes with hop count, delay, and jitter constraints exist [45]. The complexity of such algorithms is $O(N * E * e)$, where N is the hop count, E is the number of links of the network, and $e \leq E$ is the number of distinct bandwidth values among all links. Nevertheless, the theorem can tell us qualitatively the complexity of a routing algorithm: a complex algorithm in circuit-switched networks is still complex in packet networks, although it may not be NP-complete.

Fortunately, algorithms for finding routes with bandwidth and hop-count constraints are much simpler [42]. Bellman-Ford's (BF) Algorithm or Dijkstra's Algorithm can be used. For example, to find the shortest path between two nodes with bandwidth greater than 1 Mb/s, all the links with residual band-

width less than 1 Mb/s can be pruned first. BF or Dijkstra's Algorithm can then be used to compute the shortest path in the pruned network. The complexity of such algorithms is $O(N^2E)$.

Bandwidth and hop count are more useful constraints than delay and jitter, because:

- Although applications may care about delay and jitter bounds, few applications cannot tolerate occasional violation of such constraints. Therefore, there is no obvious need for routing flows with delay and jitter constraints. Besides, since delay and jitter parameters of a flow can be determined by the allocated bandwidth and the hop count of the route [48, 49], delay and jitter constraints can be mapped to bandwidth and hop-count constraints if needed.
- Many real-time applications will require a certain amount of bandwidth. The bandwidth metric is therefore useful. The hop count metric of a route is important because the more hops a flow traverses, the more resources it consumes. For example, a 1-Mb/s flow that traverses two hops consumes twice as many resources as one that traverses a single hop.

In constraint-based routing, routes can be computed on demand or precomputed for each traffic class. On-demand computations are triggered by the receipt of the QoS request of a flow. In either case, a router will have to compute its routing table more frequently with constraint-based routing than with dynamic routing. This is because, even without topology changes, routing table computation can still be triggered by significant bandwidth changes. Besides, constraint-based routing algorithms are at least as complex as dynamic routing algorithms. Therefore, the computation load of routers with constraint-based routing can be very high.

Common approaches to reduce the computation overhead of constraint-based routing include:

- Using a large-valued timer to reduce computation frequency
- Choosing bandwidth and hop count as constraints
- Using administrative policy to prune unsuitable links before computing the routing table

For example, if a flow has delay requirement, high propagation delay links such as satellite links are pruned before the routing table computation.

Constraint-Based Routing: Pros and Cons

The pros of constraint-based routing are meeting the needs for QoS requirements of flows better, and improved network utilization.

The cons of constraint-based routing are increased communication and computation overhead, increased routing table size, the fact that longer paths may consume more resources, and potential routing instability.

Of the cons, the first was addressed in an earlier section. The rest are addressed in this section.

In constraint-based routing, an essential issue is routing granularity. Routing can be destination-based, source-destination-based, class-based, traffic-trunk-based, or flow-based. Routing with finer granularity is more flexible, and thus more efficient in terms of resource utilization and more stable. However, the computation overhead and storage overhead are also higher.

Routing Table Structure and Size — Routing table structure and size depend directly on routing granularity and route metrics. Logically, we can view the routing table as a two-dimensional array. The number of rows is determined by routing granularity, and the number of columns is determined by route metrics. For example, in destination-based routing with bandwidth and hop count as route metrics, the routing table can be organized as a $K \times H$ array, where K is the number of destinations, and H is the maximum number of hops allowed for any route. The (k, h) th entry of the array contains one or

more h -hop routes for destination k . Each route also has an available bandwidth associated with it [42].

Obviously, the size of a constraint-based routing table can be far larger than the size of a normal routing table for the same network. This introduces significant storage overhead. It may also slow down the routing table lookup.

Approaches to reducing the routing table size in constraint-based routing include:

- Using coarse routing granularity
- Using hop quantization (i.e., dividing all hop-count values into a few classes to reduce the number of columns in the routing table) [44]
- Keeping the routing table only for best-effort traffic, and computing the routes for flows with QoS requests on demand [50]

The third scheme basically trades computation time for smaller storage requirements.

The Trade-off between Resource Conservation and Load Balancing — A constraint-based routing scheme can choose one of the following as the route for a destination:

- The *widest-shortest* path, that is, a path with minimum hop count and, if there are multiple such paths, the one with the largest available bandwidth.
- The *shortest-widest* path, that is, a path with the largest available bandwidth and, if there are multiple such paths, the one with the minimum hop count.
- The *shortest-distance* path. The distance of a k -hop path P is defined as

$$\text{dist}(P) = \sum_{i=1}^k \frac{1}{r_i}, \text{ where } r_i \text{ is the bandwidth of link } i.$$

Using paths other than the shortest paths consume more resources. This is not efficient when the load of the network is heavy. A trade-off must be made between resource conservation and load balancing. The first approach above is basically the same as today's dynamic routing. It emphasizes preserving network resources by choosing the shortest paths. The second approach emphasizes load balancing by choosing the widest paths. The third approach makes a trade-off between the two extremes. It favors shortest paths when network load is heavy and widest paths when network load is medium. Simulations showed that the third approach consistently outperforms the other two approaches for best-effort traffic, regardless of network topology and traffic pattern [45].

Stability — Because constraint-based routing algorithms recompute routing tables more frequently than dynamic routing algorithms do, they can introduce instability.

The stability of networks with constraint-based routing depends heavily on the routing granularity. If routing is done with coarse granularity (e.g., based solely on destination address), when the original route between two nodes becomes congested, all the traffic to that destination is shifted from the original route to an alternate route. This may cause congestion in the alternate route. Traffic may have to be shifted again [46].

The high computation overhead of constraint-based routing may also hurt the stability of the network. When a router is busy computing the routing table, it is slow in reacting to new topology changes.

To improve stability, the timer value for periodic routing table recomputation should be carefully chosen [50]. Constraint-based routing at the granularity of traffic trunk provides a good trade-off between stability and computation overhead [16]. Reducing the computation complexity of the routers also helps to improve stability.

In summary, constraint-based routing must be deployed

with caution. Otherwise, the cost of instability and increased complexity may outweigh the gain.

Constraint-based routing is similar to the dynamic/adaptive routing in telephone networks and ATM networks [51–54]. Many lessons can be learned from those works. Since constraint-based routing is a superset of today's dynamic routing, it is possible that in the future, constraint-based routing may replace dynamic routing, especially in the intradomain case. An emerging intradomain constraint-based routing protocol is QOSPF [42].

The Position of Constraint-Based Routing in the QoS Framework

In this section we describe the relationships between constraint-based routing and other components in the QoS framework.

The Relationship between Constraint-Based Routing and Differentiated Services — Constraint-based routing is to select the optimal routes for flows so that their QoS requirements are most likely to be met. It is not to replace differentiated services, but to help differentiated services be better delivered. Figure 2 shows an example in point.

The Relationship between Constraint-Based Routing and RSVP — RSVP and constraint-based routing are independent but complementary. For a router with dynamic routing, when an RSVP PATH message is received, it will be forwarded to the next hop determined by the dynamic routing protocol. The QoS requirement of the flow and the load of the networks are not considered in selecting the next hop. However, with a router running constraint-based routing, such information is considered. The next hop of the RSVP messages determined by constraint-based routing therefore may be different. In either case, the actual reservation of resources for the flow is done by RSVP. In short, constraint-based routing determines the path for RSVP messages, but does not reserve resources. RSVP reserves resources, but depends on constraint-based or dynamic routing to determine the path.

The Relationship between Constraint-Based Routing and MPLS — Given that MPLS is a forwarding scheme and constraint-based routing is a routing scheme, MPLS and constraint-based routing are, in theory, mutually independent. Constraint-based routing determines the route between two nodes based on resource information and topology information. It is useful with or without MPLS. Given the routes, MPLS uses its label distribution protocol to set up the LSPs. It does not care whether the routes are determined by constraint-based or dynamic routing.

However, when MPLS and constraint-based routing are used together, they make each other more useful. MPLS makes it possible to do constraint-based routing at traffic trunk granularity without introducing MF classification to the core routers. MPLS's per-LSP statistics provide constraint-based routing with precise information about the amount of traffic between every ingress-egress pair. Given such information, constraint-based routing can better compute the routes for setting up LSPs. In combination, MPLS and constraint-based routing provide powerful tools for traffic engineering.

A Comparison of ATM Networks to Router Networks

QoS and some sort of traffic engineering have long been provided by ATM networks. So why introduce differentiated ser-

vices and MPLS into the router networks? To answer this question, we give a brief comparison between ATM networks and router networks.

In an ATM network, QoS can be provided by allocating a certain amount of bandwidth for a specific VC. Traffic engineering is usually done by computing the routes offline and then downloading the configuration statically into the ATM switches on an hourly or daily basis. Per-permanent virtual circuit (PVC) traffic statistics of the current configuration provide accurate traffic information for computing the routes for the next configuration.

The advantages of ATM networks over router networks without differentiated services or MPLS are:

- ATM networks are currently faster in data forwarding.
- Per-PVC traffic statistics are available.
- QoS and some sort of traffic engineering are provided.

The disadvantages of ATM networks are:

- ATM cell header overhead is large.
- Routers must be used at the boundary of the network. With both switches and routers present in the network, two sets of configurations are required: one for routers and the other for switches.

With differentiated services and MPLS, router networks can also provide QoS and traffic engineering. This can be done without a big header overhead and two sets of configurations. Router networks with differentiated services and MPLS therefore provide some advantages over ATM networks [55]; but this is more or less from the perspective of router vendors.

Summary

The big picture of the emerging Internet QoS can be summarized as follows:

- Customers negotiate SLAs with ISPs. The SLAs specify what services the customers will receive. SLAs can be static or dynamic. For static SLAs, customers can transmit data at any time. For dynamic SLAs, customers must use a signaling protocol such as RSVP to request services on demand before transmitting data. The bandwidth brokers in the customer domains decide how applications share the services specified by the SLAs. The DS fields of packets are marked accordingly to indicate the desired services.
- The ingress routers of ISPs are configured with classification, policing, and remarking rules. The egress routers of ISP networks are configured with reshaping rules. Such rules may be configured manually by network administrators or dynamically by some protocol such as LDAP or RSVP. ISPs must implement admission control in order to support dynamic SLAs. Classification, marking, policing, and shaping/reshaping are only done at the boundary routers. Core routers are shielded from the signaling process. They need only implement two queues with strict priority. They process packets based solely on their DS fields.
- With MPLS, LSPs are set up between each ingress-egress pair. At the ISP ingress routers, labels and COS fields are determined from the classification and routing results. MPLS headers are then inserted into the packets. Core routers process packets based on their labels and COS fields only. Labels are removed before packets leave an MPLS domain.
- Constraint-based routing can be used to compute the routes subject to QoS and policy constraints. The goal is to meet the QoS requirements of traffic and to improve utilization of the networks.
- MPLS and constraint-based routing can be used together to control the path of traffic to avoid congestion and improve the utility of the networks.

Application Layer	
Transport Layer	Integrated Service/RSVP, Differentiated Services
Network Layer	Constraint Based Routing
	MPLS
Link Layer	

■ Table 2. The relative positions of the components in the QoS framework.

- The positions of integrated services/RSVP, differentiated services, MPLS, and constraint-based routing in the Internet network model are depicted in Table 2.

Acknowledgments

This article was written when the first author was an intern at Ascend Communications Inc. The support of Ascend Communications Inc. is gratefully acknowledged. The authors sincerely thank the anonymous reviewers for their insightful comments.

References

- [1] R. Comerford, "State of the Internet: Roundtable 4.0," *IEEE Spectrum*, Oct. 1998.
- [2] D. Ferrari and L. Delgrosso, "Charging For QoS," *IEEE/IFIP IWQOS '98* keynote paper, Napa, CA, May 1998.
- [3] P. Ferguson and G. Huston, *Quality of Service*, Wiley, 1998.
- [4] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," Internet RFC 1633, June 1994.
- [5] R. Jain, "Myths about Congestion Management in High Speed Networks," *Internetworking: Res. and Exp.*, vol. 3, 1992, pp. 101-13.
- [6] Cisco's 12000 Series, <http://www.cisco.com/warp/public/733/12000>
- [7] Ascend's GRF routers, <http://www.ascend.com/300.html>
- [8] Bay Networks' Accelar Routing Switches, <http://business5.baynetworks.com/MainBody.asp>
- [9] 3Com's switches, <http://www.3com.com/products/switches.html>
- [10] Juniper's M40 Series, <http://www.juniper.net/products/default.htm>
- [11] Lucent's PacketStar 6400 Series, <http://www.lucent.com:80/dns/products/ps6400.html>
- [12] R. Braden et al., "Resource ReSerVation Protocol (RSVP) — Version 1 Functional Specification," RFC 2205, Sept. 1997.
- [13] S. Blake et al., "An Architecture for Differentiated Services," RFC 2475, Dec. 1998.
- [14] Y. Bernet et al., "A Framework for Differentiated Services," Internet draft, draft-ietf-diffserv-framework-00.txt, May 1998.
- [15] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," Internet draft, draft-ietf-mpls-arch-01.txt, Mar. 1998.
- [16] D. Awduche et al., "Requirements for Traffic Engineering over MPLS," Internet draft, draft-ietf-mpls-traffic-eng-00.txt, Oct. 1998.
- [17] E. Crawley et al., "A Framework for QoS-based Routing in the Internet," RFC 2386, Aug. 1998.
- [18] S. Shenker, C. Partridge and R. Guerin, "Specification of Guaranteed Quality of Service," RFC 2212, Sept. 1997.
- [19] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," RFC 2211, Sept. 1997.
- [20] R. Guerin, S. Blake, and S. Herzog, "Aggregating RSVP-based QoS Requests," Internet draft, draft-guerin-aggreg-RSVP-00.txt, Nov. 1997.
- [21] T. Li and Y. Rekhter, "Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)," RFC 2430, Oct. 1998.
- [22] D. Awduche et al., "Extension to RSVP for Traffic Engineering," Internet draft, draft-swallow-mpls-RSVP-traffic-eng-00.txt, Aug. 1998.
- [23] D. Ferrari and D. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE JSAC*, vol. 8, no. 3, Apr. 1990, pp. 368-79.
- [24] A. Banerjee and B. Mah, "The Real-Time Channel Administration Protocol," *Proc. 2nd Int'l. Wksp. Network and OS Support for Digital Audio and Video*, Heidelberg: Springer-Verlag, pp. 160-70.
- [25] D. Clark, "The Design Philosophy of the DARPA Internet Protocol," *Proc. ACM SIGCOMM '88*, Aug. 1988.
- [26] J. Postel, "Service Mappings," RFC 795, Sept. 1981.
- [27] K. Nichols et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, Dec. 1998.
- [28] K. Nichols, V. Jacobson, and L. Zhang, "A Two-Bit Differentiated Services Architecture for the Internet," Internet draft, draft-nichols-diff-svc-arch-00.txt, Nov. 1997.
- [29] J. Heinanen et al., "Assured Forwarding PHB Group," Internet draft, draft-ietf-diffserv-af-03.txt, Nov. 1998.
- [30] Y. Bernet et al., "A Framework for use of RSVP with Diff-serv Networks," Internet draft, draft-ietf-diffserv-rsvp-00.txt, June 1998.

- [31] D. Clark and J. Wroclawski, "An Approach to Service Allocation in the Internet," Internet draft, draft-clark-different-svc-alloc-00.txt, July 1997.
- [32] B. Braden et al., "Recommendation on Queue Management and Congestion Avoidance in the Internet," RFC 2309, Apr. 1998.
- [33] T. Li, "CPE based VPNs using MPLS," Internet draft, draft-li-MPLS-vpn-00.txt, Oct. 1998.
- [34] H. Zhang, "Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks," *Proc. IEEE*, vol. 83, no. 10, Oct. 1995.
- [35] W. Yeang, T. Howes, and S. Kille, "Lightweight Directory Access Protocol," RFC 1777, Mar. 1995.
- [36] P. Vaananen and R. Ravikanth, "Framework for Traffic Management in MPLS Networks," Internet draft, draft-vaananen-mpls-tm-framework-00.txt, Mar. 1998.
- [37] L. Andersson et al., "Label Distribution Protocol," Internet draft, draft-ietf-mpls-ldp-02.txt, Nov. 1998.
- [38] M. Waldvoege et al., "Scalable High Speed IP Routing Lookups," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [39] S. Nilsson and G. Karlsson, "Fast Address Lookup for Internet Routers," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [40] J. Moy, "OSPF Version 2," RFC 2178, Apr. 1998.
- [41] C. Villamizar and T. Li, "IS-IS Optimized Multipath (IS-IS OMP)," Internet draft, draft-villamizar-isis-omp-00.txt, Oct. 1998.
- [42] R. Guerin et al., "QoS Routing Mechanisms and OSPF extensions," Internet draft, draft-guerin-QoS-routing-ospf-03.txt, Jan. 1998.
- [43] Z. Zhang et al., "QoS Extensions to OSPF," Internet draft, draft-zhang-qps-ospf-01, Sept. 1997.
- [44] A. Orda, "Routing with End-to-End QoS Guarantees in Broadband Networks," Tech. rep., Technion, Israel.
- [45] Q. Ma, "QoS Routing in the Integrated Services networks," Ph.D. thesis, CMU-CS-98-138, Jan. 1998.
- [46] Y. Goto, M. Ohta, and K. Araki, "Path QoS Collection for Stable Hop-by-hop QoS Routing," *Proc. INET '97*, Kuala Lumpur, Malaysia, June 1997.
- [47] Z. Wang and J. Crowcroft, "Quality of Service Routing for Supporting Multimedia Applications," *IEEE JSAC*, Sept. 1996.
- [48] A. Parekh, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks," Ph.D. thesis, MIT, Feb. 1992.
- [49] C. Partridge, *Gigabit Networking*, Addison-Wesley, 1994.
- [50] G. Apostolopoulos et al., "Quality of Service Based Routing: A Performance Perspective," *Proc. ACM SIGCOMM '98*, Vancouver, Canada, Aug. 1998, pp. 17-28.
- [51] F. Kelly, "Notes on Effective Bandwidths," in *Stochastic Networks: Theory and Applications*, F. Kelly, S. Zachary, and I. B. Ziedins, Eds., Oxford Univ. Press, 1996, pp. 141-68.
- [52] F. Kelly, "Modeling Communication Networks, Present and Future," *Phil. Trans. Royal Soc. A354*, 1996, pp. 437-63.
- [53] G. R. Ash et al., "Real-Time Network Routing in a Dynamic Class-of-Service Network," *Proc. ITC 13*, Copenhagen, Denmark, June 1991.
- [54] ATM Forum PNNI subworking group, "Private Network-Network Interface Spec. v1.0 (PNNI 1.0)," atpnni-0055.00, Mar. 1996.
- [55] Juniper Whitepaper, "Optimizing Routing Software for Reliable Internet Growth," <http://www.juniper.net/leadingedge/whitepapers/optimizing-routing-sw.fm.html>.

Additional Reading

- [1] K. Nichols et al., "Differentiated Services Operational Model and Definitions," Internet draft, draft-nichols-dsopdef-00.txt, Feb. 1998.
- [2] Z. Wang, "Routing and Congestion Control in Datagram Networks," Ph.D. thesis, Dept. of Comp. Sci., Univ. College London, Jan. 1992.

Biographies

XIPENG XIAO (xiaoxipe@cse.msu.edu) is a Ph.D. candidate of the Dept. of Computer Science and Engineering, Michigan State University. He got his B.S. and M.S. degrees in computer science in 1992 and 1995, respectively, from Zhejiang University, China. His research interests include gigabit routing and router architecture, QoS, QoS, and traffic engineering. He won the First Prize in the ACM 1998 Graduate Student Research Poster Contest with the project "The Design of a Scalable Gigabit IP Router." He worked at the Multi-Gigabit Routing Division of Ascend Communications from May 1998 to Dec. 1998. He is now a senior network traffic analysis engineer at Frontier GlobalCenter Inc.

LIONEL M. NI (ni@cse.msu.edu) earned his B.S. degree in electrical engineering from National Taiwan University in 1973 and his Ph.D. degree in electrical engineering from Purdue University, West Lafayette, Indiana, in 1980. He is a Fellow of IEEE and a professor in the Computer Science and Engineering Department at Michigan State University. He has published over 160 technical articles in refereed journals and proceedings in the areas of high-speed networks and high-performance computer systems. He has served as an editor of many professional journals. He was program director of the NSF Microelectronic Systems Architecture Program from 1995 to 1996. He has served as program chair or general chair of many professional conferences. He has received a number of awards for authoring outstanding papers and also won the Michigan State University Distinguished Faculty Award in 1994.



Education and research in Telecommunications

Sophia-Antipolis (10 km from Nice)

Founded in 1991 by the French Institute of Telecommunications (ENST-TELECOM PARIS), and the Swiss Federal Institute of Technology Lausanne (EPFL), Eurécom is a university-level education and research institute in Telecommunications: 60 academic staff and 100 students (130 in year 2000). The board includes important organizations and research centres and major companies involved in communication and information technology, integrated in the "Telecom Valley" concept. In 1994, Eurécom received from IEEE the Major Educational Innovation Award. The fast evolution of the telecommunications science demands a unique adaptive curriculum. Eurécom is proud to have its alumni already present in the most successful Universities and companies worldwide. It is now a far reaching example of tomorrow's higher education institutes.

The board of Eurécom seeks a:

General Director

Main Tasks:

- Manage and develop a scientific and pedagogical policy and planning
- Assure and develop the promotion and acceptance with the academic and industrial partners
- Manage a research and higher education institute of 75 staff (academic and administrative).

Qualifications:

- Personality with a professional vision, solid, reliable, entrepreneur and talented communicator
- Senior academic or research career in a related field
- Experience in the management of research teams and projects
- French and English

Please send your CV to Nicole Haefliger
K/F Selection
Limmatpark
8957 Spreitenbach
Switzerland

For increased speed,
you may fax ++41 56 401 65 07
or e-mail to us at:
nicole.haefliger@kornferry.com
Your application will be treated confidentially.

K/F SELECTION – A DIVISION OF KORN/FERRY INTERNATIONAL

APPENDIX 1005B

LIBRARY

CATALOG

PERIODICAL OR NEWSPAPER

IEEE network

Full Record

MARC Tags

Main title

IEEE network.

Published/Created

[New York, N.Y.] : [Institute of Electrical and Electronics Engineers], [©1987]-

Request this Item

 LC Find It

More Information

>

LCCN Permalink	https://lccn.loc.gov/87649810
Publication history	Vol. 1, no. 1 (Jan. 1987)-
Description	volumes : illustrations ; 28 cm
Current frequency	Bimonthly, 1988-
Former frequency	Quarterly, 1987
ISSN	0890-8044
Linking ISSN	0890-8044

LC classification	TK5105.5 .I324
Variant title	Institute of Electrical and Electronics Engineers network
Portion of title	Network
Running title	IEEE network magazine
Serial key title	IEEE network
Abbreviated title	IEEE netw.
Related names	IEEE Communications Society. Institute of Electrical and Electronics Engineers.
LC Subjects	Computer networks--Periodicals.
Other Subjects	Computer networks.
Form/Genre	Periodicals. Periodicals.
Browse by shelf order	TK5105.5
Notes	Title from cover. "A publication of the IEEE Communications Society."
Reproduction no./Source	IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854-4150
Additional formats	Online version: IEEE network 1558-156X (DLC) 2005212472 (OCoLC)60638386
LCCN	87649810
Invalid LCCN	sn 86011872
Dewey class no.	004.6
Postal reg no.	000827 USPS
National bib agency no.	100971621
Other system no.	(OCoLC)ocm14393593

Incorrect ISSN	1055-6877
Additional Links	http://ieeexplore.ieee.org/servlet/opac?punumber=65
Type of material	Periodical or Newspaper
Content type	text
Media type	unmediated
Carrier type	volume
Source of description	Latest issue consulted: Vol. 1, no. 2 (Apr. 1987).

Item Availability



CALL NUMBER	TK5105.5 .I324 Set 1
Request in	Jefferson or Adams Building Reading Rooms
Status	c.1 v. 35, no. 1 2021 Jan/Feb At Bindery - 09-19-2022 c.1 v. 35, no. 2-3 2021 Mar/Apr-May/June At Bindery - 09-19-2022
Older receipts	v.1 (1987), v.2:no.2-v.2:no.5 (1988:Mar.-1988:Sept.), v.3-v.17 (1989-2003), v.18:no.2 (2004:Mar./Apr.), v.18:no.5-v.22:no.2 (2004:Sept./Oct.-2008:Mar./Apr.), v.22:no.5-v.26:no.4 (2008:Sept./Oct.-2012:July/Aug.), v.27:no.1-v.32:no.2 (2013:Jan/Feb.-2018:Mar./Apr.), v.32:no.4-v.35:no.3 (2018:July/Aug.-2021:May/June)

CALL NUMBER	TK5105.5 .I324 Copy 1 Unbound issues
Request in	Newspaper & Current Periodical Reading Room (Madison LM133)
Latest receipts	v. 36, no. 3 (2022 May-Jun) v. 36, no. 2 (2022 Mar-Apr) v. 36, no. 1 (2022 Jan-Feb) v. 35, no. 6 (2021 Nov-Dec) v. 35, no. 5 (2021 Sept-Oct)

v. 35, no. 4 (2021 Jul-Aug)

v. 22, no. 4 (2008 Jul-Aug)

LIBRARY | CATALOG

← 1 of 1 →

PERIODICAL OR NEWSPAPER

IEEE network

[Full Record](#)

[MARC Tags](#)

000 03151cas a2200661 a 4500
001 11331918
005 20220812104428.0
008 861014c19879999nyubr p 0 a0eng c
010 _ |a 87649810 |z sn 86011872
016 7_ |a 100971621 |2 DNLM
022 0_ |a 0890-8044 |l 0890-8044 |y 1055-6877 |2 1
032 _ |a 000827 |b USPS
035 _ |a (OCoLC)ocm14393593
037 _ |b IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854-4150
040 _ |a NSD |b eng |c NSD |d NST |d NSD |d DLC |d NST |d DLC |d NSD |d NYG |d DLC |d NST |d OCL |d HUL |d NST |d OCL |d MYG |d CUS |d OCL |d MYG |d SYS |d IEEEEX |d NLM |d OCLCQ |d ORE |d NSD |d LVB |d CIT |d OCLCQ |d OCLCF |d OCLCO |d IUL |d OCLCA |d OCLCQ |d UMC |d OCLCO |d PUL
042 _ |a nsdp |a pcc
050 00 |a TK5105.5 |b .I324
082 00 |a 004.6 |2 19
210 0_ |a IEEE netw.
222 _0 |a IEEE network
245 00 |a IEEE network.
246 3_ |a Institute of Electrical and Electronics Engineers network
246 30 |a Network
246 17 |a IEEE network magazine
260 _ |a [New York, N.Y.] : |b [Institute of Electrical and Electronics Engineers], |c [©1987]-

300 ___ |a volumes : |b illustrations ; |c 28 cm
 310 ___ |a Bimonthly, |b 1988-
 321 ___ |a Quarterly, |b 1987
 336 ___ |a text |b txt |2 rdacontent
 337 ___ |a unmediated |b n |2 rdamedia
 338 ___ |a volume |b nc |2 rdacarrier
 362 0_ |a Vol. 1, no. 1 (Jan. 1987)-
 500 ___ |a Title from cover.
 550 ___ |a "A publication of the IEEE Communications Society."
 588 ___ |a Latest issue consulted: Vol. 1, no. 2 (Apr. 1987).
 590 ___ |a SERBIB/SERLOC merged record
 650 _0 |a Computer networks |v Periodicals.
 650 _6 |a Réseaux d'ordinateurs |v Périodiques.
 650 _7 |a Computer networks. |2 fast |0 (OCoLC)fst00872297
 655 _7 |a Periodicals. |2 fast |0 (OCoLC)fst01411641
 655 _7 |a Periodicals. |2 lcgft |0 http://id.loc.gov/authorities/genreForms/gf2014026139 |0 (uri)
 http://id.loc.gov/authorities/genreForms/gf2014026139
 710 2_ |a IEEE Communications Society.
 710 2_ |a Institute of Electrical and Electronics Engineers.
 776 08 |i Online version: |t IEEE network |x 1558-156X |w (DLC) 2005212472 |w (OCoLC)60638386
 850 ___ |a AAP |a AU |a ArU |a AzTeS |a CCC |a CLSU |a CPT |a CSt |a CU |a CU-A |a CU-I |a CU-
 RivP |a CU-S |a CU-
 SB |a CaAEU |a CaOLU |a CaOTU |a CaOWA |a CaOWtU |a CaSSU |a CoFS |a CoU |a CtY-
 KS |a DGW |a DLC |a DeU |a FU |a GAT |a ICU |a ICarbS |a IEN |a IaAS |a IaU |a InLP |a KMK |a
 KPT |a KyU |a LU |a MBU |a MH-
 GM |a MMeT |a MdU |a MiD |a MiEM |a MiU |a MnU |a MoKL |a MsU |a N |a NBuU-
 SE |a NIC |a NN |a NNC |a NNE |a NRU |a NSyU |a NcD |a NcRS |a NcU |a NhU |a NjP |a NjR |a
 NmU |a OCU |a OU |a OkU |a OrCS |a PBL |a PPD |a PPiC |a PU |a RPB-
 S |a TU |a TxCM |a TxDaM |a TxHR |a ULA |a UPB |a UU |a ViBibV |a ViU |a ViW |a VtU |a WM |a
 WaU |a WvU |a WyU
 859 41 |u http://ieeexplore.ieee.org/servlet/opac?punumber=65
 853 33 |8 1 |a v. |b no. |u 4 |v r |i (year) |j (month) |w q |x 01
 863 40 |8 1 |a <1> |i <1987> |x provisional
 853 33 |8 2 |a v. |b no. |u 6 |v r |i (year) |j (month) |w b |x 01
 863 40 |8 2 |a <2>- |i <1988>- |x provisional
 890 ___ |a IEEE network. |i 87-649810
 906 ___ |a 7 |b cbc |c serials |d 3 |e ncip |f 19 |g n-oclcserc
 920 ___ |a Keep 1
 923 ___ |d 20020726 |s Group Registration

984 ___ |a srvf
985 ___ |e eserial 200406
991 ___ |b c-GenColl |h TK5105.5 |i .I324 |w SERIALS
992 ___ |b SER |w SERLOC

[Request this Item](#)[LC Find It](#)

Item Availability



CALL NUMBER TK5105.5 .I324
Set 1

Request in Jefferson or Adams Building Reading Rooms

Status c.1 v. 35, no. 1 2021 Jan/Feb At Bindery - 09-19-2022
c.1 v. 35, no. 2-3 2021 Mar/Apr-May/June At Bindery - 09-19-2022

Older receipts v.1 (1987), v.2:no.2-v.2:no.5 (1988:Mar.-1988:Sept.), v.3-v.17 (1989-2003), v.18:no.2 (2004:Mar./Apr.), v.18:no.5-v.22:no.2 (2004:Sept./Oct.-2008:Mar./Apr.), v.22:no.5-v.26:no.4 (2008:Sept./Oct.-2012:July/Aug.), v.27:no.1-v.32:no.2 (2013:Jan/Feb.-2018:Mar./Apr.), v.32:no.4-v.35:no.3 (2018:July/Aug.-2021:May/June)

CALL NUMBER TK5105.5 .I324
Copy 1
Unbound issues

Request in Newspaper & Current Periodical Reading Room (Madison LM133)

Latest receipts v. 36, no. 3 (2022 May-Jun)
v. 36, no. 2 (2022 Mar-Apr)
v. 36, no. 1 (2022 Jan-Feb)
v. 35, no. 6 (2021 Nov-Dec)
v. 35, no. 5 (2021 Sept-Oct)
v. 35, no. 4 (2021 Jul-Aug)
v. 22, no. 4 (2008 Jul-Aug)

APPENDIX 1005C



Try the [Copyright Public Records System \(CPRS\)](#) pilot with enhanced search features and filters.

[Help](#)[Search](#)[History](#)[Titles](#)[Start Over](#)

Public Catalog

Copyright Catalog (1978 to present)

Search Request: Left Anchored Title = IEEE Network

Search Results: Displaying 17 of 40 entries



Labeled View

IEEE Network : the magazine of computer communications.

Type of Work: Serial

Title: IEEE Network : the magazine of computer communications.

Serial Publication Year: 1999

Serial Key Title: IEEE Network

Imprint: New York : IEEE, 1987-.

ISSN: 0890-8044

Description: print material.

Frequency: Quarterly; bimonthly by v. 2, no. 2, Mar. 1988.

Publication History: Vol. 1, no. 1, Jan. 1987-

Copyright Claimant: Institute of Electrical and Electronics Engineers, Inc.

Issues Registered: v. 13, no. 1, Jan-Feb99.; Pub. 1999-03-01; Reg. 1999-05-03; TX0004881643

v. 13, no. 2, Mar-Apr99.; Pub. 1999-03-29; Reg. 1999-05-03; TX0004881643

v. 13, no. 3, May-Jun99.; Pub. 1999-05-24; Reg. 1999-10-07; TX0004989332

v. 13, no. 4, Jul-Aug99.; Pub. 1999-07-20; Reg. 1999-10-07; TX0004989332

v. 13, no. 5, Sep-Oct99.; Pub. 1999-09-20; Reg. 2000-01-14; TX0005044844

v. 13, no. 6, Nov-Dec99.; Pub. 1999-11-22; Reg. 2000-01-14; TX0005044844

Variant title: Masthead title: IEEE Network magazine.

Other Title: Network

Names: [Institute of Electrical and Electronics Engineers, Inc.](#)



Save, Print and Email ([Help Page](#))

Select Download Format Full Record

Enter your email address:

[Help](#) [Search](#) [History](#) [Titles](#) [Start Over](#)

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

APPENDIX 1005D

IEEE NETWORK[®]

March/April 1999 Vol. 13 No. 2

THE MAGAZINE OF GLOBAL INFORMATION EXCHANGE

www.comsoc.org

S
INS
EL24
NTW

Transmission and Distribution of Digital Video: Field Trials and Prototype Implementations



A Publication of the IEEE Communications Society
in cooperation with the IEEE Computer Society and the Internet Society

Ex. 1039

CISCO SYSTEMS, INC. / Page 121 of 245

IEEE NETWORK[®]

THE MAGAZINE OF GLOBAL INFORMATION EXCHANGE

MARCH/APRIL 1999 VOL. 13 NO. 2

..... F E A T U R E S

8

Internet QoS: A Big Picture



A framework for the emerging Internet QoS is presented, including RSVP, differentiated services, MPLS, and constraint-based routing.

Xipeng Xiao and Lionel M. Ni

20

Admission Control for Statistical QoS: Theory and Practice



A broad set of proposed admission control algorithms, key components of future multi-service networks, is described.

Edward W. Knightly and Ness B. Shroff

30

Transmission and Distribution of Digital Video The Integration of Real-Time I/O and Network Support in the Stony Brook Video Server



How a real-time I/O subsystem interacts with the RETHER subsystem, and the performance cost of this integration, are described.

Tzi-cker Chiueh, Chitra Venkatramani, and Michael Vernick

37

AMUSE: Advanced Broadband Services Trials for Residential Users



The AMUSE project within the ACTS program is presented, and different technologies under experimentation are explored.

Marco Di Concetto, Giovanna Pavarani, Corrado Rosa, Fabio Rossi, Siegfried Paul, and Pierluigi Di Martino

46

Video Multicast over the Internet



Techniques that aim to improve the fairness among receivers by sending video in multiple (layered or replicated) streams are considered.

Xue Li, Mostafa H. Ammar, and Sanjoy Paul

61

Multimedia Service Interworking over Heterogeneous Networking Environments



The authors discuss technical issues related to delivery of multimedia services over heterogeneous networking environments.

Sehyeong Cho and Youngmee Shin

70

A CORBA-Based Quality of Service Management Framework for Distributed Multimedia Services and Applications



QoS management system for managing and controlling the QoS of a distributed multimedia system called MAESTRO is described.

James Won-Ki Hong, Jong-Seo Kim, and Jae-Kyu Park



■ THIS ISSUE

includes the second installment of a two-
Special Issue on
Transmission and
Distribution of Digital
Video.

Cover graphics:
Image Bank

Icons explained



The mortarboard indicates a tutorial article.



The light bulb indicates an article that describes a unique idea or a new approach.



The meshing gears indicate an article that describes an implementation.



The graph indicates an article that contains measurements from a real system.



The sigma indicates an article that contains theory or analysis.



The clock indicates a timely article.

..... D E P A R T M E N T S

Editor's Note	2
IEEE 802 Perspectives	4
Guest Editorial	6

Director of Magazines
Hussein Mouftah, Queen's University, Canada

Editor-in-Chief
Jörg Liebeherr, University of Virginia, USA

Senior Technical Editors
Salah Aidarous, NEC America, USA
Thomas M. Chen, Southern Methodist U., USA
Biswanath Mukherjee, Univ. of California-Davis, USA

IEEE Network Technical Editorial Board
Anindo Banerjee, USC/ISI, USA
Jon Bennett, Harvard University, USA
Ernst Biersack, Eurecom, France
Chatschik Bisdikian, IBM Research, USA
Jim Carlo, Texas Instruments, USA
Anna Charny, MIT, USA
Jon Crowcroft, University College London, UK
John Daigle, University of Mississippi, USA
Julio Escobar, SENACYT, Panama
Sally Floyd, Lawrence Berkeley National Laboratory, USA
David Greaves, University of Cambridge, UK
Paolo Giacomazzi, Politecnico di Milano, Italy
Alden Jackson, BBN, USA
Nikhil Jain, Qualcomm, USA
Tim King, British Telecom, USA
Edward Knightly, Rice University, USA
Yi-Bing (Jason) Lin, National Chiao Tung University, Taiwan

Frank Magee, AT&T Research, USA
Allison Mankin, USC/ISI
Ioannis Nikolaidis, University of Alberta, Canada
Peter O'Reilly, GTE Laboratories, USA
Guru Parulkar, Washington University, St. Louis, USA
Krzysztof Pawlikowski, University of Canterbury, New Zealand

Stephen Pink, SICS, Sweden
K. K. Ramakrishnan, AT&T Research, USA
Kazem Sohraby, Bell Labs, USA
Martha Steenstrup, BBN, USA
C.-K. Toh, Georgia Tech, USA
Joe Touch, USC/ISI, USA
Vittorio Trecordi, CEFRIEL, Italy
Bulent Yener, Bell Labs, USA
Martina Zitterbart, TU Braunschweig, Germany

Feature Editors
Jim Carlo, "IEEE 802 Perspectives," Texas Instruments, USA
Allison Mankin, "IETF Perspectives," USC/ISI, USA
John D. Spragins, "New Books and Multimedia," Clemson Univ., USA

IEEE Production Staff
Joseph Milizzo, Publications Manager
Elizabeth Wilber, Managing Editor
Catherine Kemelmacher, Production Editor
Eric Levine, Advertising Sales Manager
Joanne O'Rourke, Staff Assistant
Susan Lange, Layout Editor
Jennifer Porcello, Editorial Assistant

1999 IEEE Communications Society Officers
Thomas J. Plevyak, President
Stephen B. Weinstein, Past President
Roberto de Marca, President-Elect
Douglas N. Zuckerman, VP-Technical Activities
Curtis A. Siller, VP-Membership Services
Ron Horn, VP-Membership Development
Tetsuya Miki, VP-Society Relations
Ross C. Anderson, Treasurer
John M. Howell, Secretary

Board of Governors
The elected officers above plus Members-at-Large:
Class of 1999

Jeffrey H. Derby — Ronald W. Horn
Harold Sobol — Desmond P. Taylor
Class of 2000

Gerhard Fettweis — Paul Hartmann
Michael Kincaid — William R. Robinson
Class of 2001

Laura Cerchio, Leonard Cimini
Roberta Cohen, William Tranter

1999 IEEE Officers
Kenneth R. Laker, President
Bruce A. Eisenstein, President-Elect
Maurice Papo, Secretary
David A. Conner, Treasurer
Joseph Bolognina, Past President
Daniel J. Senese, Executive Director
David G. Daut, Director, Division III



Eliminate Barriers of Time and Space



Jörg Liebeherr

What are the most important research challenges in network research?" It is difficult to think of a question which may heat up a conversation among a group of *IEEE Network* readers more than this one. Indeed, when a group of panelists was asked to respond to this question at a recent workshop, the range of answers covered a wide spectrum.

One of the panelists, however, gave a response which I thought did not exist — a response which is as simple as it is agreeable. The panelist, Hilarie Orman from DAPRA, stated that there is only one grand challenge in network research: to eliminate all barriers of time and space.

I see several implications of Dr. Orman's response, all of which I find encouraging. For one, trying to eliminate the barriers of time and space is a worthy goal and a great contribution to society at large. Furthermore, I find it useful to identify the existence of a common thread which holds together the very diverse and large networking community. Finally, the formulation of the grand challenge implies that a ranking of importance of networking areas is not only futile, but also counter-productive. In other words, comparing the relative importance of "universal wireless access" and "Terabit networks" distracts from the fact that both serve the same greater goal.

This issue contains the second part of a Special Issue on "Transmission and Distribution of Digital Video" by guest editors Bhumip Khasnabish and Anindo Banerjee. The articles in this part focus on testbeds and trials of digital video networks. The guest editors provide an introduction to the articles in their Guest Editorial.

In addition to the guest-edited articles, this issue also contains two feature articles. Both articles shed light on the state-of-the-art of QoS in packet networks. The first article, "Internet QoS: The Big Picture" by Xipeng Xiao and Lionel M. Ni discusses how constrained-based routing, differentiated services, and label switching can constitute the building blocks of a QoS framework for the Internet. The second article, "Admission Control for Statistical QoS: Theory and Practice," reviews the state of the art of our understanding of the multiplexing gain in packet switching networks. I hope you enjoy this issue of *IEEE Network*. I appreciate any feedback sent to me via e-mail to jorg@cs.virginia.edu. Please visit our website at <http://www.comsoc.org/socstr/techcom/ntrwk> for up-to-date information on calls-for-papers of upcoming guest-edited issues.

IEEE NETWORK ISSN 0890-8044 is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc. Headquarters address: IEEE, 3 Park Avenue, 17th Floor, New York, NY 10016-5997, USA; tel: +1-212-705-8900; e-mail: ieee.network@ieee.org. Responsibility for the contents rests upon authors of signed articles and not the IEEE or its members. Unless otherwise specified, the IEEE neither endorses nor sanctions any positions or actions espoused in *IEEE Network*.

ANNUAL SUBSCRIPTION: \$22 in addition to IEEE Communications Society or any other IEEE Society member dues. Non-member prices: \$115. Single copy prices: Member, \$10; Non-member, \$20.

EDITORIAL CORRESPONDENCE: Address to: Jörg Liebeherr, Editor-in-Chief, *IEEE Network*, IEEE Communications Society, 305 East 47th Street, New York, NY 10017-2394, USA; e-mail: jorg@cs.virginia.edu

COPYRIGHT AND REPRINT PERMISSIONS: Abstracting is permitted with credit to the source. Libraries are permitted to

photocopy beyond the limits of U.S. Copyright law for private use of patrons: those articles that carry a code on the bottom of the first page provided the per copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, USA. For other copying, reprint, or republication permission, write to Director, Publishing Services, at IEEE Headquarters. All rights reserved. Copyright © 1999 by the Institute of Electrical and Electronics Engineers, Inc.

POSTMASTER: Send address changes to *IEEE Network*, IEEE, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA. Printed in USA. Periodical-class postage paid at New York, NY and at additional mailing offices.

SUBSCRIPTIONS, orders, address changes should be sent to IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA. Tel. +1-908-981-0060. GST Reg# 125634188.

ADVERTISING: Advertising is accepted at the discretion of the publisher. Address correspondence to *IEEE Network*, 305 East 47th Street, New York, NY 10017-2394, USA.



Internet QoS: A Big Picture

Xipeng Xiao and Lionel M. Ni, Michigan State University

Abstract

In this article we present a framework for the emerging Internet quality of service (QoS). All the important components of this framework — integrated services, RSVP, differentiated services, multiprotocol label switching (MPLS), and constraint-based routing — are covered. We describe what integrated services and differentiated services are, how they can be implemented, and the problems they have. We then describe why MPLS and constraint-based routing have been introduced into this framework, how they differ from and relate to each other, and where they fit into the differentiated services architecture. Two likely service architectures are presented, and the end-to-end service deliveries in these two architectures are illustrated. We also compare ATM networks to router networks with differentiated services and MPLS. Putting all these together, we give the readers a grasp of the big picture of the emerging Internet QoS.

Today's Internet only provides *best-effort service*. Traffic is processed as quickly as possible, but there is no guarantee as to timeliness or actual delivery. With the rapid transformation of the Internet into a commercial infrastructure, demands for service quality have rapidly developed [1-4].

It is becoming apparent that several service classes will likely be demanded. One service class will provide predictable Internet services for companies that do business on the Web. Such companies will be willing to pay a certain price to make their services reliable and give their users a fast feel of their Web sites. This service class may contain a single service, or it may contain *gold service*, *silver service*, and *bronze service*, with decreasing quality. Another service class will provide low-delay and low-jitter services to applications such as Internet telephony and videoconferencing. Companies will be willing to pay a premium price to run a high-quality videoconference to save travel time and cost. Finally, best-effort service will remain for those customers who only need connectivity.

Whether mechanisms are even needed to provide quality of service (QoS) is a hotly debated issue. One opinion is that fibers and wavelength-division multiplexing (WDM) will make bandwidth so abundant and cheap that QoS will be automatically delivered. The other opinion is that no matter how much bandwidth the networks can provide, new applications will be invented to consume them; therefore, mechanisms will still be needed to provide QoS. This argument is beyond the scope of this article [5]. Here we simply note that even if bandwidth will eventually become abundant and cheap, it is not going to happen soon. For now, some simple mechanisms are definitely needed in order to provide QoS on the Internet. Our view is supported by the fact that all the major router/switch vendors now provide some QoS mechanisms in their high-end products [6-11].

The Internet Engineering Task Force (IETF) has proposed

many service models and mechanisms to meet the demand for QoS. Notably among them are the integrated services/Resource Reservation Protocol (RSVP) model [4, 12], the differentiated services (DS) model [13, 14], multiprotocol label switching (MPLS) [15], traffic engineering [16], and constraint-based routing [17].

The integrated services model is characterized by resource reservation. For real-time applications, before data are transmitted, the applications must first set up paths and reserve resources. RSVP is a signaling protocol for setting up paths and reserving resources. In differentiated services, packets are marked differently to create several packet classes. Packets in different classes receive different services. MPLS is a forwarding scheme. Packets are assigned labels at the ingress of an MPLS-capable domain. Subsequent classification, forwarding, and services for the packets are based on the labels. Traffic engineering is the process of arranging how traffic flows through the network. Constraint-based routing is to find routes that are subject to some constraints, such as bandwidth or delay requirement.

Although there are many papers on each of integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing, to the best of the authors' knowledge, they are never discussed together in a single paper. As a result, it is difficult for readers to understand the relationships among them and to grasp the big picture of the QoS framework.

In this article we give an introduction to integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing. We describe how they differ from, relate to, and work with each other to deliver QoS on the Internet. Through this, we intend to present to readers a clear overview of Internet QoS.

The organization of the rest of the article is as follows. In the next two sections, we describe integrated services, RSVP,

Flow	A stream of packets with the same source IP address, source port number, destination IP address, destination port number, and protocol ID.
Service level agreement (SLA)	A service contract between a customer and a service provider that specifies the forwarding service a customer should receive. A customer may be a user organization or another provider domain (upstream domain).
Traffic profile	A description of the properties of a traffic stream, such as rate and burst size.
Differentiated services (DS) field	The field in which the differentiated services class is encoded. It is the Type of Service (TOS) octet in the IPv4 header or the traffic class octet in the IPv6 header.
Per-hop behavior (PHB)	The externally observable behavior of a packet at a DS-compliant router.
Mechanism	A specific algorithm or operation (e.g., queuing discipline) that is implemented in a router to realize a set of one or more per-hop behaviors.
Admission control	The decision process of whether to accept a request for resources (link bandwidth plus buffer space).
Classification	The process of sorting packets based on the content of packet headers according to defined rules.
Behavior aggregate (BA) classification	The process of sorting packets based only on the contents of the DS field.
Multifield (MF) classification	The process of classifying packets based on the content of multiple fields such as source address, destination address, TOS byte, protocol ID, source port number, and destination port number.
Marking	The process of setting the DS field in a packet.
Policing	The process of handling out-of-profile traffic (e.g., discarding excess packets).
Shaping	The process of delaying packets within a traffic stream to cause it to conform to some defined traffic profile.
Scheduling	The process of deciding which packet to send first in a system of multiple queues.
Queue management	Controlling the length of packet queues by dropping packets when necessary or appropriate.
Traffic trunk	An aggregation of flows with the same service class that can be put into an MPLS label-switched path.

■ Table 1. Frequently used terminologies.

and differentiated services, their characteristics, mechanisms, and problems. A likely differentiated services architecture and the complete process for delivering end-to-end services in this architecture are also presented. MPLS and a service architecture based on MPLS are then described. We then describe traffic engineering and constraint-based routing. ATM networks and router networks are compared. Finally, we summarize the article in the last section.

The frequently used terminologies in this article are defined in Table 1.


Integrated Services and RSVP

The integrated services model [4] proposes two service classes in addition to best-effort service. They are:

- *Guaranteed service* [18] for applications requiring fixed delay bound
- *Controlled-load service* [19] for applications requiring reliable and enhanced best-effort service

The philosophy of this model is that "there is an inescapable requirement for routers to be able to reserve resources in order to provide special QoS for specific user packet streams, or flows. This in turn requires flow-specific state in the routers" [4].

RSVP was invented as a signaling protocol for applications to reserve resources [12]. The signaling process is illustrated in Fig. 1. The sender sends a PATH message to the receiver specifying the characteristics of the traffic. Every intermediate



■ Figure 1. RSVP signaling.

router along the path forwards the PATH message to the next hop determined by the routing protocol. Upon receiving a PATH message, the receiver responds with a RESV message to request resources for the flow. Every intermediate router along the path can reject or accept the request of the RESV message. If the request is rejected, the router will send an error message to the receiver, and the signaling process will terminate. If the request is accepted, link bandwidth and buffer space are allocated for the flow, and the related flow state information will be installed in the router.

Recently, RSVP has been modified and extended in several ways to reserve resources for aggregation of flows, to set up explicit routes (ERs) with QoS requirements, and to do some other signaling tasks [20–22]. Whether or not this is appropriate is a hotly debated issue in the IETF and is beyond the scope of this article.

Integrated services is implemented by four components: the

signaling protocol (e.g., RSVP), the admission control routine, the classifier, and the packet scheduler. Applications requiring guaranteed or controlled-load service must set up the paths and reserve resources before transmitting their data. The admission control routines will decide whether a request for resources can be granted. When a router receives a packet, the classifier will perform a *multifield* (MF) classification and put the packet in a specific queue based on the classification result. The packet scheduler will then schedule the packet accordingly to meet its QoS requirements.

The integrated services/RSVP architecture is influenced by the work of Ferrari *et al.* [23, 24]. It represents a fundamental change to the current Internet architecture, which is founded on the concept that all flow-related state information should be in the end systems [25].

The problems with the integrated services architecture are:

- The amount of state information increases proportionally with the number of flows. This places a huge storage and processing overhead on the routers. Therefore, this architecture does not scale well in the Internet core.
- The requirement on routers is high. All routers must have RSVP, admission control, MF classification, and packet scheduling.
- Ubiquitous deployment is required for guaranteed service. Incremental deployment of controlled-load service is possible by deploying controlled-load service and RSVP functionality at the bottleneck nodes of a domain and tunneling the RSVP messages over other parts of the domain.

Differentiated Services

Because of the difficulty in implementing and deploying integrated services and RSVP, differentiated services (DS) are introduced.

An Introduction to Differentiated Services

IPv4 header contains a Type of Service (TOS) byte. Its meaning was previously defined in [26]. Applications can set three bits in the TOS byte to indicate the need for low-delay, high-throughput, or low-loss-rate service. However, choices are limited. Differentiated services defines the layout of the TOS byte (*DS field*) and a base set of packet forwarding treatments (per-hop behaviors, or PHBs) [27]. By marking the DS fields of packets differently and handling packets based on their DS fields, several differentiated service classes can be created. Therefore, differentiated services is essentially a relative-priority scheme.

In order for a customer to receive differentiated services from its Internet service provider (ISP), it must have a service level agreement (SLA) with its ISP. An SLA basically specifies the service classes supported and the amount of traffic allowed in each class. An SLA can be static or dynamic. *Static SLAs* are negotiated on a regular (e.g., monthly or yearly) basis. Customers with *dynamic SLAs* must use a signaling protocol (e.g., RSVP) to request services on demand.

Customers can mark DS fields of individual packets to indicate the desired service or have them marked by the leaf router based on MF classification.

At the ingress of the ISP networks, packets are classified, policed, and possibly shaped. The classification, policing, and shaping rules used at the ingress routers are derived from the SLAs. The amount of buffering space needed for these operations is also derived from the SLAs. When a packet enters one domain from another domain, its DS field may be remarked as determined by the SLA between the two domains.

Using the classification, policing, shaping, and scheduling mechanisms, many services can be provided, such as:

- *Premium service* for applications requiring low-delay and low-jitter service
- *Assured service* for applications requiring better reliability than best-effort service
- *Olympic service*, which provides three tiers of services: *gold*, *silver*, and *bronze*, with decreasing quality [28, 29]

Note that differentiated services only defines DS fields and PHBs. It is ISPs' responsibility to decide which services to provide.

Differentiated services is significantly different from integrated services. First, there are only a limited number of service classes indicated by the DS field. Since service is allocated in the granularity of a class, the amount of state information is proportional to the number of classes rather than the number of flows. Differentiated services is therefore more scalable. Second, sophisticated classification, marking, policing, and shaping operations are only needed at the boundary of the networks. ISP core routers need only to have *behavior aggregate* (BA) classification. Therefore, it is easier to implement and deploy differentiated services.

There is another reason the second feature is desirable for ISPs. ISP networks usually consist of boundary routers connected to customers and core routers/switches interconnecting the boundary routers. Core routers must forward packets very quickly, and therefore must be simple. Boundary routers need not forward packets very quickly because customer links are relatively slow. Therefore, they can spend more time on sophisticated classification, policing and shaping [3]. Boundary routers at the *network access points* (NAPs) are exceptions. They must forward packets very quickly and do sophisticated classification, policing, and shaping. Therefore, they must be well equipped.

In the differentiated services model, incremental deployment is possible for assured service. DS-incapable routers simply ignore the DS fields of the packets and give the assured service packets best-effort service. Since assured service packets are less likely to be dropped by DS-capable routers, the overall performance of assured service traffic will be better than that of best-effort traffic.

An End-to-End Service Architecture

In this section a service architecture for differentiated services is presented. This architecture provides assured service and premium service in addition to best-effort service. It is mainly based on the architecture proposed in [28]. Other possible service architectures also exist [30].

Assured Service — Assured service is intended for customers that need reliable services from their service providers, even in times of network congestion. Customers will have SLAs with their ISPs. The SLAs will specify the amount of bandwidth allocated for the customers. Customers are responsible for deciding how their applications share that amount of bandwidth. One possible service allocation process is described later. SLAs for assured service are usually static, meaning that customers can start data transmission whenever they want without signaling their ISPs.

Assured service can be implemented as follows. First, classification and policing are done at the ingress routers of the ISP networks. If the assured service traffic does not exceed the bit rate specified by the SLA, they are considered *in profile*; otherwise, the excess packets are considered *out of profile*. Second, all packets, *in* and *out*, are put into an *assured queue* (AQ) to avoid out of order delivery. Third, the queue is managed by a queue management scheme called *random early detection* (RED) with *In and Out* — RIO [31].

RED [32] is a queue management scheme that drops packets randomly. This will trigger the TCP flow control mechanism.

nisms at different end hosts to reduce send rates at different time. By doing so, RED can prevent the queue at the routers from overflowing, and therefore avoid the tail-drop behavior (dropping all subsequent packets when a queue overflows). Tail-drop triggers multiple TCP flows to decrease and later increase their rates simultaneously. It causes network utilization to oscillate and can hurt performance significantly. RED has proven to be useful and has been widely deployed.

RIO is a more advanced RED scheme. It basically maintains two RED algorithms, one for in packets and one for out packets. There are two thresholds for each queue. When the queue size is below the first threshold, no packets are dropped. When the queue size is between the two thresholds, only out packets are randomly dropped. When the queue size exceeds the second threshold, indicating possible network congestion, both in and out packets are randomly dropped, but out packets are dropped more aggressively. In addition to breaking the TCP flow-control synchronization, RIO prevents, to some extent, greedy flows from hurting the performance of other flows by dropping the out packets more aggressively.

Because in packets have low loss even in cases of congestion, the customers will perceive a predictable service from the network if they keep traffic conformant. When there is no congestion, out packets will also be delivered. The networks are thus better utilized.


Best-effort traffic can be treated differently, or identically, from assured service out traffic. In this article we assume that they are treated identically. Therefore, conceptually, we can consider that there is an *A-bit* in the DS field. The A-bits of assured service in packets are set to 1, while the A-bits of assured service out packets and best-effort packets are reset to 0.

Premium Service — Premium service provides low-delay and low-jitter service for customers that generate fixed peak bit rate traffic. Each customer will have an SLA with its ISP. The SLA specifies a desired peak bit-rate for a specific flow or an aggregation of flows. The customer is responsible for not exceeding the peak rate; otherwise, excess traffic will be dropped. The ISP guarantees that the contracted bandwidth will be available when traffic is sent. Premium service is suitable for Internet telephony, videoconferencing, or for creating virtual lease lines for virtual private networks (VPNs) [33].

Because premium service is more expensive than assured service, it is desirable for ISPs to support both static and dynamic SLAs. Dynamic SLAs allow customers to request premium service on demand without subscribing to it. Admission control is needed for dynamic SLAs.

Premium service can be implemented as follows. At the customer side, some entity will decide which application flow can use premium service. The leaf routers directly connected to the senders will do MF classifications and shape the traffic. Conceptually, we can consider that there is a *P-bit* in the DS field. If the P-bit of a packet is set, this packet belongs to the premium class; otherwise, the packet belongs to the assured service or best-effort class. After the shaping, the P-bits of all packets are set for the flow that is allowed to use premium service. The exit routers of the customer domain may need to reshape the traffic to make sure that the traffic does not exceed the peak rate specified by the SLA. At the provider side, the ingress routers will police the traffic. Excess traffic is dropped. All packets with the P-bit set enter a *premium queue* (PQ). Packets in the PQ will be sent before packets in the AQ.

First, by admission control the amount of premium traffic can be limited to a small percentage, say 10 percent, of the bandwidth of input links. Second, excess packets are dropped at the ingress routers of the networks. Nonconformant flows cannot impact the performance of conformant flows. Third,



■ Figure 2. Uneven distribution of premium traffic in an ISP. The shaded area is the core of the ISP.

premium packets are forwarded before packets of other classes; they can potentially use 100 percent of the bandwidth of the output links. Since most links are full-duplex, the bandwidth of the input links equals the bandwidth of the output links. Therefore, if premium traffic is distributed evenly among the links, these three factors should guarantee that the service rate of the PQ is much higher than the arrival rate. Therefore, arriving premium packets should find the PQ empty or very short most of the time. The delay or jitter experienced by premium packets should be very low. However, premium service provides no quantified guarantee on the delay or jitter bound.

However, uneven distribution of premium traffic may cause a problem for premium service. In ISP networks, aggregation of traffic from the boundary routers to a core router (e.g., CR1 in Fig. 2) is inevitable; but this is not a problem because the output link is much faster than the input links. However, aggregation of premium traffic in the core at CR4 may invalidate the assumption that the arrival rate of premium traffic is far below the service rate. Differentiated services alone cannot solve this problem. Traffic engineering/constraint-based routing must be used to avoid congestion caused by such uneven traffic distribution.

By limiting the total amount of bandwidth requested by premium traffic, the network administrators can guarantee that premium traffic will not starve the assured and best-effort traffic. Another scheme is to use Weighted Fair Queuing (WFQ) [34] between the PQ and the AQ.

Service Allocation in Customer Domains — Given an SLA, a customer domain should decide how its hosts share the services specified by the SLA. This process is called *service allocation*.

There are basically two choices:

- Each host makes its own decision as to which service to use.
- A resource controller called the *bandwidth broker* (BB) [28] makes decision for all hosts.

A BB can be a host, a router, or a software process on an exit router. It is configured with the organizational policies and manages the resources of a domain. A domain may also have backup BBs. Since all hosts must cooperate to share a limited amount of resources specified by the SLA, it is technically better to have a BB to allocate resources.

At the initial deployment stage, hosts need no DS mechanism. They simply send their packets unmarked. The exit routers marked them before sending them out to the ISPs. The packets are treated as best-effort traffic inside the customer domain. In later deployment stages, hosts may have some signaling or marking mechanisms. Before a host starts sending packets, it may decide the service class for the packets by itself, or it may consult a BB for a service class. The host may mark the packets by itself or send the packets unmarked. If the host

sends the packets unmarked, the BB must use some protocols, such as RSVP or Lightweight Directory Access Protocol (LDAP) [35], to set the classification, marking, and shaping rules at the leaf router directly connected to the sender so that the leaf router knows how to mark the sender's packets.

If the SLA between a customer and its ISP is dynamic, the BB in the customer domain must also use some signaling protocol to request resources on demand from its ISP. From now on, we assume that RSVP is used as the signaling protocol.

Resource Allocation in ISP Domains — Given the SLAs, ISPs must decide how to configure their boundary routers so that they know how to handle the incoming traffic. This process is called *resource allocation*.

For static SLAs, boundary routers can be manually configured with the classification, policing, and shaping rules. Resources are therefore statically allocated for each customer. Unused resources can be shared by other customers.

For a dynamic SLA, resource allocation is closely related to the signaling process. The BB in the customer domain uses RSVP to request resources from its ISP. At the ISP side, admission control decisions can be made in a distributed manner by the boundary routers, or centrally by a BB. If boundary routers are directly involved in the signaling process, they are configured with the corresponding classification, policing, and shaping rules when they grant a request. If a BB is involved rather than the boundary routers, the BB must configure the boundary routers when it grants a request. Such procedures will be detailed in the next section.

In both cases, the ISP core routers must be shielded from the requests to avoid the scalability problem.

Examples of End-to-End Service Delivery

Example 1: Delivery of Assured Service with a Static SLA — In Fig. 3, host S in corporate network 1 (CN1) wants to use assured service to send data to host D in corporate network 2 (CN2). CN1 has a static SLA with ISP1. The numbers inside the circles are the step numbers in the service delivery process, described below.

- 1 Host S sends a RSVP message to the local BB, CN1-BB, requesting for assured service for its traffic.
- 2 If CN1-BB grants the request, it will configure leaf router LR1 so that LR1 can set the A-bits of the packets of this flow. CN1-BB will then reply to host S; otherwise, an error message is sent to host S.
- 3 Host S sends packets to leaf router LR1.
- 4 If LR1 is configured to mark the traffic, it will set the A-bits of the packets.
- 5 Every router from LR1 (exclusive) to ER1 (inclusive) does a BA classification. Packets with the A-bit set are consid-

ered in, while packets with the A-bit reset are considered out. All packets enter the AQ. RIO is applied on the AQ.

- 6 BR1 polices the traffic. All out traffic remains out. If the in traffic exceeds its bit rate, the excess packets' A-bits will be reset. All packets enter the AQ. RIO is applied on the queue.
- 7 All routers between boundary routers BR1 and BR2 (inclusive) perform BA classifications and apply RIO on their AQs.
- 8 ER2 performs the same operations as BR1.
- 9 The packets are eventually delivered to host D.


Note that:

- If there are multiple ISPs between CN1 and CN2, steps 6 and 7 will be repeated multiple times, once per ISP.
- If CN1 does not have any SLA with ISP1, it can only send traffic as best effort. No matter how the routers in CN1 mark the DS fields of their packets, the A-bits will be reset at BR1.


Example 2: Delivery of Premium Service with a Dynamic SLA — In Fig. 4, host S in CN1 wants to use premium service to send data to host D in CN2. CN1 has a dynamic SLA with ISP1.

Phase 1: Signaling

- 1 Host S sends an RSVP PATH message to the local BB, CN1-BB.
- 2 CN1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to host S. The signaling process ends.
- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to ISP1-BB.
- 4 ISP1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to CN1-BB. Sender S will be notified.
 - If the request is accepted, ISP1-BB sends the PATH message to CN2-BB.
- 5 CN2-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to ISP1-BB. Sender S will be notified.
 - If the request is accepted, CN2-BB will use LDAP or RSVP to request the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to ISP1-BB.
- 6 When ISP1-BB receives the RESV message, it will configure the classification and policing rules on router BR1, and the policing and reshaping rules on router BR2. It will then send the RESV message to CN1-BB.
- 7 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so that if the traffic of the admitted flow is nonconformant, LR1 will shape it. CN1-BB will also set the policing and reshaping



■ Figure 3. The delivery process of assured service with a static SLA.



■ Figure 4. The delivery process of premium service with a dynamic SLA.

rules on router ER1. CN1-BB will then send the RESV message to host S.

8 When host S receives the RESV message, it can start transmitting data.

Note that:

- This signaling process is significantly different from the signaling process in the integrated services/RSVP model. First, it is the sender who requests resources, not the receiver. Second, a request can be rejected when the BB receives the PATH message from the sender. In integrated services/RSVP, a request is rejected only when a router receives the RESV message from the receiver. Third, a BB can aggregate multiple requests and make a single request to the next BB. Fourth, each domain behaves like a single node, represented by the BB. ISP core routers are not involved in this process.
- The state information installed by the BB on the boundary routers is soft state. It must be regularly refreshed, or it will time out.
- If there are multiple ISPs between CN1 and CN 2, repeat steps 4 and 6 once for each ISP.
- If the SLA between CN1 and ISP1 is static, simply skip steps 3–6 in the signaling process.

Phase 2: Data Transmission

9 Host S sends packets to leaf router LR1.

10 Leaf router LR1 performs an MF classification. If the traffic is nonconformant, LR1 will shape it. LR1 will also set the P-bits of the packets. All packets enter the PQ.

11 Each intermediate router between LR1 and ER1 performs a BA classification, puts the packets into the PQ, and sends them out.

12 ER1 performs a BA classification and reshapes the traffic to make sure that the negotiated peak rate is not exceeded. Reshaping is done for the aggregation of all flows heading toward BR1, not for each individual flow.

13 BR1 classifies and polices the premium traffic. Excess premium packets are dropped.

14 Intermediate routers between leaf routers BR1 and BR2 (inclusive) perform BA classifications. BR2 also reshapes the premium traffic.

15 ER2 classifies and polices the premium traffic. Excess premium packets are dropped.

16 The premium packets are delivered to host D.

Requirements on Routers

The requirements on routers to support premium service and assured service are summarized below.

- The leaf routers in customer domains need MF classifications, marking, and shaping.

- The ISP ingress routers need policing and remarking.
- The ISP egress routers optionally need re-shaping.
- All routers need BA classification and two queues with strict priority.
- If dynamic SLAs are supported, each customer domain will need a BB to request service on behalf of the domain and to allocate services inside the domain. Signaling and admission control mechanisms are needed in both customer and ISP domains.

If assured service is to be replaced by olympic service, the AQ must be replaced by three queues: a *gold queue*, a *silver queue*, and a *bronze queue*. WFQ can be used to schedule these queues. The rate parameters of these queues can be manually configured based on experience.

MPLS


The motivation for MPLS is to use a fixed-length label to decide packet handling. MPLS is also a useful tool for traffic engineering [16, 36].

An Introduction to MPLS

MPLS is a forwarding scheme. It evolved from Cisco's *Tag Switching*. In the open systems interconnection (OSI) seven-layer model, it is between layer 2 (L2, link layer) and layer 3 (L3, network layer).

Each MPLS packet has a header. The header contains a 20-bit label, a 3-bit Class of Service (COS) field, a 1-bit label stack indicator, and an 8-bit Time to Live (TTL) field. The MPLS header is encapsulated between the link layer header and the network layer header. An MPLS-capable router, called the *label-switched router* (LSR), examines only the label in forwarding the packet. The network protocol can be IP or others. This is why it is called *multiprotocol* label switching.

MPLS needs a protocol to distribute labels to set up *label-switched paths* (LSPs). Whether a generic *label distribution protocol* (LDP) [37] should be created or RSVP should be extended [22] for this purpose is another hotly debated issue. MPLS labels can also be piggybacked by routing protocols. An LSP is similar to an ATM virtual circuit (VC) and is unidirectional from the sender to the receiver. MPLS LSRs use the protocol to negotiate the semantics of each label, that is, how to handle a packet with a particular label from the peer. LSP setup can be control-driven (i.e., triggered by control traffic such as routing updates) or data-driven (i.e., triggered by the request of a flow or *traffic trunk*). In MPLS, a traffic trunk is an aggregation of flows with the same service class that can be put into an LSP. The LSP between two routers can be the same as the L3 hop-by-hop route, or the sender LSR can specify an *explicit route* (ER) for the LSP. The ability to set up



■ Figure 5. The signaling process of dynamic premium service in an MPLS-based architecture.

ERs is one of the most useful features of MPLS. A forwarding table indexed by labels is constructed as the result of label distribution. Each forwarding table entry specifies how to process packets carrying the indexing label.

Packets are classified and routed at the ingress LSRs of an MPLS-capable domain. MPLS headers are then inserted. When an LSR receives a labeled packet, it will use the label as the index to look up the forwarding table. This is faster than the process of parsing the routing table in search of the longest match done in IP routing [38, 39]. The packet is processed as specified by the forwarding table entry. The incoming label is replaced by the outgoing label, and the packet is switched to the next LSR. This label-switching process is similar to ATM's VCI/VPI processing. Inside an MPLS domain, packet forwarding, classification, and QoS service are determined by the labels and the COS fields. This makes core LSRs simple. Before a packet leaves an MPLS domain, its MPLS label is removed.

MPLS LSPs can be used as tunnels. After LSPs are set up, a packet's path can be completely determined by the label assigned by the ingress LSR. There is no need to enumerate every intermediate router of the tunnel. Compared to other tunneling mechanisms, MPLS is unique in that it can control the complete path of a packet without explicitly specifying the intermediate routers.

In short, MPLS is strategically significant because:

- It provides faster packet classification and forwarding.
- It provides an efficient tunneling mechanism.

These features, particularly the second one, make MPLS useful for traffic engineering [16, 36].

A Service Architecture Based on MPLS

MPLS can be used together with differentiated services to provide QoS. In such an architecture, LSPs are first configured between each ingress-egress pair. For LSP(LSR1 → LSR2) and LSP(LSR2 → LSR1), their intermediate LSRs need not be reciprocal. It is likely that for each ingress-egress pair, a separate LSP is created for each traffic class. In this case, a total number of $C*N*(N-1)/2$ LSPs are needed, where C is the number of traffic classes and N is the number of boundary routers. In order to reduce the number of LSPs, the LSPs from all ingress routers to a single egress router can be merged into a *sink tree*. The total number of sink trees needed is $C*N$. It is also possible to use a single sink tree to transmit packets of different traffic classes, and use the COS bits to differentiate packet classes. In this case, the number of sink trees is reduced to N . In this architecture, as the number of transiting flows increases, the number of flows in each LSP or sink tree also increases. But the number of LSPs or sink trees need not increase. This architecture is therefore scalable.

The operations of the routers are basically the same in this architecture as in the DS-field-based architecture described earlier. There are three differences in the processing of a packet:

- At the ingress of the ISP network, in addition to all the processing described in the DS-field-based architecture, an MPLS header is inserted into the packet.
- Core routers process the packet based on its label and COS field rather than its DS field.
- At the egress, unless interdomain LSPs are configured, the MPLS header is removed.

Note that, with such schemes, the MPLS effect is confined within the ISPs that use MPLS. Whether a particular ISP's architecture is DS-field-based or MPLS-based is transparent to other ISPs. Therefore, the DS-field based and MPLS-based architectures can easily interoperate.

Each customer domain still needs a BB to allocate services, and to request resources on behalf of the customer domain when the SLA is dynamic. But since LSPs are configured within the ISPs, resource requests can easily be hidden from the core routers by tunneling them from the ingress routers to the egress routers. Therefore, BBs may not be needed in MPLS-based ISP networks. Admission control is done in a distributed fashion by the ingress and egress routers.

Without BBs in the ISP networks, the signaling process for dynamic SLAs is slightly different from the one described earlier. It is depicted in Fig. 5 and described below. The data transmission process remains the same except for the three differences noted above.

- 1 Host S sends an RSVP PATH message to its local domain BB CN1-BB.
- 2 CN1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to host S. The signaling process ends.
- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to BR1.
- 4 BR1 decides if there are enough resources to send the traffic to egress router BR2.
 - If no, the request is denied. An error message is sent back to CN1-BB. Sender S will be notified.
 - If yes, ISP1-BB sends the PATH message through an LSP to BR2.
- 5 BR2 sends the PATH message to CN2-BB
- 6 CN2-BB decides if its domain can support the traffic.
 - If no, the request is denied. An error message is sent back to BR2. Sender S will be notified.
 - If yes, the request is accepted. CN2-BB will use LDAP or RSVP to set the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to BR2.
- 7 BR2 configures the reshaping rules for the traffic. It then sends the RESV message through an LSP to BR1.

8 BR1 configures the classification and policing rules for the traffic. It then sends the RESV message to CN1-BB.
 9 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so if the traffic of the admitted flow is nonconformant, LR1 can shape it. CN1-BB will also set the reshaping rules on router ER1. CN1-BB will then pass the RESV message to host S.
 10 Sender S starts transmitting data.
 If there are multiple ISPs between CN1 and CN2, repeat steps 4, 5, 7, and 8 once per ISP.

Traffic Engineering and Constraint-Based Routing

QoS schemes such as integrated services/RSVP and differentiated services essentially provide differentiated degradation of performance for different traffic when traffic load is heavy. When the load is light, integrated services/RSVP, differentiated services, and best-effort service differ little. Then why not avoid congestion at the first place? This is the motivation for traffic engineering.

Traffic Engineering

Network congestion can be caused by lack of network resources or uneven distribution of traffic. In the first case, all routers and links are overloaded, and the only solution is to provide more resources by upgrading the infrastructure. In the second case, some parts of the network are overloaded while other parts are lightly loaded. Uneven traffic distribution can be caused by the current dynamic routing protocols such as RIP, OSPF, and IS-IS, because they always select the shortest paths to forward packets. As a result, routers and links along the shortest path between two nodes may become congested while routers and links along a longer path are idle. The *equal-cost multipath* (ECMP) option of OSPF [40], and recently of IS-IS [41], is useful in distributing load to several shortest paths; but if there is only one shortest path, ECMP does not help. For simple networks, it may be possible for network administrators to manually configure the cost of the links so that traffic can be evenly distributed. For complex ISP networks, this is almost impossible.

Traffic engineering is the process of arranging how traffic flows through the network so that congestion caused by uneven network utilization can be avoided. *Constraint-based routing* is an important tool for making the traffic engineering process automatic.

Avoiding congestion and providing graceful degradation of performance in congestion are complementary. Traffic engineering therefore complements differentiated services.

Constraint-Based Routing

In a sentence, constraint-based routing is used to compute routes that are subject to multiple constraints.

Constraint-based routing evolves from *QoS routing*. Given the QoS request of a flow or an aggregation of flows, QoS routing returns the route that is most likely to be able to meet the QoS requirements. Constraint-based routing extends QoS routing by considering other constraints of the network such as policy. The goals of constraint-based routing are:

- To select routes that can meet certain QoS requirements
- To increase utilization of the network

While determining a route, constraint-based routing considers not only network topology, but also requirements of the flow, resource availability of the links, and possibly other policies specified by the network administrators. Therefore, constraint-based routing may find a longer but lightly loaded path

better than the heavily loaded shortest path. Network traffic is thus distributed more evenly.

In order to do constraint-based routing, routers need to distribute new link state information and to compute routes based on such information.

Distribution of Link State Information — A router needs topology and resource availability information in order to compute QoS routes. Here, resource availability information means link available bandwidth [42]. Buffer space is assumed to be sufficient and is not explicitly considered [42].

One approach to distributing bandwidth information is to extend the link state advertisements of protocols such as OSPF and IS-IS [42, 43]. Because link residual bandwidth is frequently changing, a trade-off must be made between the need for accurate information and the need to avoid frequent flooding of link state advertisements.

To reduce the frequency of link state advertisements, one possible way is to distribute them only when there are topology or significant bandwidth changes (e.g., more than 50 percent or more than 10 Mb/s) [44]. A hold-down timer should always be used to limit the frequency of such advertisements. A recommended timer value is 30 seconds [45]. An approach to limit the flooding scope of such advertisements is described in [46].

Route Computation — The routing table computation algorithms in constraint-based routing and the complexity of such algorithms depend on the metrics chosen for the routes.

Common route metrics in constraint-based routing are monetary cost, hop count, bandwidth, reliability, delay, and jitter. Routing algorithms select routes that optimize one or more of these metrics.

Metrics can be divided into three classes. Let $d(i, j)$ be a metric for link (i, j) . For any path $P = (i, j, k, \dots, l, m)$, metric d is:

- *Additive* if $d(P) = d(i, j) + d(j, k) + \dots + d(l, m)$
- *Multiplicative* if $d(P) = d(i, j) * d(j, k) * \dots * d(l, m)$
- *Concave* if $d(P) = \min\{d(i, j), d(j, k), \dots, d(l, m)\}$

According to this definition, metrics *delay*, *jitter*, *cost*, and *hop count* are additive, *reliability* ($1 - \text{loss rate}$) is multiplicative, and *bandwidth* is concave.

A well-known theorem in constraint-based routing is that computing optimal routes subject to constraints of two or more additive and/or multiplicative metrics is *NP-complete* [47]. That is, algorithms that use two or more of delay, jitter, hop count, and loss probability as metrics, and optimize them simultaneously are NP-complete. The computationally feasible combinations of metrics are bandwidth and one of those metrics.

However, the proof of NP-completeness in [47] is based on the assumptions that all the metrics are independent, and the delay and jitter of a link are known a priori. Although such assumptions may be true in circuit-switched networks, metrics bandwidth, delay, and jitter are not independent in packet networks. As a result, polynomial algorithms for computing routes with hop count, delay, and jitter constraints exist [45]. The complexity of such algorithms is $O(N * E * e)$, where N is the hop count, E is the number of links of the network, and $e \leq E$ is the number of distinct bandwidth values among all links. Nevertheless, the theorem can tell us qualitatively the complexity of a routing algorithm: a complex algorithm in circuit-switched networks is still complex in packet networks, although it may not be NP-complete.

Fortunately, algorithms for finding routes with bandwidth and hop-count constraints are much simpler [42]. Bellman-Ford's (BF) Algorithm or Dijkstra's Algorithm can be used. For example, to find the shortest path between two nodes with bandwidth greater than 1 Mb/s, all the links with residual band-

width less than 1 Mb/s can be pruned first. BF or Dijkstra's Algorithm can then be used to compute the shortest path in the pruned network. The complexity of such algorithms is $O(N^2E)$.

Bandwidth and hop count are more useful constraints than delay and jitter, because:

- Although applications may care about delay and jitter bounds, few applications cannot tolerate occasional violation of such constraints. Therefore, there is no obvious need for routing flows with delay and jitter constraints. Besides, since delay and jitter parameters of a flow can be determined by the allocated bandwidth and the hop count of the route [48, 49], delay and jitter constraints can be mapped to bandwidth and hop-count constraints if needed.
- Many real-time applications will require a certain amount of bandwidth. The bandwidth metric is therefore useful. The hop count metric of a route is important because the more hops a flow traverses, the more resources it consumes. For example, a 1-Mb/s flow that traverses two hops consumes twice as many resources as one that traverses a single hop.

In constraint-based routing, routes can be computed on demand or precomputed for each traffic class. On-demand computations are triggered by the receipt of the QoS request of a flow. In either case, a router will have to compute its routing table more frequently with constraint-based routing than with dynamic routing. This is because, even without topology changes, routing table computation can still be triggered by significant bandwidth changes. Besides, constraint-based routing algorithms are at least as complex as dynamic routing algorithms. Therefore, the computation load of routers with constraint-based routing can be very high.

Common approaches to reduce the computation overhead of constraint-based routing include:

- Using a large-valued timer to reduce computation frequency
- Choosing bandwidth and hop count as constraints
- Using administrative policy to prune unsuitable links before computing the routing table

For example, if a flow has delay requirement, high propagation delay links such as satellite links are pruned before the routing table computation.

Constraint-Based Routing: Pros and Cons

The pros of constraint-based routing are meeting the needs for QoS requirements of flows better, and improved network utilization.

The cons of constraint-based routing are increased communication and computation overhead, increased routing table size, the fact that longer paths may consume more resources, and potential routing instability.

Of the cons, the first was addressed in an earlier section. The rest are addressed in this section.

In constraint-based routing, an essential issue is routing granularity. Routing can be destination-based, source-destination-based, class-based, traffic-trunk-based, or flow-based. Routing with finer granularity is more flexible, and thus more efficient in terms of resource utilization and more stable. However, the computation overhead and storage overhead are also higher.

Routing Table Structure and Size — Routing table structure and size depend directly on routing granularity and route metrics. Logically, we can view the routing table as a two-dimensional array. The number of rows is determined by routing granularity, and the number of columns is determined by route metrics. For example, in destination-based routing with bandwidth and hop count as route metrics, the routing table can be organized as a $K \times H$ array, where K is the number of destinations, and H is the maximum number of hops allowed for any route. The (k, h) th entry of the array contains one or

more h -hop routes for destination k . Each route also has an available bandwidth associated with it [42].

Obviously, the size of a constraint-based routing table can be far larger than the size of a normal routing table for the same network. This introduces significant storage overhead. It may also slow down the routing table lookup.

Approaches to reducing the routing table size in constraint-based routing include:

- Using coarse routing granularity
- Using hop quantization (i.e., dividing all hop-count values into a few classes to reduce the number of columns in the routing table) [44]
- Keeping the routing table only for best-effort traffic, and computing the routes for flows with QoS requests on demand [50]

The third scheme basically trades computation time for smaller storage requirements.

The Trade-off between Resource Conservation and Load Balancing — A constraint-based routing scheme can choose one of the following as the route for a destination:

- The *widest-shortest* path, that is, a path with minimum hop count and, if there are multiple such paths, the one with the largest available bandwidth.
- The *shortest-widest* path, that is, a path with the largest available bandwidth and, if there are multiple such paths, the one with the minimum hop count.
- The *shortest-distance* path. The distance of a k -hop path P is defined as

$$\text{dist}(P) = \sum_{i=1}^k \frac{1}{r_i}, \text{ where } r_i \text{ is the bandwidth of link } i.$$

Using paths other than the shortest paths consume more resources. This is not efficient when the load of the network is heavy. A trade-off must be made between resource conservation and load balancing. The first approach above is basically the same as today's dynamic routing. It emphasizes preserving network resources by choosing the shortest paths. The second approach emphasizes load balancing by choosing the widest paths. The third approach makes a trade-off between the two extremes. It favors shortest paths when network load is heavy and widest paths when network load is medium. Simulations showed that the third approach consistently outperforms the other two approaches for best-effort traffic, regardless of network topology and traffic pattern [45].

Stability — Because constraint-based routing algorithms recompute routing tables more frequently than dynamic routing algorithms do, they can introduce instability.

The stability of networks with constraint-based routing depends heavily on the routing granularity. If routing is done with coarse granularity (e.g., based solely on destination address), when the original route between two nodes becomes congested, all the traffic to that destination is shifted from the original route to an alternate route. This may cause congestion in the alternate route. Traffic may have to be shifted again [46].

The high computation overhead of constraint-based routing may also hurt the stability of the network. When a router is busy computing the routing table, it is slow in reacting to new topology changes.

To improve stability, the timer value for periodic routing table recomputation should be carefully chosen [50]. Constraint-based routing at the granularity of traffic trunk provides a good trade-off between stability and computation overhead [16]. Reducing the computation complexity of the routers also helps to improve stability.

In summary, constraint-based routing must be deployed

with caution. Otherwise, the cost of instability and increased complexity may outweigh the gain.

Constraint-based routing is similar to the dynamic/adaptive routing in telephone networks and ATM networks [51–54]. Many lessons can be learned from those works. Since constraint-based routing is a superset of today's dynamic routing, it is possible that in the future, constraint-based routing may replace dynamic routing, especially in the intradomain case. An emerging intradomain constraint-based routing protocol is QOSPF [42].

The Position of Constraint-Based Routing in the QoS Framework

In this section we describe the relationships between constraint-based routing and other components in the QoS framework.

The Relationship between Constraint-Based Routing and Differentiated Services — Constraint-based routing is to select the optimal routes for flows so that their QoS requirements are most likely to be met. It is not to replace differentiated services, but to help differentiated services be better delivered. Figure 2 shows an example in point.

The Relationship between Constraint-Based Routing and RSVP — RSVP and constraint-based routing are independent but complementary. For a router with dynamic routing, when an RSVP PATH message is received, it will be forwarded to the next hop determined by the dynamic routing protocol. The QoS requirement of the flow and the load of the networks are not considered in selecting the next hop. However, with a router running constraint-based routing, such information is considered. The next hop of the RSVP messages determined by constraint-based routing therefore may be different. In either case, the actual reservation of resources for the flow is done by RSVP. In short, constraint-based routing determines the path for RSVP messages, but does not reserve resources. RSVP reserves resources, but depends on constraint-based or dynamic routing to determine the path.

The Relationship between Constraint-Based Routing and MPLS — Given that MPLS is a forwarding scheme and constraint-based routing is a routing scheme, MPLS and constraint-based routing are, in theory, mutually independent. Constraint-based routing determines the route between two nodes based on resource information and topology information. It is useful with or without MPLS. Given the routes, MPLS uses its label distribution protocol to set up the LSPs. It does not care whether the routes are determined by constraint-based or dynamic routing.

However, when MPLS and constraint-based routing are used together, they make each other more useful. MPLS makes it possible to do constraint-based routing at traffic trunk granularity without introducing MF classification to the core routers. MPLS's per-LSP statistics provide constraint-based routing with precise information about the amount of traffic between every ingress–egress pair. Given such information, constraint-based routing can better compute the routes for setting up LSPs. In combination, MPLS and constraint-based routing provide powerful tools for traffic engineering.

A Comparison of ATM Networks to Router Networks

QoS and some sort of traffic engineering have long been provided by ATM networks. So why introduce differentiated ser-

vices and MPLS into the router networks? To answer this question, we give a brief comparison between ATM networks and router networks.

In an ATM network, QoS can be provided by allocating a certain amount of bandwidth for a specific VC. Traffic engineering is usually done by computing the routes offline and then downloading the configuration statically into the ATM switches on an hourly or daily basis. Per-permanent virtual circuit (PVC) traffic statistics of the current configuration provide accurate traffic information for computing the routes for the next configuration.

The advantages of ATM networks over router networks without differentiated services or MPLS are:

- ATM networks are currently faster in data forwarding.
- Per-PVC traffic statistics are available.
- QoS and some sort of traffic engineering are provided.

The disadvantages of ATM networks are:

- ATM cell header overhead is large.
- Routers must be used at the boundary of the network. With both switches and routers present in the network, two sets of configurations are required: one for routers and the other for switches.

With differentiated services and MPLS, router networks can also provide QoS and traffic engineering. This can be done without a big header overhead and two sets of configurations. Router networks with differentiated services and MPLS therefore provide some advantages over ATM networks [55]; but this is more or less from the perspective of router vendors.

Summary

The big picture of the emerging Internet QoS can be summarized as follows:

- Customers negotiate SLAs with ISPs. The SLAs specify what services the customers will receive. SLAs can be static or dynamic. For static SLAs, customers can transmit data at any time. For dynamic SLAs, customers must use a signaling protocol such as RSVP to request services on demand before transmitting data. The bandwidth brokers in the customer domains decide how applications share the services specified by the SLAs. The DS fields of packets are marked accordingly to indicate the desired services.
- The ingress routers of ISPs are configured with classification, policing, and remarking rules. The egress routers of ISP networks are configured with reshaping rules. Such rules may be configured manually by network administrators or dynamically by some protocol such as LDAP or RSVP. ISPs must implement admission control in order to support dynamic SLAs. Classification, marking, policing, and shaping/reshaping are only done at the boundary routers. Core routers are shielded from the signaling process. They need only implement two queues with strict priority. They process packets based solely on their DS fields.
- With MPLS, LSPs are set up between each ingress–egress pair. At the ISP ingress routers, labels and COS fields are determined from the classification and routing results. MPLS headers are then inserted into the packets. Core routers process packets based on their labels and COS fields only. Labels are removed before packets leave an MPLS domain.
- Constraint-based routing can be used to compute the routes subject to QoS and policy constraints. The goal is to meet the QoS requirements of traffic and to improve utilization of the networks.
- MPLS and constraint-based routing can be used together to control the path of traffic to avoid congestion and improve the utility of the networks.

Application Layer	
Transport Layer	Integrated Service/RSVP, Differentiated Services
Network Layer	Constraint Based Routing
	MPLS
Link Layer	

■ Table 2. The relative positions of the components in the QoS framework.

- The positions of integrated services/RSVP, differentiated services, MPLS, and constraint-based routing in the Internet network model are depicted in Table 2.

Acknowledgments

This article was written when the first author was an intern at Ascend Communications Inc. The support of Ascend Communications Inc. is gratefully acknowledged. The authors sincerely thank the anonymous reviewers for their insightful comments.

References

- [1] R. Comerford, "State of the Internet: Roundtable 4.0," *IEEE Spectrum*, Oct. 1998.
- [2] D. Ferrari and L. Delgrossi, "Charging For QoS," *IEEE/IFIP IWQOS '98* keynote paper, Napa, CA, May 1998.
- [3] P. Ferguson and G. Huston, *Quality of Service*, Wiley, 1998.
- [4] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," Internet RFC 1633, June 1994.
- [5] R. Jain, "Myths about Congestion Management in High Speed Networks," *Internetworking: Res. and Exp.*, vol. 3, 1992, pp. 101-13.
- [6] Cisco's 12000 Series, <http://www.cisco.com/warp/public/733/12000>
- [7] Ascend's GRF routers, <http://www.ascend.com/300.html>
- [8] Bay Networks' Accelar Routing Switches, <http://business5.baynetworks.com/MainBody.asp>
- [9] 3Com's switches, <http://www.3com.com/products/switches.html>
- [10] Juniper's M40 Series, <http://www.juniper.net/products/default.htm>
- [11] Lucent's PacketStar 6400 Series, <http://www.lucent.com:80/dns/products/ps6400.html>
- [12] R. Braden et al., "Resource ReSerVation Protocol (RSVP) — Version 1 Functional Specification," RFC 2205, Sept. 1997.
- [13] S. Blake et al., "An Architecture for Differentiated Services," RFC 2475, Dec. 1998.
- [14] Y. Berner et al., "A Framework for Differentiated Services," Internet draft, draft-ietf-diffserv-framework-00.txt, May 1998.
- [15] E. Rosen, A. Viswanathan, and R. Collon, "Multiprotocol Label Switching Architecture," Internet draft, draft-ietf-mpls-arch-01.txt, Mar. 1998.
- [16] D. Awduche et al., "Requirements for Traffic Engineering over MPLS," Internet draft, draft-ietf-mpls-traffic-eng-00.txt, Oct. 1998.
- [17] E. Crawley et al., "A Framework for QoS-based Routing in the Internet," RFC 2386, Aug. 1998.
- [18] S. Shenker, C. Partridge and R. Guerin, "Specification of Guaranteed Quality of Service," RFC 2212, Sept. 1997.
- [19] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," RFC 2211, Sept. 1997.
- [20] R. Guerin, S. Blake, and S. Herzog, "Aggregating RSVP-based QoS Requests," Internet draft, draft-guerin-aggreg-RSVP-00.txt, Nov. 1997.
- [21] T. Li and Y. Rekhter, "Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)," RFC 2430, Oct. 1998.
- [22] D. Awduche et al., "Extension to RSVP for Traffic Engineering," Internet draft, draft-swallow-mpls-RSVP-traffic-eng-00.txt, Aug. 1998.
- [23] D. Ferrari and D. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE JSAC*, vol. 8, no. 3, Apr. 1990, pp. 368-79.
- [24] A. Banerjee and B. Mah, "The Real-Time Channel Administration Protocol," *Proc. 2nd Int'l. Wksp. Network and OS Support for Digital Audio and Video*, Heidelberg: Springer-Verlag, pp. 160-70.
- [25] D. Clark, "The Design Philosophy of the DARPA Internet Protocol," *Proc. ACM SIGCOMM '88*, Aug. 1988.
- [26] J. Postel, "Service Mappings," RFC 795, Sept. 1981.
- [27] K. Nichols et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, Dec. 1998.
- [28] K. Nichols, V. Jacobson, and L. Zhang, "A Two-Bit Differentiated Services Architecture for the Internet," Internet draft, draft-nichols-diff-svc-arch-00.txt, Nov. 1997.
- [29] J. Heinanen et al., "Assured Forwarding PHB Group," Internet draft, draft-ietf-diffserv-af-03.txt, Nov. 1998.
- [30] Y. Berner et al., "A Framework for use of RSVP with Diff-serv Networks," Internet draft, draft-ietf-diffserv-rsvp-00.txt, June 1998.

- [31] D. Clark and J. Wroclawski, "An Approach to Service Allocation in the Internet," Internet draft, draft-clark-different-svc-alloc-00.txt, July 1997.
- [32] B. Braden et al., "Recommendation on Queue Management and Congestion Avoidance in the Internet," RFC 2309, Apr. 1998.
- [33] T. Li, "CPE based VPNs using MPLS," Internet draft, draft-li-MPLS-vpn-00.txt, Oct. 1998.
- [34] H. Zhang, "Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks," *Proc. IEEE*, vol. 83, no. 10, Oct. 1995.
- [35] W. Yeong, T. Howes, and S. Kille, "Lightweight Directory Access Protocol," RFC 1777, Mar. 1995.
- [36] P. Vaananen and R. Ravikanth, "Framework for Traffic Management in MPLS Networks," Internet draft, draft-vaananen-mpls-tm-framework-00.txt, Mar. 1998.
- [37] L. Andersson et al., "Label Distribution Protocol," Internet draft, draft-ietf-mpls-ldp-02.txt, Nov. 1998.
- [38] M. Waldvogel et al., "Scalable High Speed IP Routing Lookups," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [39] S. Nilsson and G. Karlsson, "Fast Address Lookup for Internet Routers," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [40] J. Moy, "OSPF Version 2," RFC 2178, Apr. 1998.
- [41] C. Villamizar and T. Li, "IS-IS Optimized Multipath (IS-IS OMP)," Internet draft, draft-villamizar-isis-omp-00.txt, Oct. 1998.
- [42] R. Guerin et al., "QoS Routing Mechanisms and OSPF extensions," Internet draft, draft-guerin-QoS-routing-ospf-03.txt, Jan. 1998.
- [43] Z. Zhang et al., "QoS Extensions to OSPF," Internet draft, draft-zhang-qps-ospf-01, Sept. 1997.
- [44] A. Orda, "Routing with End-to-End QoS Guarantees in Broadband Networks," Tech. rep., Technion, Israel.
- [45] Q. Ma, "QoS Routing in the Integrated Services networks," Ph.D. thesis, CMU-CS-98-138, Jan. 1998.
- [46] Y. Goto, M. Ohta, and K. Araki, "Path QoS Collection for Stable Hop-by-hop QoS Routing," *Proc. INET '97*, Kuala Lumpur, Malaysia, June 1997.
- [47] Z. Wang and J. Crowcroft, "Quality of Service Routing for Supporting Multimedia Applications," *IEEE JSAC*, Sept. 1996.
- [48] A. Parekh, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks," Ph.D. thesis, MIT, Feb. 1992.
- [49] C. Partridge, *Gigabit Networking*, Addison-Wesley, 1994.
- [50] G. Apostolopoulos et al., "Quality of Service Based Routing: A Performance Perspective," *Proc. ACM SIGCOMM '98*, Vancouver, Canada, Aug. 1998, pp. 17-28.
- [51] F. Kelly, "Notes on Effective Bandwidths," in *Stochastic Networks: Theory and Applications*, F. Kelly, S. Zachary, and I. B. Ziedins, Eds., Oxford Univ. Press, 1996, pp. 141-68.
- [52] F. Kelly, "Modeling Communication Networks, Present and Future," *Phil. Trans. Royal Soc. A354*, 1996, pp. 437-63.
- [53] G. R. Ash et al., "Real-Time Network Routing in a Dynamic Class-of-Service Network," *Proc. ITC 13*, Copenhagen, Denmark, June 1991.
- [54] ATM Forum PNIN subworking group, "Private Network-Network Interface Spec. v1.0 (PNIN 1.0)," alpnin-0055.00, Mar. 1996.
- [55] Juniper Whitepaper, "Optimizing Routing Software for Reliable Internet Growth," <http://www.juniper.net/leadingedge/whitepapers/optimizing-routing-sw.fm.html>.

Additional Reading

- [1] K. Nichols et al., "Differentiated Services Operational Model and Definitions," Internet draft, draft-nichols-dsopdef-00.txt, Feb. 1998.
- [2] Z. Wang, "Routing and Congestion Control in Datagram Networks," Ph.D. thesis, Dept. of Comp. Sci., Univ. College London, Jan. 1992.

Biographies

XIPENG XIAO (xiaoxipe@cse.msu.edu) is a Ph.D. candidate of the Dept. of Computer Science and Engineering, Michigan State University. He got his B.S. and M.S. degrees in computer science in 1992 and 1995, respectively, from Zhejiang University, China. His research interests include gigabit routing and router architecture, QoS, QoS, and traffic engineering. He won the First Prize in the ACM 1998 Graduate Student Research Poster Contest with the project "The Design of a Scalable Gigabit IP Router." He worked at the Multi-Gigabit Routing Division of Ascend Communications from May 1998 to Dec. 1998. He is now a senior network traffic analysis engineer at Frontier GlobalCenter Inc.

LIONEL M. NI (ni@cse.msu.edu) earned his B.S. degree in electrical engineering from National Taiwan University in 1973 and his Ph.D. degree in electrical engineering from Purdue University, West Lafayette, Indiana, in 1980. He is a Fellow of IEEE and a professor in the Computer Science and Engineering Department at Michigan State University. He has published over 160 technical articles in refereed journals and proceedings in the areas of high-speed networks and high-performance computer systems. He has served as an editor of many professional journals. He was program director of the NSF Microelectronic Systems Architecture Program from 1995 to 1996. He has served as program chair or general chair of many professional conferences. He has received a number of awards for authoring outstanding papers and also won the Michigan State University Distinguished Faculty Award in 1994.

APPENDIX 1005E



IEEE network

AVAILABLE AS

View Online

Available Issues / Request a Copy

[Save to Lists](#)

[Citation](#)

Physical Locations

This item is also held by UW-Milwaukee

Multiple items exist. Use [Available Issues / Request a Copy](#) to see what's on the shelf.

Shared Print Repositories
BTAA Shared Print

Verona Shelving Facility
by request only
S IN5 EL24 NTW

Notes
Print subscription cancelled after: v.16 (2002).

This Location Has
v.13:no.1(1994:01/02)-v.14:no.6(2000:11/12)
v.15:no.1(2001:01/02)-v.15:no.2(2001:03/04)
v.15:no.4(2001:07/08)-v.15:no.6(2001:11/12)
v.16:no.1(2002:01/02)-v.16:no.6(2002:11/12)

Publication Details

OTHER TITLE
IEEE network magazine

FORMAT
Journals, Magazines, Newspapers

LANGUAGE

English

CONTRIBUTORS[IEEE Communications Society](#)[Institute of Electrical and Electronics Engineers](#)**PUBLICATION**

[New York, N.Y.] : Institute of Electrical and Electronics Engineers, c1987-

PUBLICATION DATES

Vol. 1, no. 1 (Jan. 1987)-

PHYSICAL DETAILS

volumes : illustrations ; 28 cm

ISSNS

0890-8044, 1558-156X

OCLC

ocm60638386, ocm14393593

 Notes

Title from cover.

Also issued online.

Subjects[Computer networks -- Periodicals.](#)**Items Related By Call Number**[Shelf Browse](#)

Sixth IEEE International
Symposium on Network Com...
IEEE International Symposium on
Ne...

TK5105.5.I323 2007

GECON 2004 : 2004 1st IEEE
International Workshop...
IEEE International Workshop on
Gri...

TK5105.5.I3235 2004

Gecon 2006 - Proceedings Of
The 3rd International...
Lee, Hing-yan

TK5105.5.I3235 2006

 Information from the Web Library Staff Details**MARC Bibliographic Record**

LEADER 02727cas a2200637 a 4500

001 9911911773602122

005 20150408141440.0

008 861014c19879999nyubr p 0 a0eng c

010 \$a 87649810 \$zsn 86011872

012 \$a3\$b3\$i8708\$k1\$l1\$m1

016 7_ \$a100971621\$2DNLM

019 \$a185064003

022 0_ \$a0890-8044\$l0890-8044\$y1055-6877\$21

032 \$a000827\$bUSPS

035 \$a(OCOLC)ocm14393593

035 \$a(OCOLC)14393593

035 \$9AJK7566UW

035 \$a1191177

035 \$a(WU)1191177-uwmadisondb

035 \$z(OCOLC)185064003

035 \$a(EXLNZ-01UWI_NETWORK)999583751502121

037 \$bIEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854-4150

040 \$aNSD\$beng\$cNSD\$dNST\$dNSD\$dDLC\$dNST\$dDLC\$dNSD\$dNYG\$dDLC\$dNST\$dOCL\$dHUL\$dNST\$dOCL\$dMYG\$dCUS\$d

042 \$ansdp\$apec

049 \$aGZMA

050 00 \$aTK5105.5\$b.I324

082 00 \$a004.6\$219

098 11 \$aS\$aIN5\$aEL24\$aNTW

210 0_ \$aIEEE netw.

222 _0 \$aIEEE network

245 00 \$aIEEE network.

246 3_ \$aInstitute of Electrical and Electronics Engineers network

246 30 \$aNetwork

246 17 \$aIEEE network magazine

260 \$a[New York, N.Y.] :\$bInstitute of Electrical and Electronics Engineers,\$c1987-

```

300      $avolumes :$billustrations ;$c28 cm
310      $aBimonthly,$b1988-
321      $aQuarterly,$b1987
336      $atext$btxt$2rdacontent
337      $aunmediated$bnc$2rdamedia
338      $avolume$bnc$2rdacarrier
362      0_      $aVol. 1, no. 1 (Jan. 1987)-
500      $aTitle from cover.
530      $aAlso issued online.
550      $a"A publication of the IEEE Communications Society."
650      _0      $aComputer networks$vPeriodicals.
650      _6      $aRéseaux d'ordinateurs$vPériodiques.
710      2_      $aIEEE Communications Society.
710      2_      $aInstitute of Electrical and Electronics Engineers.
776      08      $iOnline version:$tIEEE network (Online) $x1558-156X$w(DLC) 2005212472$w(OCOLC) 60638386
850      $AAP$aAU$aArU$aAzTeS$aCCC$aCLSU$aCPT$aCSt$aCU$aCU-A$aCU-I$aCU-RivP$aCU-S$aCU-SB$aCaAEU$aCaO
KS$aDGW$aDLC$aDeU$aFU$aGAT$aICU$aICarbs$aIEN$aIaAS$aIaU$aInLP$aKMK$aKPT$aKyU$aLU$aMBU$aMH-GM
SE$aNIC$aNN$aNNC$aNNE$aNRU$aNSyU$aNcD$aNcRS$aNcU$aNhU$aNjP$aNjR$aNmU$aOCU$aOU$aOkU$aOrCS$aPB
S$aTU$aTxCM$aTxDaM$aTxHR$aULA$aUPB$aUU$aViBlbV$aViU$aViW$aVtU$aWM$aWaU$aWvU$aWyU
891      33      $81$a v. $bno. $u4 $v r $i (year) $j (month) $wq $x01
891      40      $81$a <1> $i <1987> $xprovisional
891      33      $82$a v. $bno. $u6 $v r $i (year) $j (month) $wb $x01
891      40      $82$a <2> - $i <1988> - $xprovisional
994      $aC0$bGZM
997      $aMARCIVE

LEADER    02025nas 22005657a 4500

001      991021971584702122
005      20180213162544.0
006      m o d
007      cr un|
008      050615c19879999nyubr pso 0 a0eng c
010      $a 2005212472

```

016 7_ \$a2051708-7\$2DE-600

022 \$a1558-156X

035 \$a(OCOLC) 60638386

035 \$a(CKB) 954925556449

035 \$a(CONSER) 2005212472

035 \$a(DE-599) ZDB2051708-7

035 \$a(EXLCZ) 99954925556449

037 \$bIEEE Service Center, 445 Hoes Lane, P.O. Box 1331 Piscataway, NJ 08855-1331

042 \$ansdp\$aapcc

050 10 \$aISSN RECORD

050 _4 \$aTK5105.5\$b.I324

082 10 \$a004.6\$213

130 0_ \$aIEEE network (Online)

222 _0 \$aIEEE network\$b(Online)

245 10 \$aIEEE network\$h[electronic resource].

260 \$a[New York, N.Y. :\$bInstitute of Electrical and Electronics Engineers], \$c1987-

310 \$aBimonthly, \$b1988-

321 \$aQuarterly, \$b1987

336 \$atext\$btxt\$2rdacontent

337 \$acomputer\$bc\$2rdamedia

338 \$aonline resource\$bcr\$2rdacarrier

362 1_ \$aPrint began with v. 1, no. 1 (Jan. 1987).

538 \$aMode of access: World Wide Web.

550 \$a"A publication of the IEEE Communications Society."

588 \$aDescription based on print version record.

588 \$aLatest issue consulted: Vol. 27, issue 5 (Sept./Oct. 2013) (IEEE Xplore Digital Libra:

500 \$aRefereed/Peer-reviewed

650 _0 \$aComputer networks\$vpPeriodicals.

650 _6 \$aRéseaux d'ordinateurs\$xpériodiques.

710 2_ \$aIEEE Communications Society.

710	2_	\$aInstitute of Electrical and Electronics Engineers.
776		\$x0890-8044
906		\$aJOURNAL

MMS IDs
Document ID: 999583751502121
Network Electronic IDs: 9911025563702121
Network Physical IDs: 999583751502121
mms_mad_ids: 9911911773602122, 991021971584702122
mms_ml_ids: 997254443402124
mms_gb_ids: 991007109864002123
mms_lc_ids: 991016969520402125
mms_plt_ids: 991013491598602128
mms_rf_ids: 991013813357102129

APPENDIX 1005F

Appendix 1005-X

Early Citations to Xiao

Source Document

Xiao, X., & Ni, L. M. (1999). Internet QoS: A big picture. *IEEE Network*, 13(2), 8-18.

Selected Early Citations (verified): May 1999 – April 10, 2000 (in chronological order)

Metz, C. (1999, May). RSVP: General-purpose signaling for IP. *IEEE Internet Computing*, 3(3), 95-99.

Gunter, M., & Braun, T. (1999, October). Evaluation of bandwidth broker signaling. In *Proceedings. Seventh International Conference on Network Protocols* (pp. 145-152). IEEE.

Altinkaya, M. (2000, March). *Server probing for server and agent based active network management*. NAVAL POSTGRADUATE SCHOOL MONTEREY CA.

Huang, C. M., Kung, H. Y., & Yang, J. L. (April 10, 2000). PARK: a paused-and-run k-stream multimedia synchronization control scheme. In *Proceedings 20th IEEE International Conference on Distributed Computing Systems* (pp. 272-279). IEEE.

Any time

Since 2022

Since 2021

Since 2018

Custom range...

Sort by relevance

Sort by date

☒ Create alert

Internet QoS: A big picture

☐ Search within citing articles

Traffic Engineering with MPLS in the Internet

X Xiao, A Hannan, B Bailey, LM Ni - IEEE network, 2000 - ieexplore.ieee.org

This article discusses traffic engineering with multiprotocol label switching (MPLS) in an Internet service provider's network. We first review MPLS, constraint-based routing, and ...

☆ Save [77](#) Cite Cited by 740 Related articles All 18 versions[\[PDF\] utdallas.edu](#)

Multi-constrained optimal path selection

T Korkmaz, M Krunz - ... Twentieth Annual Joint Conference of the ..., 2001 - ieexplore.ieee.org

Providing quality-of-service (QoS) guarantees in packet networks gives rise to several challenging issues. One of them is how to determine a feasible path that satisfies a set of ...

☆ Save [77](#) Cite Cited by 630 Related articles All 12 versions[\[PDF\] psu.edu](#)

Algebra and algorithms for QoS path computation and hop-by-hop routing in the Internet

JL Sobrinho - ... IEEE INFOCOM 2001. Conference on Computer ..., 2001 - ieexplore.ieee.org

Prompted by the advent of QoS routing in the Internet, we investigate the properties that path weight functions must have so that hop-by-hop routing is possible and optimal paths can be ...

☆ Save [77](#) Cite Cited by 403 Related articles All 13 versions[\[PDF\] psu.edu](#)

Host-to-host congestion control for TCP

A Afanasyev, N Tilley, P Reiher... - ... surveys & tutorials, 2010 - ieexplore.ieee.org

The Transmission Control Protocol (TCP) carries most Internet traffic, so performance of the Internet depends to a great extent on how well TCP works. Performance characteristics of a ...

☆ Save [77](#) Cite Cited by 370 Related articles All 20 versions[\[PDF\] etf.unsa.ba](#)

[PDF] QoS support in mobile ad hoc networks

K Wu, J Harms - Crossing Boundaries-the GSA Journal of ..., 2001 - manet.eurecom.fr

[\[PDF\] eurecom.fr](#)

Articles

About 102 results (0.02 sec)

Any time

Since 2022

Since 2021

Since 2018

Custom range...

1999 — 2000

Search

Sort by relevance

Sort by date

☒ Create alert**Internet QoS: A big picture**☐ Search within citing articles**Traffic Engineering with MPLS in the Internet****X Xiao, A Hannan, B Bailey, LM Ni** - IEEE network, 2000 - [ieeexplore.ieee.org](#)

This article discusses traffic engineering with multiprotocol label switching (MPLS) in an Internet service provider's network. We first review MPLS, constraint-based routing, and ...

☆ Save [PDF](#) Cite Cited by 738 Related articles All 18 versions[\[PDF\] utdallas.edu](#)**Internet quality of service: An overview****W Zhao, D Olshefski, HG Schulzrinne** - 2000 - [academiccommons.columbia.edu](#)

This paper presents an overview of Internet QoS, covering motivation and considerations for adding QoS to Internet, the definition of Internet QoS, traffic and service specifications ...

☆ Save [PDF](#) Cite Cited by 202 Related articles All 13 versions [»»](#)[\[PDF\] columbia.edu](#)**An efficient algorithm for finding a path subject to two additive constraints****T Korkmaz, M Krunz, S Tragoudas** - ACM SIGMETRICS Performance ..., 2000 - [dl.acm.org](#)

One of the key issues in providing end-to-end quality-of-service guarantees in packet networks is how to determine a feasible route that satisfies a set of constraints while ...

☆ Save [PDF](#) Cite Cited by 95 Related articles All 10 versions[\[PDF\] acm.org](#)**Providing quality of service in the Internet****X Xiao** - 2000 - [search.proquest.com](#)

This dissertation presents a framework for providing quality of service (QoS) in the Internet.

This framework consists of traffic directing and load balancing at the application layer ...

☆ Save [PDF](#) Cite Cited by 65 Related articles All 4 versions [»»](#)[\[PDF\] coteor.net.bo](#)**Evaluation of bandwidth broker signaling****M Gunter, T Braun** - Proceedings. Seventh International ..., 1999 - [ieeexplore.ieee.org](#)

The differentiated services (DiffServ) architecture for the Internet implements a scalable mechanism for quality-of-service (QoS) provisioning. Bandwidth brokers represent the ...

☆ Save [PDF](#) Cite Cited by 62 Related articles All 15 versions[\[PDF\] psu.edu](#)

APPENDIX 1005G



Internet QoS: A Big Picture

Xipeng Xiao and Lionel M. Ni, Michigan State University

Abstract

In this article we present a framework for the emerging Internet quality of service (QoS). All the important components of this framework — integrated services, RSVP, differentiated services, multiprotocol label switching (MPLS), and constraint-based routing — are covered. We describe what integrated services and differentiated services are, how they can be implemented, and the problems they have. We then describe why MPLS and constraint-based routing have been introduced into this framework, how they differ from and relate to each other, and where they fit into the differentiated services architecture. Two likely service architectures are presented, and the end-to-end service deliveries in these two architectures are illustrated. We also compare ATM networks to router networks with differentiated services and MPLS. Putting all these together, we give the readers a grasp of the big picture of the emerging Internet QoS.

Today's Internet only provides *best-effort service*. Traffic is processed as quickly as possible, but there is no guarantee as to timeliness or actual delivery. With the rapid transformation of the Internet into a commercial infrastructure, demands for service quality have rapidly developed [1–4].

It is becoming apparent that several service classes will likely be demanded. One service class will provide predictable Internet services for companies that do business on the Web. Such companies will be willing to pay a certain price to make their services reliable and give their users a fast feel of their Web sites. This service class may contain a single service, or it may contain *gold service*, *silver service*, and *bronze service*, with decreasing quality. Another service class will provide low-delay and low-jitter services to applications such as Internet telephony and videoconferencing. Companies will be willing to pay a premium price to run a high-quality videoconference to save travel time and cost. Finally, best-effort service will remain for those customers who only need connectivity.

Whether mechanisms are even needed to provide quality of service (QoS) is a hotly debated issue. One opinion is that fibers and wavelength-division multiplexing (WDM) will make bandwidth so abundant and cheap that QoS will be automatically delivered. The other opinion is that no matter how much bandwidth the networks can provide, new applications will be invented to consume them; therefore, mechanisms will still be needed to provide QoS. This argument is beyond the scope of this article [5]. Here we simply note that even if bandwidth will eventually become abundant and cheap, it is not going to happen soon. For now, some simple mechanisms are definitely needed in order to provide QoS on the Internet. Our view is supported by the fact that all the major router/switch vendors now provide some QoS mechanisms in their high-end products [6–11].

The Internet Engineering Task Force (IETF) has proposed

many service models and mechanisms to meet the demand for QoS. Notably among them are the integrated services/Resource Reservation Protocol (RSVP) model [4, 12], the differentiated services (DS) model [13, 14], multiprotocol label switching (MPLS) [15], traffic engineering [16], and constraint-based routing [17].

The integrated services model is characterized by resource reservation. For real-time applications, before data are transmitted, the applications must first set up paths and reserve resources. RSVP is a signaling protocol for setting up paths and reserving resources. In differentiated services, packets are marked differently to create several packet classes. Packets in different classes receive different services. MPLS is a forwarding scheme. Packets are assigned labels at the ingress of an MPLS-capable domain. Subsequent classification, forwarding, and services for the packets are based on the labels. Traffic engineering is the process of arranging how traffic flows through the network. Constraint-based routing is to find routes that are subject to some constraints, such as bandwidth or delay requirement.

Although there are many papers on each of integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing, to the best of the authors' knowledge, they are never discussed together in a single paper. As a result, it is difficult for readers to understand the relationships among them and to grasp the big picture of the QoS framework.

In this article we give an introduction to integrated services, RSVP, differentiated services, MPLS, traffic engineering, and constraint-based routing. We describe how they differ from, relate to, and work with each other to deliver QoS on the Internet. Through this, we intend to present to readers a clear overview of Internet QoS.

The organization of the rest of the article is as follows. In the next two sections, we describe integrated services, RSVP,

Flow	A stream of packets with the same source IP address, source port number, destination IP address, destination port number, and protocol ID.
Service level agreement (SLA)	A service contract between a customer and a service provider that specifies the forwarding service a customer should receive. Customer may be a user organization or another provider domain (upstream domain).
Traffic profile	A description of the properties of a traffic stream, such as rate and burst size.
Differentiated services (DS) field	The field in which the differentiated service class is encoded. It is the Type of Service (TOS) octet in the IP header or the traffic class field in the IPv6 header.
Per-hop behavior (PHB)	The externally observable behavior of a packet at a DS-compliant router.
Mechanism	A specific algorithm or operation (e.g., queuing discipline) that is implemented in a router to realize a set of one or more per-hop behaviors.
Admission control	The decision process of whether to accept a request for resources (link bandwidth plus buffer space).
Classification	The process of sorting packets based on the content of packet headers according to defined rules.
Behavior aggregate (BA) classification	The process of sorting packets based only on the contents of the DS field.
Multifield (MF) classification	The process of classifying packets based on the content of multiple fields such as source address, destination address, TOS byte, protocol ID, source port number, and destination port number.
Marking	The process of setting the DS field in a packet.
Policing	The process of handling out-of-profile traffic (e.g., discarding excess packets).
Shaping	The process of delaying packets within a traffic stream to cause it to conform to some defined traffic profile.
Scheduling	The process of deciding which packet to send first in a system of multiple queues.
Queue management	Controlling the length of packet queues by dropping packets when necessary or appropriate.
Traffic trunk	An aggregation of flows with the same service class that can be put into an MPLS label-switched path.

■ Table 1. Frequently used terminologies.

and differentiated services, their characteristics, mechanisms, and problems. A likely differentiated services architecture and the complete process for delivering end-to-end services in this architecture are also presented. MPLS and a service architecture based on MPLS are then described. We then describe traffic engineering and constraint-based routing. ATM networks and router networks are compared. Finally, we summarize the article in the last section.

The frequently used terminologies in this article are defined in Table 1.


Integrated Services and RSVP

The integrated services model [4] proposes two service classes in addition to best-effort service. They are:

- *Guaranteed service* [18] for applications requiring fixed delay bound
- *Controlled-load service* [19] for applications requiring reliable and enhanced best-effort service

The philosophy of this model is that "there is an inescapable requirement for routers to be able to reserve resources in order to provide special QoS for specific user packet streams, or flows. This in turn requires flow-specific state in the routers" [4].

RSVP was invented as a signaling protocol for applications to reserve resources [12]. The signaling process is illustrated in Fig. 1. The sender sends a PATH message to the receiver specifying the characteristics of the traffic. Every intermediate



■ Figure 1. RSVP signaling.

router along the path forwards the PATH message to the next hop determined by the routing protocol. Upon receiving a PATH message, the receiver responds with a RESV message to request resources for the flow. Every intermediate router along the path can reject or accept the request of the RESV message. If the request is rejected, the router will send an error message to the receiver, and the signaling process will terminate. If the request is accepted, link bandwidth and buffer space are allocated for the flow, and the related flow state information will be installed in the router.

Recently, RSVP has been modified and extended in several ways to reserve resources for aggregation of flows, to set up explicit routes (ERs) with QoS requirements, and to do some other signaling tasks [20–22]. Whether or not this is appropriate is a hotly debated issue in the IETF and is beyond the scope of this article.

Integrated services is implemented by four components: the

signaling protocol (e.g., RSVP), the admission control routine, the classifier, and the packet scheduler. Applications requiring guaranteed or controlled-load service must set up the paths and reserve resources before transmitting their data. The admission control routines will decide whether a request for resources can be granted. When a router receives a packet, the classifier will perform a *multifield* (MF) classification and put the packet in a specific queue based on the classification result. The packet scheduler will then schedule the packet accordingly to meet its QoS requirements.

The integrated services/RSVP architecture is influenced by the work of Ferrari *et al.* [23, 24]. It represents a fundamental change to the current Internet architecture, which is founded on the concept that all flow-related state information should be in the end systems [25].

The problems with the integrated services architecture are:

- The amount of state information increases proportionally with the number of flows. This places a huge storage and processing overhead on the routers. Therefore, this architecture does not scale well in the Internet core.
- The requirement on routers is high. All routers must have RSVP, admission control, MF classification, and packet scheduling.
- Ubiquitous deployment is required for guaranteed service. Incremental deployment of controlled-load service is possible by deploying controlled-load service and RSVP functionality at the bottleneck nodes of a domain and tunneling the RSVP messages over other parts of the domain.

Differentiated Services

Because of the difficulty in implementing and deploying integrated services and RSVP, differentiated services (DS) are introduced.

An Introduction to Differentiated Services

IPv4 header contains a Type of Service (TOS) byte. Its meaning was previously defined in [26]. Applications can set three bits in the TOS byte to indicate the need for low-delay, high-throughput, or low-loss-rate service. However, choices are limited. Differentiated services defines the layout of the TOS byte (*DS field*) and a base set of packet forwarding treatments (per-hop behaviors, or PHBs) [27]. By marking the DS fields of packets differently and handling packets based on their DS fields, several differentiated service classes can be created. Therefore, differentiated services is essentially a relative-priority scheme.

In order for a customer to receive differentiated services from its Internet service provider (ISP), it must have a service level agreement (SLA) with its ISP. An SLA basically specifies the service classes supported and the amount of traffic allowed in each class. An SLA can be static or dynamic. *Static SLAs* are negotiated on a regular (e.g., monthly or yearly) basis. Customers with *dynamic SLAs* must use a signaling protocol (e.g., RSVP) to request services on demand.

Customers can mark DS fields of individual packets to indicate the desired service or have them marked by the leaf router based on MF classification.

At the ingress of the ISP networks, packets are classified, policed, and possibly shaped. The classification, policing, and shaping rules used at the ingress routers are derived from the SLAs. The amount of buffering space needed for these operations is also derived from the SLAs. When a packet enters one domain from another domain, its DS field may be remarked as determined by the SLA between the two domains.

Using the classification, policing, shaping, and scheduling mechanisms, many services can be provided, such as:

- *Premium service* for applications requiring low-delay and low-jitter service
- *Assured service* for applications requiring better reliability than best-effort service
- *Olympic service*, which provides three tiers of services: *gold*, *silver*, and *bronze*, with decreasing quality [28, 29]

Note that differentiated services only defines DS fields and PHBs. It is ISPs' responsibility to decide which services to provide.

Differentiated services is significantly different from integrated services. First, there are only a limited number of service classes indicated by the DS field. Since service is allocated in the granularity of a class, the amount of state information is proportional to the number of classes rather than the number of flows. Differentiated services is therefore more scalable. Second, sophisticated classification, marking, policing, and shaping operations are only needed at the boundary of the networks. ISP core routers need only to have *behavior aggregate* (BA) classification. Therefore, it is easier to implement and deploy differentiated services.

There is another reason the second feature is desirable for ISPs. ISP networks usually consist of boundary routers connected to customers and core routers/switches interconnecting the boundary routers. Core routers must forward packets very quickly, and therefore must be simple. Boundary routers need not forward packets very quickly because customer links are relatively slow. Therefore, they can spend more time on sophisticated classification, policing and shaping [3]. Boundary routers at the *network access points* (NAPs) are exceptions. They must forward packets very quickly and do sophisticated classification, policing, and shaping. Therefore, they must be well equipped.

In the differentiated services model, incremental deployment is possible for assured service. DS-incapable routers simply ignore the DS fields of the packets and give the assured service packets best-effort service. Since assured service packets are less likely to be dropped by DS-capable routers, the overall performance of assured service traffic will be better than that of best-effort traffic.

An End-to-End Service Architecture

In this section a service architecture for differentiated services is presented. This architecture provides assured service and premium service in addition to best-effort service. It is mainly based on the architecture proposed in [28]. Other possible service architectures also exist [30].

Assured Service — Assured service is intended for customers that need reliable services from their service providers, even in times of network congestion. Customers will have SLAs with their ISPs. The SLAs will specify the amount of bandwidth allocated for the customers. Customers are responsible for deciding how their applications share that amount of bandwidth. One possible service allocation process is described later. SLAs for assured service are usually static, meaning that customers can start data transmission whenever they want without signaling their ISPs.

Assured service can be implemented as follows. First, classification and policing are done at the ingress routers of the ISP networks. If the assured service traffic does not exceed the bit rate specified by the SLA, they are considered *in profile*; otherwise, the excess packets are considered *out of profile*. Second, all packets, *in* and *out*, are put into an *assured queue* (AQ) to avoid out of order delivery. Third, the queue is managed by a queue management scheme called *random early detection* (RED) with *In* and *Out* — *RIO* [31].

RED [32] is a queue management scheme that drops packets randomly. This will trigger the TCP flow control mecha-

nisms at different end hosts to reduce send rates at different time. By doing so, RED can prevent the queue at the routers from overflowing, and therefore avoid the tail-drop behavior (dropping all subsequent packets when a queue overflows). Tail-drop triggers multiple TCP flows to decrease and later increase their rates simultaneously. It causes network utilization to oscillate and can hurt performance significantly. RED has proven to be useful and has been widely deployed.

RIO is a more advanced RED scheme. It basically maintains two RED algorithms, one for in packets and one for out packets. There are two thresholds for each queue. When the queue size is below the first threshold, no packets are dropped. When the queue size is between the two thresholds, only out packets are randomly dropped. When the queue size exceeds the second threshold, indicating possible network congestion, both in and out packets are randomly dropped, but out packets are dropped more aggressively. In addition to breaking the TCP flow-control synchronization, RIO prevents, to some extent, greedy flows from hurting the performance of other flows by dropping the out packets more aggressively.

Because in packets have low loss even in cases of congestion, the customers will perceive a predictable service from the network if they keep traffic conformant. When there is no congestion, out packets will also be delivered. The networks are thus better utilized.


Best-effort traffic can be treated differently, or identically, from assured service out traffic. In this article we assume that they are treated identically. Therefore, conceptually, we can consider that there is an *A-bit* in the DS field. The A-bits of assured service in packets are set to 1, while the A-bits of assured service out packets and best-effort packets are reset to 0.

Premium Service — Premium service provides low-delay and low-jitter service for customers that generate fixed peak bit rate traffic. Each customer will have an SLA with its ISP. The SLA specifies a desired peak bit-rate for a specific flow or an aggregation of flows. The customer is responsible for not exceeding the peak rate; otherwise, excess traffic will be dropped. The ISP guarantees that the contracted bandwidth will be available when traffic is sent. Premium service is suitable for Internet telephony, videoconferencing, or for creating virtual lease lines for virtual private networks (VPNs) [33].

Because premium service is more expensive than assured service, it is desirable for ISPs to support both static and dynamic SLAs. Dynamic SLAs allow customers to request premium service on demand without subscribing to it. Admission control is needed for dynamic SLAs.

Premium service can be implemented as follows. At the customer side, some entity will decide which application flow can use premium service. The leaf routers directly connected to the senders will do MF classifications and shape the traffic. Conceptually, we can consider that there is a *P-bit* in the DS field. If the P-bit of a packet is set, this packet belongs to the premium class; otherwise, the packet belongs to the assured service or best-effort class. After the shaping, the P-bits of all packets are set for the flow that is allowed to use premium service. The exit routers of the customer domain may need to reshape the traffic to make sure that the traffic does not exceed the peak rate specified by the SLA. At the provider side, the ingress routers will police the traffic. Excess traffic is dropped. All packets with the P-bit set enter a *premium queue* (PQ). Packets in the PQ will be sent before packets in the AQ.

First, by admission control the amount of premium traffic can be limited to a small percentage, say 10 percent, of the bandwidth of input links. Second, excess packets are dropped at the ingress routers of the networks. Nonconformant flows cannot impact the performance of conformant flows. Third,



■ Figure 2. Uneven distribution of premium traffic in an ISP. The shaded area is the core of the ISP.

premium packets are forwarded before packets of other classes; they can potentially use 100 percent of the bandwidth of the output links. Since most links are full-duplex, the bandwidth of the input links equals the bandwidth of the output links. Therefore, if premium traffic is distributed evenly among the links, these three factors should guarantee that the service rate of the PQ is much higher than the arrival rate. Therefore, arriving premium packets should find the PQ empty or very short most of the time. The delay or jitter experienced by premium packets should be very low. However, premium service provides no quantified guarantee on the delay or jitter bound.

However, uneven distribution of premium traffic may cause a problem for premium service. In ISP networks, aggregation of traffic from the boundary routers to a core router (e.g., CR1 in Fig. 2) is inevitable; but this is not a problem because the output link is much faster than the input links. However, aggregation of premium traffic in the core at CR4 may invalidate the assumption that the arrival rate of premium traffic is far below the service rate. Differentiated services alone cannot solve this problem. Traffic engineering/constraint-based routing must be used to avoid congestion caused by such uneven traffic distribution.

By limiting the total amount of bandwidth requested by premium traffic, the network administrators can guarantee that premium traffic will not starve the assured and best-effort traffic. Another scheme is to use Weighted Fair Queuing (WFQ) [34] between the PQ and the AQ.

Service Allocation in Customer Domains — Given an SLA, a customer domain should decide how its hosts share the services specified by the SLA. This process is called *service allocation*.

There are basically two choices:

- Each host makes its own decision as to which service to use.
- A resource controller called the *bandwidth broker* (BB) [28] makes decision for all hosts.

A BB can be a host, a router, or a software process on an exit router. It is configured with the organizational policies and manages the resources of a domain. A domain may also have backup BBs. Since all hosts must cooperate to share a limited amount of resources specified by the SLA, it is technically better to have a BB to allocate resources.

At the initial deployment stage, hosts need no DS mechanism. They simply send their packets unmarked. The exit routers marked them before sending them out to the ISPs. The packets are treated as best-effort traffic inside the customer domain. In later deployment stages, hosts may have some signaling or marking mechanisms. Before a host starts sending packets, it may decide the service class for the packets by itself, or it may consult a BB for a service class. The host may mark the packets by itself or send the packets unmarked. If the host

sends the packets unmarked, the BB must use some protocols, such as RSVP or Lightweight Directory Access Protocol (LDAP) [35], to set the classification, marking, and shaping rules at the leaf router directly connected to the sender so that the leaf router knows how to mark the sender's packets.

If the SLA between a customer and its ISP is dynamic, the BB in the customer domain must also use some signaling protocol to request resources on demand from its ISP. From now on, we assume that RSVP is used as the signaling protocol.

Resource Allocation in ISP Domains — Given the SLAs, ISPs must decide how to configure their boundary routers so that they know how to handle the incoming traffic. This process is called *resource allocation*.

For static SLAs, boundary routers can be manually configured with the classification, policing, and shaping rules. Resources are therefore statically allocated for each customer. Unused resources can be shared by other customers.

For a dynamic SLA, resource allocation is closely related to the signaling process. The BB in the customer domain uses RSVP to request resources from its ISP. At the ISP side, admission control decisions can be made in a distributed manner by the boundary routers, or centrally by a BB. If boundary routers are directly involved in the signaling process, they are configured with the corresponding classification, policing, and shaping rules when they grant a request. If a BB is involved rather than the boundary routers, the BB must configure the boundary routers when it grants a request. Such procedures will be detailed in the next section.

In both cases, the ISP core routers must be shielded from the requests to avoid the scalability problem.

Examples of End-to-End Service Delivery

Example 1: Delivery of Assured Service with a Static SLA — In Fig. 3, host S in corporate network 1 (CN1) wants to use assured service to send data to host D in corporate network 2 (CN2). CN1 has a static SLA with ISP1. The numbers inside the circles are the step numbers in the service delivery process, described below.

- 1 Host S sends a RSVP message to the local BB, CN1-BB, requesting for assured service for its traffic.
- 2 If CN1-BB grants the request, it will configure leaf router LR1 so that LR1 can set the A-bits of the packets of this flow. CN1-BB will then reply to host S; otherwise, an error message is sent to host S.
- 3 Host S sends packets to leaf router LR1.
- 4 If LR1 is configured to mark the traffic, it will set the A-bits of the packets.
- 5 Every router from LR1 (exclusive) to ER1 (inclusive) does a BA classification. Packets with the A-bit set are consid-

ered in, while packets with the A-bit reset are considered out. All packets enter the AQ. RIO is applied on the AQ.

- 6 BR1 polices the traffic. All out traffic remains out. If the in traffic exceeds its bit rate, the excess packets' A-bits will be reset. All packets enter the AQ. RIO is applied on the queue.

- 7 All routers between boundary routers BR1 and BR2 (inclusive) perform BA classifications and apply RIO on their AQs.

- 8 ER2 performs the same operations as BR1.

- 9 The packets are eventually delivered to host D.

Note that:

- If there are multiple ISPs between CN1 and CN2, steps 6 and 7 will be repeated multiple times, once per ISP.
- If CN1 does not have any SLA with ISP1, it can only send traffic as best effort. No matter how the routers in CN1 mark the DS fields of their packets, the A-bits will be reset at BR1.

Example 2: Delivery of Premium Service with a Dynamic SLA — In Fig. 4, host S in CN1 wants to use premium service to send data to host D in CN2. CN1 has a dynamic SLA with ISP1.

Phase 1: Signaling

- 1 Host S sends an RSVP PATH message to the local BB, CN1-BB.

- 2 CN1-BB makes an admission control decision.

- If the request is denied, an error message is sent back to host S. The signaling process ends.

- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to ISP1-BB.

- 4 ISP1-BB makes an admission control decision.

- If the request is denied, an error message is sent back to CN1-BB. Sender S will be notified.

- If the request is accepted, ISP1-BB sends the PATH message to CN2-BB.


- 5 CN2-BB makes an admission control decision.

- If the request is denied, an error message is sent back to ISP1-BB. Sender S will be notified.


- If the request is accepted, CN2-BB will use LDAP or RSVP to set the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to ISP1-BB.

- 6 When ISP1-BB receives the RESV message, it will configure the classification and policing rules on router BR1, and the policing and reshaping rules on router BR2. It will then send the RESV message to CN1-BB.

- 7 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so that if the traffic of the admitted flow is nonconformant, LR1 will shape it. CN1-BB will also set the policing and reshaping



■ Figure 3. The delivery process of assured service with a static SLA.



■ Figure 4. The delivery process of premium service with a dynamic SLA.

rules on router ER1. CN1-BB will then send the RESV message to host S.

8 When host S receives the RESV message, it can start transmitting data.

Note that:

- This signaling process is significantly different from the signaling process in the integrated services/RSVP model. First, it is the sender who requests resources, not the receiver. Second, a request can be rejected when the BB receives the PATH message from the sender. In integrated services/RSVP, a request is rejected only when a router receives the RESV message from the receiver. Third, a BB can aggregate multiple requests and make a single request to the next BB. Fourth, each domain behaves like a single node, represented by the BB. ISP core routers are not involved in this process.
- The state information installed by the BB on the boundary routers is soft state. It must be regularly refreshed, or it will time out.
- If there are multiple ISPs between CN1 and CN 2, repeat steps 4 and 6 once for each ISP.
- If the SLA between CN1 and ISP1 is static, simply skip steps 3–6 in the signaling process.

Phase 2: Data Transmission

9 Host S sends packets to leaf router LR1.

10 Leaf router LR1 performs an MF classification. If the traffic in nonconformant, LR1 will shape it. LR1 will also set the P-bits of the packets. All packets enter the PQ.

11 Each intermediate router between LR1 and ER1 performs a BA classification, puts the packets into the PQ, and sends them out.

12 ER1 performs a BA classification and reshapes the traffic to make sure that the negotiated peak rate is not exceeded. Reshaping is done for the aggregation of all flows heading toward BR1, not for each individual flow.

13 BR1 classifies and polices the premium traffic. Excess premium packets are dropped.

14 Intermediate routers between leaf routers BR1 and BR2 (inclusive) perform BA classifications. BR2 also reshapes the premium traffic.

15 ER2 classifies and polices the premium traffic. Excess premium packets are dropped.

16 The premium packets are delivered to host D.

Requirements on Routers

The requirements on routers to support premium service and assured service are summarized below.

- The leaf routers in customer domains need MF classifications, marking, and shaping.

- The ISP ingress routers need policing and remarking.
- The ISP egress routers optionally need re-shaping.
- All routers need BA classification and two queues with strict priority.
- If dynamic SLAs are supported, each customer domain will need a BB to request service on behalf of the domain and to allocate services inside the domain. Signaling and admission control mechanisms are needed in both customer and ISP domains.

If assured service is to be replaced by olympic service, the AQ must be replaced by three queues: a *gold queue*, a *silver queue*, and a *bronze queue*. WFO can be used to schedule these queues. The rate parameters of these queues can be manually configured based on experience.

MPLS


The motivation for MPLS is to use a fixed-length label to decide packet handling. MPLS is also a useful tool for traffic engineering [16, 36].

An Introduction to MPLS

MPLS is a forwarding scheme. It evolved from Cisco's *Tag Switching*. In the open systems interconnection (OSI) seven-layer model, it is between layer 2 (L2, link layer) and layer 3 (L3, network layer).

Each MPLS packet has a header. The header contains a 20-bit label, a 3-bit Class of Service (COS) field, a 1-bit label stack indicator, and an 8-bit Time to Live (TTL) field. The MPLS header is encapsulated between the link layer header and the network layer header. An MPLS-capable router, called the *label-switched router* (LSR), examines only the label in forwarding the packet. The network protocol can be IP or others. This is why it is called *multiprotocol label switching*.

MPLS needs a protocol to distribute labels to set up *label-switched paths* (LSPs). Whether a generic *label distribution protocol* (LDP) [37] should be created or RSVP should be extended [22] for this purpose is another hotly debated issue. MPLS labels can also be piggybacked by routing protocols. An LSP is similar to an ATM virtual circuit (VC) and is unidirectional from the sender to the receiver. MPLS LSRs use the protocol to negotiate the semantics of each label, that is, how to handle a packet with a particular label from the peer. LSP setup can be control-driven (i.e., triggered by control traffic such as routing updates) or data-driven (i.e., triggered by the request of a flow or *traffic trunk*). In MPLS, a traffic trunk is an aggregation of flows with the same service class that can be put into an LSP. The LSP between two routers can be the same as the L3 hop-by-hop route, or the sender LSR can specify an *explicit route* (ER) for the LSP. The ability to set up



■ Figure 5. The signaling process of dynamic premium service in an MPLS-based architecture.

ERs is one of the most useful features of MPLS. A forwarding table indexed by labels is constructed as the result of label distribution. Each forwarding table entry specifies how to process packets carrying the indexing label.

Packets are classified and routed at the ingress LSRs of an MPLS-capable domain. MPLS headers are then inserted. When an LSR receives a labeled packet, it will use the label as the index to look up the forwarding table. This is faster than the process of parsing the routing table in search of the longest match done in IP routing [38, 39]. The packet is processed as specified by the forwarding table entry. The incoming label is replaced by the outgoing label, and the packet is switched to the next LSR. This label-switching process is similar to ATM's VCI/VPI processing. Inside an MPLS domain, packet forwarding, classification, and QoS service are determined by the labels and the COS fields. This makes core LSRs simple. Before a packet leaves an MPLS domain, its MPLS label is removed.

MPLS LSPs can be used as tunnels. After LSPs are set up, a packet's path can be completely determined by the label assigned by the ingress LSR. There is no need to enumerate every intermediate router of the tunnel. Compared to other tunneling mechanisms, MPLS is unique in that it can control the complete path of a packet without explicitly specifying the intermediate routers.

In short, MPLS is strategically significant because:

- It provides faster packet classification and forwarding.
- It provides an efficient tunneling mechanism.

These features, particularly the second one, make MPLS useful for traffic engineering [16, 36].

A Service Architecture Based on MPLS

MPLS can be used together with differentiated services to provide QoS. In such an architecture, LSPs are first configured between each ingress-egress pair. For LSP(LSR1 → LSR2) and LSP(LSR2 → LSR1), their intermediate LSRs need not be reciprocal. It is likely that for each ingress-egress pair, a separate LSP is created for each traffic class. In this case, a total number of $C*N*(N-1)/2$ LSPs are needed, where C is the number of traffic classes and N is the number of boundary routers. In order to reduce the number of LSPs, the LSPs from all ingress routers to a single egress router can be merged into a *sink tree*. The total number of sink trees needed is $C*N$. It is also possible to use a single sink tree to transmit packets of different traffic classes, and use the COS bits to differentiate packet classes. In this case, the number of sink trees is reduced to N . In this architecture, as the number of transiting flows increases, the number of flows in each LSP or sink tree also increases. But the number of LSPs or sink trees need not increase. This architecture is therefore scalable.

The operations of the routers are basically the same in this architecture as in the DS-field-based architecture described earlier. There are three differences in the processing of a packet:

- At the ingress of the ISP network, in addition to all the processing described in the DS-field-based architecture, an MPLS header is inserted into the packet.
- Core routers process the packet based on its label and COS field rather than its DS field.
- At the egress, unless interdomain LSPs are configured, the MPLS header is removed.

Note that, with such schemes, the MPLS effect is confined within the ISPs that use MPLS. Whether a particular ISP's architecture is DS-field-based or MPLS-based is transparent to other ISPs. Therefore, the DS-field based and MPLS-based architectures can easily interoperate.

Each customer domain still needs a BB to allocate services, and to request resources on behalf of the customer domain when the SLA is dynamic. But since LSPs are configured within the ISPs, resource requests can easily be hidden from the core routers by tunneling them from the ingress routers to the egress routers. Therefore, BBs may not be needed in MPLS-based ISP networks. Admission control is done in a distributed fashion by the ingress and egress routers.

Without BBs in the ISP networks, the signaling process for dynamic SLAs is slightly different from the one described earlier. It is depicted in Fig. 5 and described below. The data transmission process remains the same except for the three differences noted above.

- 1 Host S sends an RSVP PATH message to its local domain BB CN1-BB.
- 2 CN1-BB makes an admission control decision.
 - If the request is denied, an error message is sent back to host S. The signaling process ends.
- 3 The request is accepted by CN1-BB. CN1-BB sends the PATH message to BR1.
- 4 BR1 decides if there are enough resources to send the traffic to egress router BR2.
 - If no, the request is denied. An error message is sent back to CN1-BB. Sender S will be notified.
 - If yes, ISP1-BB sends the PATH message through an LSP to BR2.
- 5 BR2 sends the PATH message to CN2-BB.
- 6 CN2-BB decides if its domain can support the traffic.
 - If no, the request is denied. An error message is sent back to BR2. Sender S will be notified.
 - If yes, the request is accepted. CN2-BB will use LDAP or RSVP to set the classification and policing rules on router ER2. CN2-BB will then send an RSVP RESV message to BR2.
- 7 BR2 configures the reshaping rules for the traffic. It then sends the RESV message through an LSP to BR1.

8 BR1 configures the classification and policing rules for the traffic. It then sends the RESV message to CN1-BB.
 9 When CN1-BB receives the RESV message, it will set the classification and shaping rules on router LR1, so if the traffic of the admitted flow is nonconformant, LR1 can shape it. CN1-BB will also set the reshaping rules on router ER1. CN1-BB will then pass the RESV message to host S.
 10 Sender S starts transmitting data.
 If there are multiple ISPs between CN1 and CN2, repeat steps 4, 5, 7, and 8 once per ISP.

Traffic Engineering and Constraint-Based Routing

QoS schemes such as integrated services/RSVP and differentiated services essentially provide differentiated degradation of performance for different traffic when traffic load is heavy. When the load is light, integrated services/RSVP, differentiated services, and best-effort service differ little. Then why not avoid congestion at the first place? This is the motivation for traffic engineering.

Traffic Engineering

Network congestion can be caused by lack of network resources or uneven distribution of traffic. In the first case, all routers and links are overloaded, and the only solution is to provide more resources by upgrading the infrastructure. In the second case, some parts of the network are overloaded while other parts are lightly loaded. Uneven traffic distribution can be caused by the current dynamic routing protocols such as RIP, OSPF, and IS-IS, because they always select the shortest paths to forward packets. As a result, routers and links along the shortest path between two nodes may become congested while routers and links along a longer path are idle. The *equal-cost multipath* (ECMP) option of OSPF [40], and recently of IS-IS [41], is useful in distributing load to several shortest paths; but if there is only one shortest path, ECMP does not help. For simple networks, it may be possible for network administrators to manually configure the cost of the links so that traffic can be evenly distributed. For complex ISP networks, this is almost impossible.

Traffic engineering is the process of arranging how traffic flows through the network so that congestion caused by uneven network utilization can be avoided. *Constraint-based routing* is an important tool for making the traffic engineering process automatic.

Avoiding congestion and providing graceful degradation of performance in congestion are complementary. Traffic engineering therefore complements differentiated services.

Constraint-Based Routing

In a sentence, constraint-based routing is used to compute routes that are subject to multiple constraints.

Constraint-based routing evolves from *QoS routing*. Given the QoS request of a flow or an aggregation of flows, QoS routing returns the route that is most likely to be able to meet the QoS requirements. Constraint-based routing extends QoS routing by considering other constraints of the network such as policy. The goals of constraint-based routing are:

- To select routes that can meet certain QoS requirements
- To increase utilization of the network

While determining a route, constraint-based routing considers not only network topology, but also requirements of the flow, resource availability of the links, and possibly other policies specified by the network administrators. Therefore, constraint-based routing may find a longer but lightly loaded path

better than the heavily loaded shortest path. Network traffic is thus distributed more evenly.

In order to do constraint-based routing, routers need to distribute new link state information and to compute routes based on such information.

Distribution of Link State Information — A router needs topology and resource availability information in order to compute QoS routes. Here, resource availability information means link available bandwidth [42]. Buffer space is assumed to be sufficient and is not explicitly considered [42].

One approach to distributing bandwidth information is to extend the link state advertisements of protocols such as OSPF and IS-IS [42, 43]. Because link residual bandwidth is frequently changing, a trade-off must be made between the need for accurate information and the need to avoid frequent flooding of link state advertisements.

To reduce the frequency of link state advertisements, one possible way is to distribute them only when there are topology or significant bandwidth changes (e.g., more than 50 percent or more than 10 Mb/s) [44]. A hold-down timer should always be used to limit the frequency of such advertisements. A recommended timer value is 30 seconds [45]. An approach to limit the flooding scope of such advertisements is described in [46].

Route Computation — The routing table computation algorithms in constraint-based routing and the complexity of such algorithms depend on the metrics chosen for the routes.

Common route metrics in constraint-based routing are monetary cost, hop count, bandwidth, reliability, delay, and jitter. Routing algorithms select routes that optimize one or more of these metrics.

Metrics can be divided into three classes. Let $d(i, j)$ be a metric for link (i, j) . For any path $P = (i, j, k, \dots, l, m)$, metric d is:

- *Additive* if $d(P) = d(i, j) + d(j, k) + \dots + d(l, m)$
- *Multiplicative* if $d(P) = d(i, j) * d(j, k) * \dots * d(l, m)$
- *Concave* if $d(P) = \min\{d(i, j), d(j, k), \dots, d(l, m)\}$

According to this definition, metrics *delay*, *jitter*, *cost*, and *hop count* are additive, *reliability* ($1 - \text{loss rate}$) is multiplicative, and *bandwidth* is concave.

A well-known theorem in constraint-based routing is that computing optimal routes subject to constraints of two or more additive and/or multiplicative metrics is *NP-complete* [47]. That is, algorithms that use two or more of delay, jitter, hop count, and loss probability as metrics, and optimize them simultaneously are NP-complete. The computationally feasible combinations of metrics are bandwidth and one of those metrics.

However, the proof of NP-completeness in [47] is based on the assumptions that all the metrics are independent, and the delay and jitter of a link are known a priori. Although such assumptions may be true in circuit-switched networks, metrics bandwidth, delay, and jitter are not independent in packet networks. As a result, polynomial algorithms for computing routes with hop count, delay, and jitter constraints exist [45]. The complexity of such algorithms is $O(N * E * e)$, where N is the hop count, E is the number of links of the network, and $e \leq E$ is the number of distinct bandwidth values among all links. Nevertheless, the theorem can tell us qualitatively the complexity of a routing algorithm: a complex algorithm in circuit-switched networks is still complex in packet networks, although it may not be NP-complete.

Fortunately, algorithms for finding routes with bandwidth and hop-count constraints are much simpler [42]. Bellman-Ford's (BF) Algorithm or Dijkstra's Algorithm can be used. For example, to find the shortest path between two nodes with bandwidth greater than 1 Mb/s, all the links with residual band-

width less than 1 Mb/s can be pruned first. BF or Dijkstra's Algorithm can then be used to compute the shortest path in the pruned network. The complexity of such algorithms is $O(N^2E)$.

Bandwidth and hop count are more useful constraints than delay and jitter, because:

- Although applications may care about delay and jitter bounds, few applications cannot tolerate occasional violation of such constraints. Therefore, there is no obvious need for routing flows with delay and jitter constraints. Besides, since delay and jitter parameters of a flow can be determined by the allocated bandwidth and the hop count of the route [48, 49], delay and jitter constraints can be mapped to bandwidth and hop-count constraints if needed.
- Many real-time applications will require a certain amount of bandwidth. The bandwidth metric is therefore useful. The hop count metric of a route is important because the more hops a flow traverses, the more resources it consumes. For example, a 1-Mb/s flow that traverses two hops consumes twice as many resources as one that traverses a single hop.

In constraint-based routing, routes can be computed on demand or precomputed for each traffic class. On-demand computations are triggered by the receipt of the QoS request of a flow. In either case, a router will have to compute its routing table more frequently with constraint-based routing than with dynamic routing. This is because, even without topology changes, routing table computation can still be triggered by significant bandwidth changes. Besides, constraint-based routing algorithms are at least as complex as dynamic routing algorithms. Therefore, the computation load of routers with constraint-based routing can be very high.

Common approaches to reduce the computation overhead of constraint-based routing include:

- Using a large-valued timer to reduce computation frequency
- Choosing bandwidth and hop count as constraints
- Using administrative policy to prune unsuitable links before computing the routing table

For example, if a flow has delay requirement, high propagation delay links such as satellite links are pruned before the routing table computation.

Constraint-Based Routing: Pros and Cons

The pros of constraint-based routing are meeting the needs for QoS requirements of flows better, and improved network utilization.

The cons of constraint-based routing are increased communication and computation overhead, increased routing table size, the fact that longer paths may consume more resources, and potential routing instability.

Of the cons, the first was addressed in an earlier section. The rest are addressed in this section.

In constraint-based routing, an essential issue is routing granularity. Routing can be destination-based, source-destination-based, class-based, traffic-trunk-based, or flow-based. Routing with finer granularity is more flexible, and thus more efficient in terms of resource utilization and more stable. However, the computation overhead and storage overhead are also higher.

Routing Table Structure and Size — Routing table structure and size depend directly on routing granularity and route metrics. Logically, we can view the routing table as a two-dimensional array. The number of rows is determined by routing granularity, and the number of columns is determined by route metrics. For example, in destination-based routing with bandwidth and hop count as route metrics, the routing table can be organized as a $K \times H$ array, where K is the number of destinations, and H is the maximum number of hops allowed for any route. The (k, h) th entry of the array contains one or

more h -hop routes for destination k . Each route also has an available bandwidth associated with it [42].

Obviously, the size of a constraint-based routing table can be far larger than the size of a normal routing table for the same network. This introduces significant storage overhead. It may also slow down the routing table lookup.

Approaches to reducing the routing table size in constraint-based routing include:

- Using coarse routing granularity
- Using hop quantization (i.e., dividing all hop-count values into a few classes to reduce the number of columns in the routing table) [44]
- Keeping the routing table only for best-effort traffic, and computing the routes for flows with QoS requests on demand [50]

The third scheme basically trades computation time for smaller storage requirements.

The Trade-off between Resource Conservation and Load Balancing — A constraint-based routing scheme can choose one of the following as the route for a destination:

- The *widest-shortest* path, that is, a path with minimum hop count and, if there are multiple such paths, the one with the largest available bandwidth.
- The *shortest-widest* path, that is, a path with the largest available bandwidth and, if there are multiple such paths, the one with the minimum hop count.
- The *shortest-distance* path. The distance of a k -hop path P is defined as

$$\text{dist}(P) = \sum_{i=1}^k \frac{1}{r_i}, \text{ where } r_i \text{ is the bandwidth of link } i.$$

Using paths other than the shortest paths consume more resources. This is not efficient when the load of the network is heavy. A trade-off must be made between resource conservation and load balancing. The first approach above is basically the same as today's dynamic routing. It emphasizes preserving network resources by choosing the shortest paths. The second approach emphasizes load balancing by choosing the widest paths. The third approach makes a trade-off between the two extremes. It favors shortest paths when network load is heavy and widest paths when network load is medium. Simulations showed that the third approach consistently outperforms the other two approaches for best-effort traffic, regardless of network topology and traffic pattern [45].

Stability — Because constraint-based routing algorithms recompute routing tables more frequently than dynamic routing algorithms do, they can introduce instability.

The stability of networks with constraint-based routing depends heavily on the routing granularity. If routing is done with coarse granularity (e.g., based solely on destination address), when the original route between two nodes becomes congested, all the traffic to that destination is shifted from the original route to an alternate route. This may cause congestion in the alternate route. Traffic may have to be shifted again [46].

The high computation overhead of constraint-based routing may also hurt the stability of the network. When a router is busy computing the routing table, it is slow in reacting to new topology changes.

To improve stability, the timer value for periodic routing table recomputation should be carefully chosen [50]. Constraint-based routing at the granularity of traffic trunk provides a good trade-off between stability and computation overhead [16]. Reducing the computation complexity of the routers also helps to improve stability.

In summary, constraint-based routing must be deployed

with caution. Otherwise, the cost of instability and increased complexity may outweigh the gain.

Constraint-based routing is similar to the dynamic/adaptive routing in telephone networks and ATM networks [51–54]. Many lessons can be learned from those works. Since constraint-based routing is a superset of today's dynamic routing, it is possible that in the future, constraint-based routing may replace dynamic routing, especially in the intradomain case. An emerging intradomain constraint-based routing protocol is QOSPF [42].

The Position of Constraint-Based Routing in the QoS Framework

In this section we describe the relationships between constraint-based routing and other components in the QoS framework.

The Relationship between Constraint-Based Routing and Differentiated Services — Constraint-based routing is to select the optimal routes for flows so that their QoS requirements are most likely to be met. It is not to replace differentiated services, but to help differentiated services be better delivered. Figure 2 shows an example in point.

The Relationship between Constraint-Based Routing and RSVP — RSVP and constraint-based routing are independent but complementary. For a router with dynamic routing, when an RSVP PATH message is received, it will be forwarded to the next hop determined by the dynamic routing protocol. The QoS requirement of the flow and the load of the networks are not considered in selecting the next hop. However, with a router running constraint-based routing, such information is considered. The next hop of the RSVP messages determined by constraint-based routing therefore may be different. In either case, the actual reservation of resources for the flow is done by RSVP. In short, constraint-based routing determines the path for RSVP messages, but does not reserve resources. RSVP reserves resources, but depends on constraint-based or dynamic routing to determine the path.

The Relationship between Constraint-Based Routing and MPLS — Given that MPLS is a forwarding scheme and constraint-based routing is a routing scheme, MPLS and constraint-based routing are, in theory, mutually independent. Constraint-based routing determines the route between two nodes based on resource information and topology information. It is useful with or without MPLS. Given the routes, MPLS uses its label distribution protocol to set up the LSPs. It does not care whether the routes are determined by constraint-based or dynamic routing.

However, when MPLS and constraint-based routing are used together, they make each other more useful. MPLS makes it possible to do constraint-based routing at traffic trunk granularity without introducing MF classification to the core routers. MPLS's per-LSP statistics provide constraint-based routing with precise information about the amount of traffic between every ingress–egress pair. Given such information, constraint-based routing can better compute the routes for setting up LSPs. In combination, MPLS and constraint-based routing provide powerful tools for traffic engineering.

A Comparison of ATM Networks to Router Networks

QoS and some sort of traffic engineering have long been provided by ATM networks. So why introduce differentiated ser-

vices and MPLS into the router networks? To answer this question, we give a brief comparison between ATM networks and router networks.

In an ATM network, QoS can be provided by allocating a certain amount of bandwidth for a specific VC. Traffic engineering is usually done by computing the routes offline and then downloading the configuration statically into the ATM switches on an hourly or daily basis. Per-permanent virtual circuit (PVC) traffic statistics of the current configuration provide accurate traffic information for computing the routes for the next configuration.

The advantages of ATM networks over router networks without differentiated services or MPLS are:

- ATM networks are currently faster in data forwarding.
- Per-PVC traffic statistics are available.
- QoS and some sort of traffic engineering are provided.

The disadvantages of ATM networks are:

- ATM cell header overhead is large.
- Routers must be used at the boundary of the network. With both switches and routers present in the network, two sets of configurations are required: one for routers and the other for switches.

With differentiated services and MPLS, router networks can also provide QoS and traffic engineering. This can be done without a big header overhead and two sets of configurations. Router networks with differentiated services and MPLS therefore provide some advantages over ATM networks [55]; but this is more or less from the perspective of router vendors.

Summary

The big picture of the emerging Internet QoS can be summarized as follows:

- Customers negotiate SLAs with ISPs. The SLAs specify what services the customers will receive. SLAs can be static or dynamic. For static SLAs, customers can transmit data at any time. For dynamic SLAs, customers must use a signaling protocol such as RSVP to request services on demand before transmitting data. The bandwidth brokers in the customer domains decide how applications share the services specified by the SLAs. The DS fields of packets are marked accordingly to indicate the desired services.
- The ingress routers of ISPs are configured with classification, policing, and remarking rules. The egress routers of ISP networks are configured with reshaping rules. Such rules may be configured manually by network administrators or dynamically by some protocol such as LDAP or RSVP. ISPs must implement admission control in order to support dynamic SLAs. Classification, marking, policing, and shaping/reshaping are only done at the boundary routers. Core routers are shielded from the signaling process. They need only implement two queues with strict priority. They process packets based solely on their DS fields.
- With MPLS, LSPs are set up between each ingress–egress pair. At the ISP ingress routers, labels and COS fields are determined from the classification and routing results. MPLS headers are then inserted into the packets. Core routers process packets based on their labels and COS fields only. Labels are removed before packets leave an MPLS domain.
- Constraint-based routing can be used to compute the routes subject to QoS and policy constraints. The goal is to meet the QoS requirements of traffic and to improve utilization of the networks.
- MPLS and constraint-based routing can be used together to control the path of traffic to avoid congestion and improve the utility of the networks.

Application Layer	
Transport Layer	Integrated Service/RSVP, Differentiated Services
Network Layer	Constraint Based Routing
	MPLS
Link Layer	

■ Table 2. The relative positions of the components in the QoS framework.

- The positions of integrated services/RSVP, differentiated services, MPLS, and constraint-based routing in the Internet network model are depicted in Table 2.

Acknowledgments

This article was written when the first author was an intern at Ascend Communications Inc. The support of Ascend Communications Inc. is gratefully acknowledged. The authors sincerely thank the anonymous reviewers for their insightful comments.

References

- [1] R. Comerford, "State of the Internet: Roundtable 4.0," *IEEE Spectrum*, Oct. 1998.
- [2] D. Ferrari and L. Delgrossi, "Charging For QoS," *IEEE/IFIP IWQOS '98* keynote paper, Napa, CA, May 1998.
- [3] P. Ferguson and G. Huston, *Quality of Service*, Wiley, 1998.
- [4] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," *Internet RFC 1633*, June 1994.
- [5] R. Jain, "Myths about Congestion Management in High Speed Networks," *Internetworking: Res. and Exp.*, vol. 3, 1992, pp. 101-13.
- [6] Cisco's 12000 Series, <http://www.cisco.com/warp/public/733/12000>
- [7] Ascend's GRF routers, <http://www.ascend.com/300.html>
- [8] Bay Networks' Accelar Routing Switches, <http://business5.baynetworks.com/MainBody.asp>
- [9] 3Com's switches, <http://www.3com.com/products/switches.html>
- [10] Juniper's M40 Series, <http://www.juniper.net/products/default.htm>
- [11] Lucent's PacketStar 6400 Series, <http://www.lucent.com:80/dns/products/ps6400.html>
- [12] R. Braden et al., "Resource ReSerVation Protocol (RSVP) — Version 1 Functional Specification," *RFC 2205*, Sept. 1997.
- [13] S. Blake et al., "An Architecture for Differentiated Services," *RFC 2475*, Dec. 1998.
- [14] Y. Bernet et al., "A Framework for Differentiated Services," *Internet draft, draft-ietf-diffserv-framework-00.txt*, May 1998.
- [15] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," *Internet draft, draft-ietf-mpls-arch-01.txt*, Mar. 1998.
- [16] D. Awduche et al., "Requirements for Traffic Engineering over MPLS," *Internet draft, draft-ietf-mpls-traffic-eng-00.txt*, Oct. 1998.
- [17] E. Crawley et al., "A Framework for QoS-based Routing in the Internet," *RFC 2386*, Aug. 1998.
- [18] S. Shenker, C. Partridge, and R. Guerin, "Specification of Guaranteed Quality of Service," *RFC 2212*, Sept. 1997.
- [19] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," *RFC 2211*, Sept. 1997.
- [20] R. Guerin, S. Blake, and S. Herzog, "Aggregating RSVP-based QoS Requests," *Internet draft, draft-guerin-aggreg-RSVP-00.txt*, Nov. 1997.
- [21] T. Li and Y. Rekhter, "Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)," *RFC 2430*, Oct. 1998.
- [22] D. Awduche et al., "Extension to RSVP for Traffic Engineering," *Internet draft, draft-swallow-mpls-RSVP-traffic-eng-00.txt*, Aug. 1998.
- [23] D. Ferrari and D. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE JSAC*, vol. 8, no. 3, Apr. 1990, pp. 368-79.
- [24] A. Banerjee and B. Mah, "The Real-Time Channel Administration Protocol," *Proc. 2nd Int'l. Wksp. Network and OS Support for Digital Audio and Video*, Heidelberg: Springer-Verlag, pp. 160-70.
- [25] D. Clark, "The Design Philosophy of the DARPA Internet Protocol," *Proc. ACM SIGCOMM '88*, Aug. 1988.
- [26] J. Postel, "Service Mappings," *RFC 795*, Sept. 1981.
- [27] K. Nichols et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," *RFC 2474*, Dec. 1998.
- [28] K. Nichols, V. Jacobson, and L. Zhang, "A Two-Bit Differentiated Services Architecture for the Internet," *Internet draft, draft-nichols-diff-svc-arch-00.txt*, Nov. 1997.
- [29] J. Heinanen et al., "Assured Forwarding PHB Group," *Internet draft, draft-ietf-diffserv-af-03.txt*, Nov. 1998.
- [30] Y. Bernet et al., "A Framework for use of RSVP with Diff-serv Networks," *Internet draft, draft-ietf-diffserv-rsvp-00.txt*, June 1998.

- [31] D. Clark and J. Wroclawski, "An Approach to Service Allocation in the Internet," *Internet draft, draft-clark-different-svc-alloc-00.txt*, July 1997.
- [32] B. Braden et al., "Recommendation on Queue Management and Congestion Avoidance in the Internet," *RFC 2309*, Apr. 1998.
- [33] T. Li, "CPE based VPNs using MPLS," *Internet draft, draft-li-MPLS-vpn-00.txt*, Oct. 1998.
- [34] H. Zhang, "Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks," *Proc. IEEE*, vol. 83, no. 10, Oct. 1995.
- [35] W. Yeong, T. Howes, and S. Kille, "Lightweight Directory Access Protocol," *RFC 1777*, Mar. 1995.
- [36] P. Vaananen and R. Ravikanth, "Framework for Traffic Management in MPLS Networks," *Internet draft, draft-vaananen-mpls-tm-framework-00.txt*, Mar. 1998.
- [37] L. Andersson et al., "Label Distribution Protocol," *Internet draft, draft-ietf-mpls-ldp-02.txt*, Nov. 1998.
- [38] M. Waldvogel et al., "Scalable High Speed IP Routing Lookups," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [39] S. Nilsson and G. Karlsson, "Fast Address Lookup for Internet Routers," *Proc. ACM SIGCOMM '97*, Cannes, France, Sept. 1997; <http://www.acm.org/sigcomm/sigcomm97>.
- [40] J. Moy, "OSPF Version 2," *RFC 2178*, Apr. 1998.
- [41] C. Villamizar and T. Li, "IS-IS Optimized Multipath (IS-IS OMP)," *Internet draft, draft-villamizar-isis-omp-00.txt*, Oct. 1998.
- [42] R. Guerin et al., "QoS Routing Mechanisms and OSPF extensions," *Internet draft, draft-guerin-QoS-routing-ospf-03.txt*, Jan. 1998.
- [43] Z. Zhang et al., "QoS Extensions to OSPF," *Internet draft, draft-zhang-aps-ospf-01*, Sept. 1997.
- [44] A. Orda, "Routing with End-to-End QoS Guarantees in Broadband Networks," *Tech. rep.*, Technion, Israel.
- [45] Q. Ma, "QoS Routing in the Integrated Services networks," Ph.D. thesis, CMU-CS-98-138, Jan. 1998.
- [46] Y. Goto, M. Ohta, and K. Araki, "Path QoS Collection for Stable Hop-by-hop QoS Routing," *Proc. INET '97*, Kuala Lumpur, Malaysia, June 1997.
- [47] Z. Wang and J. Crowcroft, "Quality of Service Routing for Supporting Multimedia Applications," *IEEE JSAC*, Sept. 1996.
- [48] A. Parekh, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks," Ph.D. thesis, MIT, Feb. 1992.
- [49] C. Partridge, *Gigabit Networking*, Addison-Wesley, 1994.
- [50] G. Apostolopoulos et al., "Quality of Service Based Routing: A Performance Perspective," *Proc. ACM SIGCOMM '98*, Vancouver, Canada, Aug. 1998, pp. 17-28.
- [51] F. Kelly, "Notes on Effective Bandwidths," in *Stochastic Networks: Theory and Applications*, F. Kelly, S. Zachary, and I. B. Ziedins, Eds., Oxford Univ. Press, 1996, pp. 141-68.
- [52] F. Kelly, "Modeling Communication Networks, Present and Future," *Phil. Trans. Royal Soc. A354*, 1996, pp. 437-63.
- [53] G. R. Ash et al., "Real-Time Network Routing in a Dynamic Class-of-Service Network," *Proc. ITC 13*, Copenhagen, Denmark, June 1991.
- [54] ATM Forum PN11 subworking group, "Private Network-Network Interface Spec. v1.0 (PN11.1.0)," *afpn11-0055.00*, Mar. 1996.
- [55] Juniper Whitepaper, "Optimizing Routing Software for Reliable Internet Growth," <http://www.juniper.net/leadingedge/whitepapers/optimizing-routing-sw.fm.html>.

Additional Reading

- [1] K. Nichols et al., "Differentiated Services Operational Model and Definitions," *Internet draft, draft-nichols-dsopdef-00.txt*, Feb. 1998.
- [2] Z. Wang, "Routing and Congestion Control in Datagram Networks," Ph.D. thesis, Dept. of Comp. Sci., Univ. College London, Jan. 1992.

Biographies

XIPENG XIAO (xiaoxipe@cse.msu.edu) is a Ph.D. candidate of the Dept. of Computer Science and Engineering, Michigan State University. He got his B.S. and M.S. degrees in computer science in 1992 and 1995, respectively, from Zhejiang University, China. His research interests include gigabit routing and router architecture, QoS, QoS, and traffic engineering. He won the First Prize in the ACM 1998 Graduate Student Research Poster Contest with the project "The Design of a Scalable Gigabit IP Router." He worked at the Multi-Gigabit Routing Division of Ascend Communications from May 1998 to Dec. 1998. He is now a senior network traffic analysis engineer at Frontier GlobalCenter Inc.


LIONEL M. NI (ni@cse.msu.edu) earned his B.S. degree in electrical engineering from National Taiwan University in 1973 and his Ph.D. degree in electrical engineering from Purdue University, West Lafayette, Indiana, in 1980. He is a Fellow of IEEE and a professor in the Computer Science and Engineering Department at Michigan State University. He has published over 160 technical articles in refereed journals and proceedings in the areas of high-speed networks and high-performance computer systems. He has served as an editor of many professional journals. He was program director of the NSF Microelectronic Systems Architecture Program from 1995 to 1996. He has served as program chair or general chair of many professional conferences. He has received a number of awards for authoring outstanding papers and also won the Michigan State University Distinguished Faculty Award in 1994.

APPENDIX 1007A

Journal of High Speed Networks

Volume 7 Numbers 3,4 1998

SPECIAL ISSUE:
Multimedia Networking



LAST NUMBER OF THIS VOLUME

Editor-in-chief: Prof. D. Sidhu

JOURNAL OF HIGH SPEED NETWORKS

Editor-in-Chief

Prof. Deepinder Sidhu
Maryland Center for
Telecommunications Research
University of Maryland-Baltimore
County
Baltimore, MD 21228-5398
USA
Tel.: +1 410 455 3028
Fax: +1 410 455 3969
E-mail: sidhu@cs.umbc.edu or
jhsn@cs.umbc.edu
www: <http://www.cs.umbc.edu/~jhsn>

Editorial Board

A. Acampora
University of California, USA

S. Banerjee
Stevens Institute of Technology, USA

F.-S. Choa
University of Maryland-BC, USA

A. Chung
DePaul University, USA

R. Cruz
University of California, USA

J.J. Garcia-Luna-Aceves
University of California, USA

R. Guérin
IBM T.J. Watson Research Center, USA

Z. Haas
Cornell University, USA

R. Jain
Ohio State University, USA

C. Johnston
NEC USA-C&C Research
Laboratories, USA

S. Keshav
Cornell University, USA

L. Kleinrock
University of California, USA

A. Krishna
IBM T.J. Watson Research Center
USA

S. Kumar
Northwestern University, USA

L. Landweber
University of Wisconsin, USA

B. Makrucki
IBM, USA

D. Mitra
AT&T Bell Laboratories, USA

B. Mukherjee
University of California, USA

S. Pink
SICS, Sweden

C. Qiao
State University of New York, USA

N. Ransom
Bell South Telecommunications, USA

J.M. Smith
University of Pennsylvania, USA

O. Spaniol
Technical University Aachen
Germany

M. Steenstrup
BBN Systems and Technologies, USA

J. Sterbenz
GTE Laboratories, USA

F. Tobagi
Stanford University, USA

S.K. Tripathi
University of California, USA

A. Varma
University of California, USA

R. Watson
Lawrence Livermore National
Laboratory, USA

I. Weerakoon
Hughes Network Systems, USA

E.W. Zegura
Georgia Tech, USA

Aims and scope

The *Journal of High Speed Networks* will serve as an archive for papers describing original results of lasting significance in both the theory and practice of high speed networks. The Journal publishes contributed papers on original research, survey papers on topics of current interest, technical notes, and short communications to report progress on long-term projects. Submissions to the Journal will be refereed consistent with the review process of leading technical journals. The Journal is published quarterly.

The Journal encourages submissions on all aspects of high speed networking. The topics covered will include, but will not be limited to: the design and analysis of high speed networks; protocols for high speed communications; high speed LANs, MANs, and WANs; broadband ISDN and services; optical networks; control and management of high speed networks; high speed switches, interfaces and controllers; routing, flow control and congestion control in high speed networks; formal approaches to high speed protocol development; network management; multimedia applications; applications of high speed networks; experiences with operational high speed networks.

©1998 IOS Press. All rights reserved

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher, IOS Press, Van Diemenstraat 94, 1013 CN Amsterdam, The Netherlands. No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, instructions or ideas contained in the material herein. Although all advertising material is expected to conform to ethical standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

Special regulations for readers in the USA. This journal has been registered with the Copyright Clearance Center, Inc. Consent is given for copying of articles for personal or internal use, or for the personal use of specific clients. This consent is given on the condition that the copier pays through the Center the per-copy fee stated in the code on the first page of each article for copying beyond that permitted by Sections 107 or 108 of the U.S. Copyright Law. The appropriate fee should be forwarded with a copy of the first page of the article to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA. If no code appears in an article, the author has not given broad consent to copy and permission to copy must be obtained directly from the author. This consent does not extend to other kinds of copying, such as for general distribution, resale, advertising and promotion purposes, or for creating new collective works. Special written permission must be obtained from the publisher for such copying.

Published quarterly

0926-6801/98/\$8.00

Printed in the Netherlands

Guest-editorial

UNIVERSITY OF
MICHIGAN

MAR 18 1999

MEDIA UNION
LIBRARY

Multimedia Networking

For this issue on 'Multimedia Networking', we received over twenty-five papers, of which eight were selected in the first round of review. The revised papers went through a second of review, after which seven were selected for publication. The papers selected for this issue identify new directions and challenges and provide a broad view spanning large number of issues that are of interest to the multimedia networking community.

The paper by Mukherjee, Reininger and Sengupta proposes VBR⁺ service class that provides all the functionalities of the ATM VBR service class, with the added capability of dynamic resource renegotiation. The paper by Schelén and Pink presents an advanced reservation scheme for network resources. Nahrstedt, Chu and Narayan propose QoS-aware resource management at the end-host with the applications running the show! The paper by Smith, Mutka and Rover presents a feedback rate control scheme for multimedia conferencing over the Internet. Feng, Lam and Liu propose a bandwidth smoothing algorithm that can effectively reduce the network resource requirements for the delivery of compressed video streams. The paper by Kravets, Calvert and Schwan proposes a communication layer for distributed interactive applications that adapts itself based on both application requirements as well as network resource availability. Malan, Jahanian and Subramanian propose an attribute-based routing substrate for push applications.

As you will see, the papers selected for this issue propose innovative solutions to some of the technical obstacles facing the multimedia entablement of the Internet. Although, many problems still remain, we believe that the seven papers in this issue will take us a step forward in the right direction.

Enjoy!

Debanjan Saha
IBM T.J. Watson Research Center
Hawthorne, NY 10532
USA
Tel.: +1 914 784 7194
Fax: +1 914 784 6205
E-mail: debanjan@watson.ibm.com

Satish Tripathi
Bourns College of Engineering
University of California
Riverside, CA 92521
USA
Tel.: +1 909 787 6374
Fax: +1 909 787 3188
E-mail: tripathi@engr.ucr.edu

Resource sharing in advance reservation agents

Olov Schelén* and Stephen Pink

Computer Science and Electrical Engineering, Luleå University of Technology, SE-971 87 Luleå, Sweden
Tel.: +46 920 91000; Fax: +46 920 72801; E-mail: {olov,steve}@cdt.luth.se

Abstract. We present an architecture where clients can make end-to-end resource reservations through agents. For each domain in the network, there is an agent responsible for immediate and advance admission control. Reservations from different sources to the same destination domain are aggregated as their paths merge toward the destination.

We show that network resources can be shared between immediate and advance reservations without being pre-partitioned. Admission control for immediate reservations use information about resources to be allocated for advance reservations in the near future. An important parameter in the admission control algorithm is the so called *lookahead time*, i.e., the point at which we actually start making resources available for approaching advance reservations by rejecting immediate requests. In our model, preemption of immediate reservations is made in cases where the admission control cannot make resources available through rejection of immediate requests. The risk of preemption can be varied by changing the lookahead time. We explore, with simulations, the effects of providing advance reservations with this model. The results show the cost in terms of resource utilization, rejection probability and preemption probability.

Keywords: QoS, advance reservations, resource sharing, agents, admission control, simulation, virtual leased lines

1. Introduction

Real-time applications such as audio and video conferencing may need resource reservations in the network to perform well. However, many of the current resource reservation proposals allow only immediate reservations, i.e., reservations at the time the session begins, while real-time events often are scheduled and advertised in advance. Clients making immediate reservations for scheduled events will have to judge the appropriate time for making reservation requests to be accepted in due time. When resources are scarce, clients may start making reservations earlier than needed to make sure that they obtain resources in due time. This leads to sub-optimal scheduling of network resources. By offering reservations in advance, clients obtain better support for planned activities and providers obtain information for scheduling resources over time.

Support for advance reservations is also important between different providers. For scalability reasons, it is not desirable that each end-to-end admission request results in call-specific state for each link along the path. Therefore, providers will negotiate bulk resources among each other for carrying aggregated admission-controlled traffic over several domains. In this context, advance reservations are suitable both for providing time-limited contracts and for providing bandwidth that varies at given time intervals, e.g., over hours of the day and days of the week.

We have designed a reservation architecture where immediate and advance admission control is performed by *agents* [15,17]. The provided service corresponds to *virtual leased lines* between domains. An advantage with an agent-based approach is that reservations can be provided without having to maintain reservation state in routers. Reservations can be made for both senders and receivers. In addition, reservations can be made for remote locations, thereby supporting third party reservations and nomadic computing, i.e., reservations for places where the host will be moving in the future. Our architecture can aggregate admission requests along the path towards a

*Corresponding author.

destination. This results in admission control for the aggregate being performed for individual links. In Section 2, we give the background on our agent architecture.

The focus of this paper is on resource management for individual links. In Section 3, we describe admission control for immediate and advance reservations [16]. The objective is to share resources between immediate and advance reservations. In Section 3.1, we describe some architectural considerations about locating admission control in agents. In Section 3.2, our service model is described. In this model resources for advance reservations will be made available by *rejecting* immediate requests for a short period before resources are to be allocated, or ultimately by *preempting* service for flows that were granted immediate access. Preemption means that a flow of packets using the reservation loses its service quality, i.e., the flow will still be given best-effort service (preemption is defined in Section 3.2). In Section 3.3, we define the notions *bookahead* and *lookahead*. In Section 3.4, we define the admission control criteria. In Section 4 we show, with simulations, the effects of advance reservations in terms of resource utilization, rejection probability and preemption probability.

2. The agent architecture

In [17] and [15], we present an architecture where clients can make resource reservations through *agents*. For each domain in the network there is an agent responsible for *admission control* (Fig. 1). Each agent knows the topology and static link resources in its domain. The agent is an end-system that is configured for passively participating in a link state routing protocol (e.g., OSPF) where each participating router has an identical topological database over the domain. Thus, an agent obtains the topological database with little signaling overhead. Link state advertisements are sent from the nearest router to the agent. Agents retrieve link properties, such as static bandwidths, by querying routers seen in the topology map. For this, we use a network management protocol (e.g., SNMP). Queries are done at startup and when topology changes are detected by the routing protocol. Routers need no extensions to allow agents to build a resource map through the management protocol.

Packets using reserved resources are *marked* by applications or edge routers. Core routers handle marked packets by classifying them into a small number of predefined service types. Similar ideas have been proposed in the *differentiated services* area [2,12]. In the differentiated services model, packets are marked to obtain either lower drop probability or higher strict priority. Here, we abstract from these details and use the term *priority packets* to denote marked packets, and the term *priority traffic* for all priority packets collectively. A key point of differentiated services and our agent architecture is to avoid signaling state and expensive per-packet processing in routers. To allow resources to be managed in separate agents, we rely on router mechanisms for service differentiation on the packet level.

Independently of which packet scheduling model is used for differentiating service in routers, there must be an admission control architecture to make sure that the amount of priority traffic is low enough to receive good quality of service. One issue is whether the admission control architecture should reserve resources over the links that are actually going to be used, or just limit the overall rights to send priority traffic. A scheme that does not consider




Fig. 1. Reservation agents and their domains.

exactly where traffic will go cannot give reliable commitments without very low utilization. Thus, it is likely that such a scheme would need to gamble and accept occasional service loss.

In this work we focus on agent-based admission control where the objective is to administer resources exactly, i.e., admission control in an agent maintains information about reserved resources on each link in its domain. Agents perform immediate or advance admission control without involving the routers. The architecture offers a service comparable to unidirectional *virtual leased lines*. Reserved resources can be used by one or many applications sending data between the domains connected by a virtual leased line.

There are other frameworks, i.e., IETF Integrated services [7,20] and RSVP [21], that can set up both unicast and multicast reservations end-to-end. However, the provided service is not the same as for virtual leased lines, since reservations are made for individual data flows (i.e., application data streams). In that model, per-flow reservations are maintained in the routers. This results in scaling problems due to reservation state and packet classification overhead in routers. This model inherently relies on per-flow semantics and therefore it is hard to perform aggregation over different administrative domains. Thus, RSVP and Integrated Services may be too heavy-weight for the backbones.

In differentiated service, on the other hand, there is low overhead in routers. Our goal is to design an architecture that performs scalable resource management. Currently, our service model provides only unicast virtual leased links. When resources are managed by agents according to this service model, there are opportunities to aggregate reservations as they merge toward a destination domain. In this section, we give an overview of the admission control architecture and focus on the aggregation issues.

2.1. Direction of admission requests

An admission request can be directed to any agent, e.g., an agent that manages a user's account. Each agent sets up reservations between any two points in the network by invoking other agents. Thus there is *no difference* between reservations for *sending* and *receiving* data, or reservations made by a *third-party*.

Each reservation request contains a bandwidth to be reserved, a source and a destination address. The source address determines the point where the reservation starts (i.e., the point where traffic using the reservation will enter) and the destination address determines the point where it ends. An agent receiving a request first considers whether the starting point is in its domain. If the starting point is *not* in its domain, it finds an agent closer to the starting point and repeats the request with that agent. Technically speaking, it is possible to directly identify the agent that is responsible for the source domain, but in practice clients may prefer using a particular agent. Also, agents may prefer having business relations with only a few adjacent agents. This would require that requests are made in several steps. The result is that accounting information is aggregated, i.e., the mechanism helps keep down the number of separate charges. In Fig. 2, it is shown how a request is issued by node Dx for reserving resources from Ax to Dx.

If an agent finds that the starting point given in the request is in its domain, admission control is performed on the links from the starting point towards the destination, defined by the routing protocol. This is possible since all routers and the agent have the same topology map of the domain. Thus, an agent can perform admission control on the links from the starting point to an edge router in its domain.

An agent cannot perform admission control beyond its domain. Therefore, if the destination is outside of the agent's domain it must request a reservation with the neighboring agent, giving the edge point where traffic will cross the borders as the source address of the reservation (Fig. 2). Without distinguishing between reservations to different destinations, it would be impossible for agents to grant reservations to other domains. This is because there is no way to make sure that there are sufficient reservations along any path in other domains further downstream (unless we accept very low utilization for priority traffic).

2.2. Aggregation

If reservations were always done 'on demand' there would be signaling between all neighboring agents along the path for each end-to-end admission request. Our architecture supports several ways of aggregating reservation




Fig. 2. A receiver Dx requesting a reservation to receive data from Ax.

state. Aggregation takes place in the agents because routers are not part of the admission control process. The motivation for aggregation may be less when reservation state is stored in agents with plenty of memory and disk. However, a per-flow reservation model would have scaling problems both in terms of state and processing cost in agents. Therefore, it is important to show that agents can provide exact resource reservations with good scaling properties.

In [17] we describe how aggregation is done. The key ideas are the following: first, reservations from different sources to the same destination are aggregated as their paths merge toward the destination. Second, an agent in charge of resources at the final destination can generalize reservations for specific end points so that they are valid for any end point in the destination domain, thereby allowing more aggregation.

2.2.1. Merging reservations

To grant requests spanning many routing domains, there must be a commitment from all agents along the path involved in the request. When an agent is making a reservation with a neighboring agent, the source address included in the request determines the point (router interface) where the traffic will cross the borders, i.e., enter the domain of the neighboring agent. The neighboring agent does not care where traffic came from originally. Thus, priority packets can use reserved resources as long as the destination address matches the reservation. This allows agents to aggregate traffic with different sources into one single reservation made with a neighboring agent.

An end-to-end resource reservation can be seen as a number of consecutive *funnels* (Fig. 3). Priority packets poured into any of these consecutive funnels use reserved resources along the path to the destination. Agents may aggregate several requests with different source addresses into the same funnel, as long as all the requests specify the same destination. Agent C (in Fig. 3) sets up a funnel starting at its edge by reserving bulk resources (40 units) with neighboring agent D to a distant destination and then aggregate requests for resources between domains C and D into the established funnel. This bulk reservation implies that downstream agents only keep state for the aggregate. The agent that decides to aggregate into a bulk reservation keeps information about its individual commitments constituting the aggregate. Aggregation by merging reservations towards each destination can be done with full control over the resources. The funnels form a so called *sink-tree* with the root towards the destination. At each merging point, the reserved bandwidth on the outgoing link will be equal to, or larger than, the sum of the reserved bandwidth on the incoming links.

2.2.2. Prefix aggregation

Requests that specify different destinations cannot easily be aggregated into a funnel in the general case. If we allow funnels to split into different directions further downstream, upstream agents must keep information about which destinations are involved and how resources are divided between them. This means that the aggregate no longer can be seen as one unit. However, if a set of destinations can be identified by one common identifier, and




Fig. 3. Funnels and aggregate reservations for one destination.




Fig. 4. Funnels and generalized destinations.

there are sufficient resources available in each branch to service the whole aggregate, then split outs can be scalable and exact.

With this argument, we allow aggregation to be increased by *generalizing* a reservation as valid for all destinations within the same destination domain. This means that a funnel provides resources to all endpoints within the destination domain, i.e., resources in the destination domain must be sufficient to ensure that priority packets get expected service independently of how they are routed within that domain. It is only the agent responsible for the destination domain that can judge if resources are sufficient for allowing generalization. Consider the example in Fig. 4. Agent D can judge if the resources in D are sufficient to let incoming funnels be valid for the whole domain. Thus, it must be agent D that decides when a reservation request for a particular destination node within D can be replaced with a reservation for any node in domain D.

In IP networks an *address prefix* is often used to represent a domain (individual nodes within the domain are distinguished by the remaining suffix). An address prefix is represented by an IP address and an associated mask telling how many bits are part of that prefix. An agent generalizes a request by replying with an *address prefix* to be used as a destination address for the reservation instead of the original address included in the request (Fig. 4). The aggregation that results from generalizing a reservation is also called *prefix aggregation*.

Clients or agents that have been granted a generalized reservation can use it for any destination matching the prefix of the reservation. Agents providing generalized reservations must handle the committed bandwidth of priority traffic to any destination in its domain. A conservative agent would therefore generalize reservations up to an aggregated value equal to the narrowest link. In practice we believe that this can be useful for well provisioned local area networks.

In the Internet, agents must be able to handle occasional topology changes (e.g., node and link failures), as well as pathological changes (e.g., route-flaps) [13]. Agents listen to the routing protocol and detect routing changes on average as fast as any router in the domain. When changes occur, agents can schedule resources along new paths.

3. Admission control for immediate and advance reservations

To obtain a modular design, advance reservations should, as far as possible, be managed separately from immediate reservations. One way of separating the admission control for immediate and advance reservations would be to partition the resources so that one part is available for immediate reservations and another part for advance reservations. However, this would result in two logically separate networks with the problem of partitioning resources according to current demand. Therefore, we have chosen to share resources between immediate and advance admission control. To allocate resources safely we have extended the immediate admission control to consider the near future concerning the resources that will be allocated for advance reservations. There are three major consequences of these results. First, the resources can be managed efficiently without suffering from badly chosen resource partition sizes. Second, the immediate admission control scheme can be kept fast and simple by not having to consider advance reservations in the far future. Third, we have the architectural freedom of locating the admission control for immediate and advance reservations at separate places in the network without suffering undue communication overhead.

3.1. Parameter-based or measurement-based?

Admission control algorithms can be classified as either parameter-based or measurement-based. The metrics used in *parameter-based* admission control are *traffic specifications* and static link resources. By aggregating the given metrics, such a scheme can offer absolute or statistical performance bounds [6,7], often enforced by the packet scheduler (e.g., WFQ or EDF). Parameter-based admission control can be used for immediate reservations and, by adding timing information, for advance reservations.

The metrics used in *measurement-based* admission control involve also *measurements* of used resources [10,20]. Such a scheme can detect reserved resources that are not utilized and grant those resources to other requests. The benefit of measurement-based admission control is better resource utilization. On the other hand, there is no guarantee that reserved resources are available when needed. The scheme may also be sensitive to synchronization effects. Suppose there are many events scheduled at a given time and many clients start reserving immediate resources just before that time. As there is little traffic, more reservations will be allowed. When all events start, the network will be over-allocated with poor service as a result.

In our agent scheme, we use only parameter-based admission control for a number of reasons. First, measurement-based admission control can be efficiently implemented only in routers or switches (where traffic measurement can be done directly). When admission control is in the agent, the signaling overhead for traffic measurements may be prohibitive. Second, traffic measurements provide little information about the future and are therefore not suitable for advance reservations [4]. Third, we want to provide a service model where reservations give a trusted level of service for people that are prepared to pay for it. The problems of lower utilization for reserved traffic are mitigated in a packet-switched network as unused resources always can be used by best-effort traffic. Thus the overall network utilization does not suffer very much from parameter-based admission control.

3.2. Service model

To support advance reservations, information about time must be added to the reservation requests. Therefore, advance reservation requests include a *duration* interval: $I = [t_s, t_e]$, where t_s is the starting time of the requested service and t_e is the finishing time of that service. This information could be automatically collected from a site where scheduled events are advertised, or generated by the client making a reservation. This is indifferent to our architecture.

Immediate reservations are open ended and do not specify a duration. A consequence of allowing open ended immediate reservations together with advance reservations is that there may be cases where resources are over-allocated. This is resolved by *preempting* immediate reservations to free resources for advance reservations. Preemption is used to ensure that there are always resources to honor advance reservations. Flows that lose reserved resources through preemption will be given best-effort service. We have chosen *not* to include a service for immediate requests that are *guaranteed* never to be preempted. Instead, we aim at keeping the preemption probability low, so that the user need not worry about preemption when asking for admission. A fundamental reason for this decision is that indefinite non-preemptible service would block resources indefinitely and no advance reservations of those resources could, therefore, be made. Furthermore, resources that are reserved in advance for some time in the far future cannot be given to indefinite sessions in the meantime unless we have the option of preempting those sessions when resources are needed. Thus, we have the following service model:

- Durations must be specified for advance reservations. A duration includes the start time and the finish time for the requested service.
- Immediate reservations do not specify durations. Such requests are serviced immediately (if possible) and can be preempted later by flows that reserved in advance.
- Preemption means that a flow loses its service quality; however, it will still be serviced at a best-effort level.
- In order to manage resources safely there may be limits imposed on minimum and maximum bookahead times for advance reservations.
- The admission control for advance reservations may support finding the earliest point for which a requested service can be granted.

There are other proposals for advance reservations that do allow indefinite (open ended) resource reservations. To provide indefinite sessions without having preemption, resources could be set aside for this service exclusively. In the Tenet suite [5], resources are partitioned to support advance reservations and non-preemptible immediate sessions simultaneously. With that solution there will be some fragmentation due to it being hard to maintain the partitioning such that the the demand for the two different services are met at each instant. In our architecture we can support indefinite sessions by partitioning resources. However, in this paper, we will investigate whether or not we can avoid some problems with resource fragmentation by sharing resources between immediate and advance reservations and instead allow for statistically rare cases of preemption.

There are also proposals relying on statistical guarantees of immediate and advance reservations where both immediate and advance reservations are preemptible. In [18] a model is presented where the the call with the shortest bookahead time is preempted in case of congestion. The model relies on the theory of effective bandwidth to calculate the chance of a call being active. In our model we deal with aggregated bulk reservations rather than per-call reservations. For our model have chosen to use deterministic guarantees in the admission control for advance reservations. Our belief is that advance reservations, once accepted, should not be preempted. Preemption of advance reservations is acceptable only when there are failures in the basic infrastructure. This consideration is made in a context where there is a lot of unplanned network traffic that use immediate reservations and an even larger amount of bursty traffic that uses no reservation at all. In this context resources will be well utilized.

3.3. Bookahead and lookahead

When immediate reservations and advance reservations share resources and the duration of each immediate reservation is unknown, there is a risk of over-allocation that can be resolved with preemption. The probability of

preemption can be controlled by either monitoring immediate requests when admitting advance reservations, or vice versa, by monitoring advance reservations when admitting immediate requests.

In practice, however, the risk of preemption for immediate requests can not be controlled very well by advance admission control. When advance reservations are requested far in advance, the information obtained from monitoring present traffic or present immediate reservations will not be useful [4]. Most of the active flows would have finished at the crucial time anyway. It is only when advance reservations are requested with very short notice that it would be valuable to know about present traffic. Since advance reservations are provided to meet the demand for scheduled sessions primarily, we believe that most reservations will be requested on a notice much longer than the average session length. We denote the time from making an advance reservation until the resources are to be available as the *bookahead time*. The bookahead time is set by the clients asking for admission and varies from request to request. In our admission control we may impose a minimum bookahead time. Then, advance admission control need not consider the immediate reservations state at all.

The risk of preemption can be effectively controlled by immediate admission control. By monitoring advance reservations for the near future when admitting immediate requests, we can control the risk of preemption. If there are insufficient resources available in the near future, the request could be rejected instead of serviced. It is up to the provider to decide whether rejection or preemption is preferable in these cases. Rejecting immediate requests because the resources are reserved in the near future means that the overall utilization will go down to make room for the advance reservations. We denote the pre-allocation time, i.e., the time for starting to set aside resources for advance reservations as the *lookahead time* into the advance reservation state. Setting the lookahead time is an internal issue to the admission control scheme. As shown in simulations (Section 4), the lookahead time can be tuned to obtain reasonable utilization and preemption probability.

3.4. Admission control criteria

Figure 5 is a snap-shot of admitted flows in a time/bandwidth diagram. Flows a, b and c are immediate reservations. For immediate reservations we add up the aggregate bandwidth. When a new admission request α arrives, admission is granted if the new reservation would not cause the aggregate bandwidth to exceed the link bandwidth. For advance reservations, the admission conditions consider other advance reservations that overlap with the new flow (d, e, g, h). The conditions are checked at all points where new flows begin (t_s , t_x and t_y). For immediate requests, the admission conditions consider the current aggregate of immediate reservations and the advance reservations that are active within the lookahead time (d, e, f). The following notation is used in describing our admission criteria:




Fig. 5. Snap-shot of reservations.

R	Aggregate bandwidth reserved by current immediate reservations,
$A(t)$	Aggregate bandwidth reserved in advance at time t ,
$look$	The lookahead time used at immediate admission,
μ	The link bandwidth.

Immediate reservations. When a client requests an immediate reservation α with bandwidth r^α , the admission control algorithm performs the following checks:

- Determine if the bandwidth usage, after adding the new load r^α , will exceed the available link bandwidth μ during the lookahead time

$$\mu > \max_{t \in look} (r^\alpha + R + A(t)). \quad (1)$$

Advance reservations. When a client requests an advance reservation α with bandwidth and duration (r^α, I^α) , the admission control algorithm performs the following checks:

- Check whether the starting time is acceptable as a bookahead time. The minimum bookahead time is configurable to any value larger than the lookahead time.
- Check whether the finishing time is acceptable. It should be larger than the starting time. Admission control may also set a limit on how far in the future reservations can be granted.
- Determine whether the aggregate bandwidth of advance reservations exceeds the link bandwidth at some point during the requested duration:

$$\mu > \max_{t \in I^\alpha} (r^\alpha + A(t)). \quad (2)$$

3.5. Reservation state

As described in Section 2.2, our model has mechanisms to minimize the amount of reservation state. Nevertheless, it is important to have data structures that manage reservations efficiently and allow quick admission control. For immediate admission control, we need to keep information about the current reserved aggregate bandwidth. We also need access to the advance reservation state within the lookahead time. For advance admission control, we need to keep the aggregated bandwidth reserved in advance over time. For this we have two alternative data structures. Our first data structure keeps exact starting times and finishing times for each admission request. The amount of state is proportional to the number of reservations maintained. The other data structure uses time slots where each reservation is extended to occupy a number of time slots. For each slot, the aggregate reserved bandwidth can be found. This data structure has fixed size and depends only on the maximum bookahead time allowed and the slot length.

Also, we must be able to find a victim for preemption among the immediate reservations. For this we need information about individual immediate reservations. However, potential victims can be chosen when they are admitted. It is enough to record a small number of potential victims corresponding to the amount of resources to preempt.

4. Simulations

The objective of our simulations is to explore the effects of allowing reservations in advance. The performance measures are *total utilization*, *rejection probability*, and *preemption probability*. We vary the *fraction* of the total link resources used for *advance reservations* and the *lookahead time* into advance reservation state at admission control for immediate requests. Admission decisions are parameter-based, i.e., based on traffic specifications and static link resources. For simplicity we use the requested bandwidth as the only traffic parameter. This study is for admission control over a single link only.

4.1. Simulation parameters

- The duration of immediate requests, i.e., the call hold time, is exponentially distributed with a mean of 300 time units.
- The duration of advance requests, is also exponentially distributed with a mean of 300 time units.
- The bookahead time for advance requests, i.e., how far in advance the advance requests are made, is exponentially distributed with a mean of 20 000 time units, where the tail is cut at 60 000 units (i.e., there is a maximum bookahead time).
- The total link capacity is 45 000 resource units.
- The resources asked for in each request are uniformly distributed in [1 . . . 3000] resource units.
- The inter-arrival time of admission requests is exponentially distributed with a mean of 11.76 time units. This number was selected to obtain a reasonable rejection probability, where around 6% of the requests are rejected. The average offered load is 85% of total capacity.
- When preemption is performed, the immediate request with the largest requested capacity is picked as a victim (which would encourage making requests in advance when large capacity is needed).

Note that the same parameters for duration and resource requirements are used for both immediate and advance reservations. This is to isolate the impact of reserving in advance versus reserving immediately. Should we have different parameters for the two classes, other effects would dominate the outcome in terms of utilization, rejection probability and preemption probability (the impact of having other parameters is discussed in Section 4.6). The bookahead time is chosen to be sufficiently large so it has little impact on the currently active immediate reservations.

The simulations are run many times for each parameter set, with different seeds for the random variables. The 90% confidential intervals are computed and presented as error bars in the plots (although the confidence is often so tight that the error bars appear as dots).

4.2. Advance reservation requests

We vary the fraction of the admission requests that are made in advance between zero and one. The upper bound on resources available for advance reservations equals the total link capacity, resulting in advance reservations being seldom rejected when a small fraction of the requests are made in advance (Fig. 6).

Figure 7 shows that the average duration of the granted requests goes down as there is a higher fraction of advance reservations. When resources are scarce, advance reservations will result in some fragmentation in time, and there is a greater chance that requests for short durations are accepted. This effect does not occur for immediate requests since admission is made without considering durations (in fact, durations are unknown).

A similar effect can be observed concerning the average bandwidth given to the granted requests. If the offered load is increased, the average bandwidth for granted requests will go down as the rejection probability goes up. For immediate reservations clients see a considerable decrease in the chance of succeeding with large reservations as the offered load is increased. However, for advance reservations the rejection probability depends on the bookahead




Fig. 6. Rejection probability for advance requests.




Fig. 7. Average duration for accepted advance requests.

time. Therefore, clients can increase their chance of succeeding with large requests by reserving further in advance. This effect is clearly seen in our simulations. However, as the bookahead time is chosen by clients it is not easy to draw any conclusions on *how* early one should make an admission request. The appropriate bookahead time depends on when other clients choose to book ahead. The only effect that can be claimed for sure is that the chance of having a request granted for a given duration increases strictly with the bookahead time.

4.3. Total utilization

We define the *total utilization* as being the ratio of the total reserved bandwidth versus the total bandwidth of the link. This metric considers how large part of the total resources that are reserved. The metric does not include increased link utilization due to best-effort traffic, nor drops in link utilization due to reserved resources not being utilized.

Figure 8 shows that the total utilization goes down as the fraction of advance reservations increases. There are two reasons for this. First, we start allocating resources ahead of the starting point of the durations for advance reservations. The interval during which resources are allocated ahead of time is denoted the *lookahead* time. Technically speaking the lookahead time is used by admission control for *immediate* requests to judge whether there are any advance reservations needing the resources within that lookahead time. If so, the immediate request will be rejected. Thus, having a longer lookahead time results in more immediate requests being rejected and thereby a lower utilization. Having a shorter lookahead time means that the network will more often be over-allocated, resulting in more immediate requests being preempted (as we will show in Section 4.4).

The second reason for lower utilization is that fragmentation in time is introduced when advance reservations are allowed. There is little chance that other advance requests ask for the small fragments that may be available between already-granted advance reservations. Consequently, a decrease in utilization can be observed as the rate of advance reservations is increased, even though the lookahead time is zero. When the lookahead is zero, admission control for immediate requests does not consider any resources reserved in advance.

This may sound like bad news for providing advance reservations, but in fact similar drops in total utilization would be observed as soon as there are scheduled events introduced in the network, even though users were to make only immediate reservations. Users would still be likely to request resources before events begin to make sure that they have resources allocated in time. The only difference is that the control is with the clients instead of with the network. The degree of over-allocation in time that clients choose would be a consequence of social behavior in an environment with scarce resources. In some situations we can expect that the utilization drops in an uncontrolled way. When there are scheduled events, advance reservations can improve the service and the utilization by ensuring that resources are allocated at an appropriate time.

Measurement-based immediate admission control could alleviate the problem of clients over-allocating in time by measuring traffic on the link. However, we cannot get any useful measurements for scheduled events until the traffic starts. Thus, we must choose between staying with the traffic specification or gradually losing the resource allocation as the measurement procedure finds that there is no traffic. When there are several events scheduled at the same time, there is a risk that allocated resources are insufficient to give acceptable service.

4.4. Rejection and preemption probability

Figure 9 shows that the rejection probability for immediate reservations increases as the fraction of advance reservations is increased. However, the change in rejection probability for immediate requests is primarily an effect of picking the victims among a smaller set of immediate requests as the rate of advance reservations increases. If we instead monitor the total rejection probability (Fig. 10) over both immediate and advance requests, we get a much lower risk of rejection (note that the scales are different). Perhaps this is a better reflection of overall user satisfaction.

We can see that the rejection probability is about the same when we have only immediate requests as it is when we have only advance requests, i.e., around 6%. The decrease in total utilization (Fig. 8) is explained by the fact that the average duration for accepted requests goes down (due to fragmentation in time) with increasing demand for advance reservations (Fig. 7).




Fig. 8. Total utilization.




Fig. 9. Rejection probability for immediate requests.




Fig. 10. Total rejection probability (for immediate and advance requests).




Fig. 11. Preemption probability for immediate requests.




Fig. 12. Total preemption probability (for immediate and advance requests).

Figure 11 shows that the preemption probability increases as the fraction of advance reservations is increased. Again, part of the explanation is that all preemption victims are picked within the immediate requests and as the fraction of immediate requests goes down we get a higher risk of preemption for this service class. In Fig. 12, we show the preemption probability over the total number of served reservations.

It is clear that increasing the lookahead time makes admission control more conservative, resulting in lower preemption probability at the cost of lower utilization and higher rejection probability. By decreasing the lookahead time, we allow the system to be more over-booked resulting in higher preemption probability. The fact that the system is allowed to be over-booked explains why the total rejection probability (Fig. 10) goes down when we use zero lookahead and there is a suitable fraction of advance reservations.

4.5. Lookahead time

Given a desirable target for preemption probability, the lookahead time can be conservatively computed by considering the bandwidth that must be allocated for each advance reservation and the amount of bandwidth that is returned per time-unit by immediate reservations ending. However, at reasonable load, it is likely that there are some unallocated resources that should be considered as well. In our simulations, between 75 and 79% of the resources are allocated on average, which explains why it is possible to allocate resources for advance reservations with a total preemption probability below 5%, even when using zero lookahead.

In previous examples we kept the lookahead times constant as the fraction of advance reservations was increased. We have also tried to use exponentially distributed lookahead times but we found a strict performance degradation compared with constant lookahead time. If there are problems in meeting the requirements for high utilization and low preemption probability, we may be forced to have a limit imposed on resources available for reservations in advance. As an example, we now fix the advance request rate to 40% of the total number of requests (i.e., the average is bounded and there is no firm bound) and we plot the performance measures as functions of the lookahead time.

Figures 13–15 show the effects. It is clear that the preemption probability falls steeply as the lookahead time is increased. In most systems, like cellular phone systems, preemption is considered bad and should be avoided. In a




Fig. 13. Rejection probability for immediate requests, at 40% advance request ratio.




Fig. 14. Preemption probability for immediate requests, at 40% advance request ratio.




Fig. 15. Utilization for immediate requests, at 40% advance request ratio.

connectionless packet switched network, preemption is not as bad since generally the packets are still served best-effort. We believe, however, that the relative weight of having low preemption is still high. Provided that we have the average duration for resource reservations (assuming exponentially distributed session lengths), the fraction of advance reservations, and the individual weights (set by the provider) for utilization, blocking and preemption, we can find a suitable lookahead time.

With a 40% advance reservation request rate, we can, for example, use a lookahead of 50 time units, resulting in 77.6% of the resources being allocated under a 1.1% immediate request preemption probability (i.e., 0.44% total preemption probability). By using a lookahead of 100 time units, we get a utilization of 76.9% at a preemption probability of 0.35% for immediate requests. This is to be compared with 79% of the resources being allocated when there are no advance reservations. The offered load is in both cases 85% of the total capacity and the total rejection probability is around 7%.

If there were reservation agents without supporting lookahead in the immediate admission control algorithm, the results would be comparable with having zero lookahead (the network must support preemption and non-blocking resource allocation that can be used by the advance reservation agent). In that case, we have 5% preemption probability for immediate requests when 40% of the admission requests are made in advance.

4.6. Evaluation

Our simulations show the cost of advance reservations when they use the same holding-times and have the same resource demands as immediate reservations. We do this to isolate the impact of reserving in advance versus reserving immediately. By using different parameters for advance and immediate reservations, utilization increases when there are more advance reservations. Given a certain offered load, it is easy to show that there is a benefit from making large reservations in advance. It is generally easier to obtain a large utilization if we start packing the largest requests into the time-bandwidth diagram shown in Fig. 5. In reality, we expect that users needing large resources will tend to make reservations in advance to avoid being blocked.

5. Conclusion

In this paper we claim that resources can be more accurately allocated over time by supporting advance reservations for scheduled events. The price of supporting advance admission control is a small but controlled decrease in network utilization. This is compared with users making immediate reservations for their scheduled events. When resources are scarce, users tend to ask for admission ahead of time to make sure that they have resources granted in due time for their own scheduled event. Users make their decision with little information about other scheduled events and with little concern about the overall utilization of the network. In some situations, the effective utilization and the probability of having reservation requests granted can drop without control.

Our simulations show the effects of providing advance reservations. We assume that we allocate resources at some point for each reservation and the utilization is measured in terms of how much we can allocate. We do not consider drops in link utilization due to reserved resources not being used.

In our model, resources are shared between advance and immediate admission control, i.e., there is no pre-partitioning of resources. The most important parameter in the admission control algorithm is the so-called *lookahead* time, i.e., the point at which we actually start making the resources in the routers available for approaching advance reservations. We show how a suitable lookahead time can be found through simulations.

The effects of advance reservations can be expressed in terms of rejection probability, preemption probability, and total utilization. In the simulations presented, we can allocate 79% of the link resources when there are no advance reservations and around 77% of the resources when 40% of the admission requests are made in advance, and there is a 0.4% preemption probability for immediate requests. The offered load is in both cases 85% of the total capacity and the total rejection probability is around 7%.

Acknowledgements

This work was supported by a grant from the Center for Distance Spanning Technology (CDT), Luleå, Sweden, and Ericsson Radio Systems, Sweden.

References

- [1] A. Banerjee, D. Ferrari, B.A. Mah and M. Moran, The tenet real-time protocol suite: design, implementation, and experiences, Technical Report, TR-94-059, University of California at Berkeley, Berkeley, California, November, 1994.
- [2] D.D. Clark and W. Fang, Explicit allocation of best effort packet delivery service, Technical Report, MIT Lab for Computer Science, Boston, Massachusetts, November, 1997.
- [3] D.D. Clark, S. Shenker and L. Zhang, Supporting real-time applications in an integrated services packet network: architecture and mechanism, in: *Proc. of SIGCOMM*, Baltimore, Maryland, 1992, pp. 14–26; ACM, *Computer Communication Review* 22(4).
- [4] M. Degermark, T. Köhler, S. Pink and O. Schelén, Advance reservations for predictive service, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 3–14.

- [5] D. Ferrari, A. Gupta and G. Ventre, Distributed advance reservation of real-time connections, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 15–26.
- [6] V. Firoiu, J. Kurose and D. Towsley, Efficient admission control for edf schedulers, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1996.
- [7] R. Guerin, C. Partridge and S. Shenker, Specification of guaranteed quality of service, Request for Comments (Proposed Standard) 2212, Internet Engineering Task Force, October, 1997.
- [8] A. Gupta, W. Howe, M. Moran and Q. Nguyen, Resource sharing for multi-party real-time communication, in: *Proc. of IEEE Infocom*, Boston, Massachusetts, 1995.
- [9] S. Herzog, Rsvp extensions for policy control, Internet Draft, Internet Engineering Task Force, November, 1996, Work in progress.
- [10] S. Jamin, P.B. Danzig, S. Shenker and L. Zhang, A measurement-based admission control algorithm for integrated services packet networks, in: *Proc. of SIGCOMM*, Cambridge, Massachusetts, 1995.
- [11] D.J. Mitzel, D. Estrin, S. Shenker and L. Zhang, A study of reservation dynamics in integrated services packet networks, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1996.
- [12] K. Nichols, V. Jacobson and L. Zhang, A two-bit differentiated services architecture for the Internet, Internet draft, Bay Networks, LBNL and UCLA, November, 1997.
- [13] V. Paxson, End-to-end routing behavior in the Internet, in: *Proc. of SIGCOMM*, Palo Alto, California, 1996, ACM, pp. 25–38.
- [14] W. Reinhardt, Advance reservation of network resources for multimedia applications, in: *2nd International Workshop on Advanced Tele-services and High Speed Communication Architectures (IWACA '94)*, Heidelberg, Germany, 1994.
- [15] O. Schelén and S. Pink, An agent-based architecture for advance reservations, in: *IEEE 22nd Annual Conference on Computer Networks (LCN '97)*, Minneapolis, Minnesota, 1997.
- [16] O. Schelén and S. Pink, Sharing resources through advance reservation agents, in: *Proc. of IFIP Fifth International Workshop on Quality of Service (IWQoS '97)*, New York, 1997.
- [17] O. Schelén and S. Pink, Aggregating resource reservations over multiple routin domains, in: *Proc. of IFIP Sixth International Workshop on Quality of Service (IWQoS '98)*, Napa Valley, California, 1998.
- [18] D. Wischnik and A. Greenberg, Admission control for booking ahead shared resources, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1998.
- [19] L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller and H. Wittig, Issues of reserving resources in advance, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 27–37.
- [20] J. Wroclawski, Specification of the controlled-load network element service, Request for Comments (Proposed Standard) 2211, Internet Engineering Task Force, October, 1997.
- [21] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, RSVP: a new resource ReSerVation protocol, *IEEE Network Magazine* 7(5) (1993), 8–18.

APPENDIX 1007B

UNIVERSITY OF MICHIGAN LIBRARY

EXPLORE ▾



SEARCH

[Account](#) [My Favorites](#) [Log in](#)[ANN ARBOR](#) [FLINT](#)

Keyword ▾

isn:27344595

[Advanced](#)

❖ **Search Tip:** Enter one or more keywords. Use quotes to search for a phrase (e.g. solar power; polar bears; "systems of oppression"). See tips about [Basic Keyword Searching](#).

[Everything](#) [Catalog](#) [Articles](#) [Databases](#) [Online Journals](#) [Guides and More](#)[Catalog](#) / Record[Serial](#) [Journal](#)

Journal of high speed networks. □

Record info:

Format

[Serial](#)[Journal](#)

Other Titles

[High speed networks](#)[JHSN](#)

Contributors

[OCLC FirstSearch Electronic Collections Online.](#)[IOS Press](#)

Published/Created

Amsterdam ; Washington, [D.C.] : IOS Press, 1992-

Language

English

Related Items

Electronic serial mode of access: World Wide Web.

Numbering

Vol. 1, no. 1 (1992)-

Current Publication Frequency

Quarterly

Numbering Peculiarities Note

Some issues called also: Special issue.



Ask a Librarian



Note

Title from cover.

Physical Description

v. : ill. ; 24 cm.

ISSN

0926-6801
1875-8940 (online)

OCLC Number

27344595

Subjects (LCSH)

[Telecommunication systems -- Periodicals.](#)
[Computer networks -- Periodicals.](#)

Academic Discipline

[Social Sciences](#) > [Information Science](#) > [Human Computer Interaction Engineering](#) > [Computer Science Engineering](#) > [Electrical Engineering](#)
[Social Sciences](#) > [Information Science](#) > [Human Computer Interaction](#)

The University of Michigan Library aims to describe library materials in a way that respects the people and communities who create, use, and are represented in our collections. Report harmful or offensive language in catalog records, finding aids, or elsewhere in our collections anonymously through our [metadata feedback form](#). More information at [Remediation of Harmful Language](#).

Available at:

 Online Resources · 2 items

 Offsite Shelving · 14 items

About location

- 1-18 : 1992-2012

Action	Description	Status	Call Number
Get This	v.1-2 1992-1993	✓ On shelf	TK 5101 .A1 J68
Get This	v.3 1994	✓ On shelf	TK 5101 .A1 J68
Get This	v.4 1995	✓ On shelf	TK 5101 .A1 J68
Get This	v.5 1996	✓ On shelf	TK 5101 .A1 J68
Get This	v.6-7 1997-1998	✓ On shelf	TK 5101 .A1 J68
Get This	v.8 1999	✓ On shelf	TK 5101 .A1 J68
Get This	v.9 2000	✓ On shelf	TK 5101 .A1 J68
Get This	v.10-11 2001-2002	✓ On shelf	

 Ask a Librarian ^

Action	Description	Status	Call Number
Get This	v.12-13 2002-2004	✓ On shelf	TK 5101 .A1 J68
Get This	v.14 2005	✓ On shelf	TK 5101 .A1 J68
Show all 14			

Actions Select what to do with this record.

 Email  Text  Citation  Export Citation (EndNote)  Export Citation (RIS)  Copy link

Link copied!

MARC Data

LDR	00000nas a2200361 a 4500
001	990030687990106381
005	20210109123733.0
008	960717c19929999ne qr p 0 0eng d
035	a (MiU)003068799MIU01
010	a sn 96023015
022	a 0926-6801 0 http://worldcat.org/issn/0926-6801
022	a 1875-8940 (online) 0 http://worldcat.org/issn/1875-8940
030	a JHSNEB
035	a (0CoLC)ocm27344595 0 http://www.worldcat.org/oclc/27344595
035	a (SFX)954928598064
040	a NED c NED d CAS d EYM
042	a lcd
050 1 4	a TK5105
245 0 0	a Journal of high speed networks.
246 3 0	a High speed networks



Ask a Librarian



246	1 3	a JHSN
260		a Amsterdam ; a Washington, [D.C.] : b IOS Press, c c1992-
300		a v. : b ill. ; c 24 cm.
310		a Quarterly
362	0	a Vol. 1, no. 1 (1992)-
500		a Title from cover.
515		a Some issues called also: Special issue.
580		a Electronic serial mode of access: World Wide Web.
650	0	a Telecommunication systems v Periodicals. 0 http://id.loc.gov/authorities/subjects/sh2010115961
650	0	a Computer networks v Periodicals. 0 http://id.loc.gov/authorities/subjects/sh2008101495
710	2	a OCLC FirstSearch Electronic Collections Online.
710	2	a IOS Press 0 http://id.loc.gov/authorities/names/no2015091156 0 http://viaf.org/viaf/sourceID/LC no2015091156
856	4 1	3 1992- z Access to the IOS Press online version restricted; authentication may be required: u http://content.iospress.com/journals/journal-of-high-speed-networks m IOS Press
856	4 1	3 1998-2007 z Access to the ECO online version restricted; authentication may be required: u http://firstsearch.oclc.org/journal=0926-6801;ECOIP m ECO
997		a MARS
998		c sersolserial-dld ERLoad20190425merge
998		c sc4
958		a MiU
958		a MiFliC
959		a (notis)ULARC8948
CNV		a wasess_mult
995		a 20



Ask a Librarian



BIB	u 2021-08-17 11:58:59 US/Eastern c 2021-06-21 09:44:58 US/Eastern s false
E56	m IOS Press l ELEC t static u https://na04.alma.exlibrisgroup.com/view/uresolver/01UMICH_INST /openurl?u.ignore_date_coverage=true&portfolio_pid=531072345370006381& Force_direct=true a 2021-06-24 01:54:58 v https://umich.alma.exlibrisgroup.com/view/uresolver/01UMICH_INST /openurl?&rfr_id=info:sid/primo.exlibrisgroup.com& svc_dat=CT0?u.ignore_date_coverage=true?u.ignore_date_coverage=true& rft.mms_id=990030687990106381 z Access to the IOS Press online version restricted; authentication may be required: q true s Available 8 531072345370006381
E56	m ECO l ELEC t static u https://na04.alma.exlibrisgroup.com/view/uresolver/01UMICH_INST /openurl?u.ignore_date_coverage=true&portfolio_pid=531072345350006381& Force_direct=true a 2021-06-24 01:54:58 v https://umich.alma.exlibrisgroup.com/view/uresolver/01UMICH_INST /openurl?&rfr_id=info:sid/primo.exlibrisgroup.com& svc_dat=CT0?u.ignore_date_coverage=true?u.ignore_date_coverage=true& rft.mms_id=990030687990106381 z Access to the ECO online version restricted; authentication may be required: q true s Available 8 531072345350006381
852 0	b OFFS a MiU c ELLS h TK 5101 .A1 J68 8 22828774800006381
866 3 2	a 1-18 8 22828774800006381
866 3 2	a 1992-2012 8 22828774800006381
974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015089899101 e ELLS 7 23828774660006381 z v.18 2011-2012 p 01 r 2012-01-31 05:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015048234895 e ELLS 7 23828774740006381 z v.8 1999 p 01 r 2001-03-01 05:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015048106903 e ELLS 7 23828774750006381 z v.6-7 1997-1998 p 01 r 1999-10-15 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015040328216 e ELLS 7 23828774760006381 z v.5 1996 p 01 r 1998-09-03 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015058750392 e ELLS 7 23828774710006381 z v.12-13 2002-2004 p 01 r 2004-12-01 05:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS



Ask a Librarian



974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015039859445 e ELLS 7 23828774700006381 z v.14 2005 p 01 r 2005-08-30 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015058323836 e ELLS 7 23828774720006381 z v.10-11 2001-2002 p 01 r 2004-09-02 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015035257669 e ELLS 7 23828774790006381 z v.1-2 1992-1993 p 01 r 1996-07-25 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015035258972 e ELLS 7 23828774780006381 z v.3 1994 p 01 r 1996-07-30 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015035257115 e ELLS 7 23828774770006381 z v.4 1995 p 01 r 1996-07-25 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015087155480 e ELLS 7 23828774680006381 z v.16 2007 p 01 r 2007-04-24 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSUE a 39015049124608 e ELLS 7 23828774730006381 z v.9 2000 p 01 r 2001-11-13 05:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015091172562 e ELLS 7 23828774670006381 z v.17 2008-2010 p 01 r 2008-09-23 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS
974	8 22828774800006381 f 1 c ELLS m ISSBD a 39015065053871 e ELLS 7 23828774690006381 z v.15 2006 p 01 r 2006-04-07 06:59:00 US/Eastern h TK 5101 .A1 J68 d OFFS b OFFS

[Give feedback about this page](#) 



Ask a Librarian



[Home](#)

[Accessibility](#)

[Tips for Using Library Search](#)

[Get research help](#)

[About Library Search](#)

[Make an I.L.L. Request](#)

©2022 Regents of the University of Michigan. For details and exceptions, see the [Copyright Policy](#).



Ask a Librarian



APPENDIX 1007C

TK5105
5
0687
v. 7
no. 3-4
ENGI
DISPLAY

ISSN 0926-6801

Journal of **High Speed Networks**

Volume 7 Numbers 3,4 1998

SPECIAL ISSUE:
Multimedia Networking

The Library - UC Berkeley
Received on: 04-12-99
Journal of high speed
networks

LAST NUMBER OF THIS VOLUME

Editor-in-chief: Prof. D. Sidhu

JOURNAL OF HIGH SPEED NETWORKS

Editor-in-Chief

Prof. Deepinder Sidhu
Maryland Center for
Telecommunications Research
University of Maryland-Baltimore
County
Baltimore, MD 21228-5398
USA
Tel.: +1 410 455 3028
Fax: +1 410 455 3969
E-mail: sidhu@cs.umbc.edu or
jhsn@cs.umbc.edu
www: <http://www.cs.umbc.edu/~jhsn>

Editorial Board

A. Acampora
University of California, USA
S. Banerjee
Stevens Institute of Technology, USA
F.-S. Choa
University of Maryland-BC, USA
A. Chung
DePaul University, USA
R. Cruz
University of California, USA
J.J. Garcia-Luna-Aceves
University of California, USA
R. Guérin
IBM T.J. Watson Research Center, USA

Z. Haas
Cornell University, USA
R. Jain
Ohio State University, USA
C. Johnston
NEC USA-C&C Research
Laboratories, USA
S. Keshav
Cornell University, USA
L. Kleinrock
University of California, USA
A. Krishna
IBM T.J. Watson Research Center
USA
S. Kumar
Northwestern University, USA
L. Landweber
University of Wisconsin, USA
B. Makrucki
IBM, USA
D. Mitra
AT&T Bell Laboratories, USA
B. Mukherjee
University of California, USA
S. Pink
SICS, Sweden
C. Qiao
State University of New York, USA

N. Ransom
Bell South Telecommunications, USA
J.M. Smith
University of Pennsylvania, USA
O. Spaniol
Technical University Aachen
Germany
M. Steenstrup
BBN Systems and Technologies, USA
J. Sterbenz
GTE Laboratories, USA
F. Tobagi
Stanford University, USA
S.K. Tripathi
University of California, USA
A. Varma
University of California, USA
R. Watson
Lawrence Livermore National
Laboratory, USA
I. Weerakoon
Hughes Network Systems, USA
E.W. Zegura
Georgia Tech, USA

Aims and scope

The *Journal of High Speed Networks* will serve as an archive for papers describing original results of lasting significance in both the theory and practice of high speed networks. The Journal publishes contributed papers on original research, survey papers on topics of current interest, technical notes, and short communications to report progress on long-term projects. Submissions to the Journal will be refereed consistent with the review process of leading technical journals. The Journal is published quarterly.

The Journal encourages submissions on all aspects of high speed networking. The topics covered will include, but will not be limited to: the design and analysis of high speed networks; protocols for high speed communications; high speed LANs, MANs, and WANs; broadband ISDN and services; optical networks; control and management of high speed networks; high speed switches, interfaces and controllers; routing, flow control and congestion control in high speed networks; formal approaches to high speed protocol development; network management; multimedia applications; applications of high speed networks; experiences with operational high speed networks.

©1998 IOS Press. All rights reserved

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher, IOS Press, Van Diemenstraat 94, 1013 CN Amsterdam, The Netherlands. No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, instructions or ideas contained in the material herein. Although all advertising material is expected to conform to ethical standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

Special regulations for readers in the USA. This journal has been registered with the Copyright Clearance Center, Inc. Consent is given for copying of articles for personal or internal use, or for the personal use of specific clients. This consent is given on the condition that the copier pays through the Center the per-copy fee stated in the code on the first page of each article for copying beyond that permitted by Sections 107 or 108 of the U.S. Copyright Law. The appropriate fee should be forwarded with a copy of the first page of the article to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA. If no code appears in an article, the author has not given broad consent to copy and permission to copy must be obtained directly from the author. This consent does not extend to other kinds of copying, such as for general distribution, resale, advertising and promotion purposes, or for creating new collective works. Special written permission must be obtained from the publisher for such copying.

Published quarterly

0926-6801/98/\$8.00

Printed in the Netherlands

JOURNAL OF HIGH SPEED NETWORKS

Volume 7, Numbers 3,4, 1998

Abstracted / Indexed in: COMPENDEX PLUS, CompuMath Citation Index, Computer Abstracts, Computer Literature Index, Current Contents, EBSCO's database, Ei page One, Engineering Index, INSPEC database, MasterFile, Materials Science Citation Index, SciSearch, Research Alert

CONTENTS

<i>D. Saha and S. Tripathi</i> Guest-editorial: Multimedia Networking	187
<i>S. Mukherjee, D. Reininger and B. Sengupta</i> An adaptive connection admission control policy for VBR ⁺ service class	189
<i>O. Schelén and S. Pink</i> Resource sharing in advance reservation agents	213
<i>K. Nahrstedt, H. Chu and S. Narayan</i> QoS-aware resource management for distributed multimedia applications	229
<i>H.M. Smith, M.W. Mutka and D.T. Rover</i> A feedback based rate control algorithm for multicast transmitted video conferencing	259
<i>W. Feng, C.C. Lam and M. Liu</i> Implementing fast bandwidth smoothing algorithms for the delivery of compressed video streams	281
<i>R. Kravets, K. Calvert and K. Schwan</i> Payoff adaptation of communication for distributed interactive applications	301
<i>G.R. Malan, F. Jahanian and S. Subramanian</i> Attribute-based data dissemination for Internet applications	319
Author Index Volume 7	339



VISIT OUR WEBSITE ON [HTTP://WWW.IOSPRESS.NL](http://WWW.IOSPRESS.NL)



Resource sharing in advance reservation agents

Olov Schelén* and Stephen Pink

Computer Science and Electrical Engineering, Luleå University of Technology, SE-971 87 Luleå, Sweden
Tel.: +46 920 91000; Fax: +46 920 72801; E-mail: {olov,steve}@cdt.luth.se

Abstract. We present an architecture where clients can make end-to-end resource reservations through agents. For each domain in the network, there is an agent responsible for immediate and advance admission control. Reservations from different sources to the same destination domain are aggregated as their paths merge toward the destination.

We show that network resources can be shared between immediate and advance reservations without being pre-partitioned. Admission control for immediate reservations use information about resources to be allocated for advance reservations in the near future. An important parameter in the admission control algorithm is the so called *lookahead time*, i.e., the point at which we actually start making resources available for approaching advance reservations by rejecting immediate requests. In our model, preemption of immediate reservations is made in cases where the admission control cannot make resources available through rejection of immediate requests. The risk of preemption can be varied by changing the lookahead time. We explore, with simulations, the effects of providing advance reservations with this model. The results show the cost in terms of resource utilization, rejection probability and preemption probability.

Keywords: QoS, advance reservations, resource sharing, agents, admission control, simulation, virtual leased lines

1. Introduction

Real-time applications such as audio and video conferencing may need resource reservations in the network to perform well. However, many of the current resource reservation proposals allow only immediate reservations, i.e., reservations at the time the session begins, while real-time events often are scheduled and advertised in advance. Clients making immediate reservations for scheduled events will have to judge the appropriate time for making reservation requests to be accepted in due time. When resources are scarce, clients may start making reservations earlier than needed to make sure that they obtain resources in due time. This leads to sub-optimal scheduling of network resources. By offering reservations in advance, clients obtain better support for planned activities and providers obtain information for scheduling resources over time.

Support for advance reservations is also important between different providers. For scalability reasons, it is not desirable that each end-to-end admission request results in call-specific state for each link along the path. Therefore, providers will negotiate bulk resources among each other for carrying aggregated admission-controlled traffic over several domains. In this context, advance reservations are suitable both for providing time-limited contracts and for providing bandwidth that varies at given time intervals, e.g., over hours of the day and days of the week.

We have designed a reservation architecture where immediate and advance admission control is performed by agents [15,17]. The provided service corresponds to *virtual leased lines* between domains. An advantage with an agent-based approach is that reservations can be provided without having to maintain reservation state in routers. Reservations can be made for both senders and receivers. In addition, reservations can be made for remote locations, thereby supporting third party reservations and nomadic computing, i.e., reservations for places where the host will be moving in the future. Our architecture can aggregate admission requests along the path towards a

*Corresponding author.

destination. This results in admission control for the aggregate being performed for individual links. In Section 2, we give the background on our agent architecture.

The focus of this paper is on resource management for individual links. In Section 3, we describe admission control for immediate and advance reservations [16]. The objective is to share resources between immediate and advance reservations. In Section 3.1, we describe some architectural considerations about locating admission control in agents. In Section 3.2, our service model is described. In this model resources for advance reservations will be made available by *rejecting* immediate requests for a short period before resources are to be allocated, or ultimately by *preempting* service for flows that were granted immediate access. Preemption means that a flow of packets using the reservation loses its service quality, i.e., the flow will still be given best-effort service (preemption is defined in Section 3.2). In Section 3.3, we define the notions *bookahead* and *lookahead*. In Section 3.4, we define the admission control criteria. In Section 4 we show, with simulations, the effects of advance reservations in terms of resource utilization, rejection probability and preemption probability.

2. The agent architecture

In [17] and [15], we present an architecture where clients can make resource reservations through *agents*. For each domain in the network there is an agent responsible for *admission control* (Fig. 1). Each agent knows the topology and static link resources in its domain. The agent is an end-system that is configured for passively participating in a link state routing protocol (e.g., OSPF) where each participating router has an identical topological database over the domain. Thus, an agent obtains the topological database with little signaling overhead. Link state advertisements are sent from the nearest router to the agent. Agents retrieve link properties, such as static bandwidths, by querying routers seen in the topology map. For this, we use a network management protocol (e.g., SNMP). Queries are done at startup and when topology changes are detected by the routing protocol. Routers need no extensions to allow agents to build a resource map through the management protocol.

Packets using reserved resources are *marked* by applications or edge routers. Core routers handle marked packets by classifying them into a small number of predefined service types. Similar ideas have been proposed in the *differentiated services* area [2,12]. In the differentiated services model, packets are marked to obtain either lower drop probability or higher strict priority. Here, we abstract from these details and use the term *priority packets* to denote marked packets, and the term *priority traffic* for all priority packets collectively. A key point of differentiated services and our agent architecture is to avoid signaling state and expensive per-packet processing in routers. To allow resources to be managed in separate agents, we rely on router mechanisms for service differentiation on the packet level.

Independently of which packet scheduling model is used for differentiating service in routers, there must be an admission control architecture to make sure that the amount of priority traffic is low enough to receive good quality of service. One issue is whether the admission control architecture should reserve resources over the links that are actually going to be used, or just limit the overall rights to send priority traffic. A scheme that does not consider




Fig. 1. Reservation agents and their domains.

exactly where traffic will go cannot give reliable commitments without very low utilization. Thus, it is likely that such a scheme would need to gamble and accept occasional service loss.

In this work we focus on agent-based admission control where the objective is to administer resources exactly, i.e., admission control in an agent maintains information about reserved resources on each link in its domain. Agents perform immediate or advance admission control without involving the routers. The architecture offers a service comparable to unidirectional *virtual leased lines*. Reserved resources can be used by one or many applications sending data between the domains connected by a virtual leased line.

There are other frameworks, i.e., IETF Integrated services [7,20] and RSVP [21], that can set up both unicast and multicast reservations end-to-end. However, the provided service is not the same as for virtual leased lines, since reservations are made for individual data flows (i.e., application data streams). In that model, per-flow reservations are maintained in the routers. This results in scaling problems due to reservation state and packet classification overhead in routers. This model inherently relies on per-flow semantics and therefore it is hard to perform aggregation over different administrative domains. Thus, RSVP and Integrated Services may be too heavy-weight for the backbones.

In differentiated service, on the other hand, there is low overhead in routers. Our goal is to design an architecture that performs scalable resource management. Currently, our service model provides only unicast virtual leased links. When resources are managed by agents according to this service model, there are opportunities to aggregate reservations as they merge toward a destination domain. In this section, we give an overview of the admission control architecture and focus on the aggregation issues.

2.1. Direction of admission requests

An admission request can be directed to any agent, e.g., an agent that manages a user's account. Each agent sets up reservations between any two points in the network by invoking other agents. Thus there is *no difference* between reservations for *sending* and *receiving* data, or reservations made by a *third-party*.

Each reservation request contains a bandwidth to be reserved, a source and a destination address. The source address determines the point where the reservation starts (i.e., the point where traffic using the reservation will enter) and the destination address determines the point where it ends. An agent receiving a request first considers whether the starting point is in its domain. If the starting point is *not* in its domain, it finds an agent closer to the starting point and repeats the request with that agent. Technically speaking, it is possible to directly identify the agent that is responsible for the source domain, but in practice clients may prefer using a particular agent. Also, agents may prefer having business relations with only a few adjacent agents. This would require that requests are made in several steps. The result is that accounting information is aggregated, i.e., the mechanism helps keep down the number of separate charges. In Fig. 2, it is shown how a request is issued by node Dx for reserving resources from Ax to Dx.

If an agent finds that the starting point given in the request is in its domain, admission control is performed on the links from the starting point towards the destination, defined by the routing protocol. This is possible since all routers and the agent have the same topology map of the domain. Thus, an agent can perform admission control on the links from the starting point to an edge router in its domain.

An agent cannot perform admission control beyond its domain. Therefore, if the destination is outside of the agent's domain it must request a reservation with the neighboring agent, giving the edge point where traffic will cross the borders as the source address of the reservation (Fig. 2). Without distinguishing between reservations to different destinations, it would be impossible for agents to grant reservations to other domains. This is because there is no way to make sure that there are sufficient reservations along any path in other domains further downstream (unless we accept very low utilization for priority traffic).

2.2. Aggregation

If reservations were always done 'on demand' there would be signaling between all neighboring agents along the path for each end-to-end admission request. Our architecture supports several ways of aggregating reservation




Fig. 2. A receiver Dx requesting a reservation to receive data from Ax.

state. Aggregation takes place in the agents because routers are not part of the admission control process. The motivation for aggregation may be less when reservation state is stored in agents with plenty of memory and disk. However, a per-flow reservation model would have scaling problems both in terms of state and processing cost in agents. Therefore, it is important to show that agents can provide exact resource reservations with good scaling properties.

In [17] we describe how aggregation is done. The key ideas are the following: first, reservations from different sources to the same destination are aggregated as their paths merge toward the destination. Second, an agent in charge of resources at the final destination can generalize reservations for specific end points so that they are valid for any end point in the destination domain, thereby allowing more aggregation.

2.2.1. Merging reservations

To grant requests spanning many routing domains, there must be a commitment from all agents along the path involved in the request. When an agent is making a reservation with a neighboring agent, the source address included in the request determines the point (router interface) where the traffic will cross the borders, i.e., enter the domain of the neighboring agent. The neighboring agent does not care where traffic came from originally. Thus, priority packets can use reserved resources as long as the destination address matches the reservation. This allows agents to aggregate traffic with different sources into one single reservation made with a neighboring agent.

An end-to-end resource reservation can be seen as a number of consecutive *funnels* (Fig. 3). Priority packets poured into any of these consecutive funnels use reserved resources along the path to the destination. Agents may aggregate several requests with different source addresses into the same funnel, as long as all the requests specify the same destination. Agent C (in Fig. 3) sets up a funnel starting at its edge by reserving bulk resources (40 units) with neighboring agent D to a distant destination and then aggregate requests for resources between domains C and D into the established funnel. This bulk reservation implies that downstream agents only keep state for the aggregate. The agent that decides to aggregate into a bulk reservation keeps information about its individual commitments constituting the aggregate. Aggregation by merging reservations towards each destination can be done with full control over the resources. The funnels form a so called *sink-tree* with the root towards the destination. At each merging point, the reserved bandwidth on the outgoing link will be equal to, or larger than, the sum of the reserved bandwidth on the incoming links.

2.2.2. Prefix aggregation

Requests that specify different destinations cannot easily be aggregated into a funnel in the general case. If we allow funnels to split into different directions further downstream, upstream agents must keep information about which destinations are involved and how resources are divided between them. This means that the aggregate no longer can be seen as one unit. However, if a set of destinations can be identified by one common identifier, and




Fig. 3. Funnels and aggregate reservations for one destination.




Fig. 4. Funnels and generalized destinations.

there are sufficient resources available in each branch to service the whole aggregate, then split outs can be scalable and exact.

With this argument, we allow aggregation to be increased by *generalizing* a reservation as valid for all destinations within the same destination domain. This means that a funnel provides resources to all endpoints within the destination domain, i.e., resources in the destination domain must be sufficient to ensure that priority packets get expected service independently of how they are routed within that domain. It is only the agent responsible for the destination domain that can judge if resources are sufficient for allowing generalization. Consider the example in Fig. 4. Agent D can judge if the resources in D are sufficient to let incoming funnels be valid for the whole domain. Thus, it must be agent D that decides when a reservation request for a particular destination node within D can be replaced with a reservation for any node in domain D.

In IP networks an *address prefix* is often used to represent a domain (individual nodes within the domain are distinguished by the remaining suffix). An address prefix is represented by an IP address and an associated mask telling how many bits are part of that prefix. An agent generalizes a request by replying with an *address prefix* to be used as a destination address for the reservation instead of the original address included in the request (Fig. 4). The aggregation that results from generalizing a reservation is also called *prefix aggregation*.

Clients or agents that have been granted a generalized reservation can use it for any destination matching the prefix of the reservation. Agents providing generalized reservations must handle the committed bandwidth of priority traffic to any destination in its domain. A conservative agent would therefore generalize reservations up to an aggregated value equal to the narrowest link. In practice we believe that this can be useful for well provisioned local area networks.

In the Internet, agents must be able to handle occasional topology changes (e.g., node and link failures), as well as pathological changes (e.g., route-flaps) [13]. Agents listen to the routing protocol and detect routing changes on average as fast as any router in the domain. When changes occur, agents can schedule resources along new paths.

3. Admission control for immediate and advance reservations

To obtain a modular design, advance reservations should, as far as possible, be managed separately from immediate reservations. One way of separating the admission control for immediate and advance reservations would be to partition the resources so that one part is available for immediate reservations and another part for advance reservations. However, this would result in two logically separate networks with the problem of partitioning resources according to current demand. Therefore, we have chosen to share resources between immediate and advance admission control. To allocate resources safely we have extended the immediate admission control to consider the near future concerning the resources that will be allocated for advance reservations. There are three major consequences of these results. First, the resources can be managed efficiently without suffering from badly chosen resource partition sizes. Second, the immediate admission control scheme can be kept fast and simple by not having to consider advance reservations in the far future. Third, we have the architectural freedom of locating the admission control for immediate and advance reservations at separate places in the network without suffering undue communication overhead.

3.1. Parameter-based or measurement-based?

Admission control algorithms can be classified as either parameter-based or measurement-based. The metrics used in *parameter-based* admission control are *traffic specifications* and static link resources. By aggregating the given metrics, such a scheme can offer absolute or statistical performance bounds [6,7], often enforced by the packet scheduler (e.g., WFQ or EDF). Parameter-based admission control can be used for immediate reservations and, by adding timing information, for advance reservations.

The metrics used in *measurement-based* admission control involve also *measurements* of used resources [10,20]. Such a scheme can detect reserved resources that are not utilized and grant those resources to other requests. The benefit of measurement-based admission control is better resource utilization. On the other hand, there is no guarantee that reserved resources are available when needed. The scheme may also be sensitive to synchronization effects. Suppose there are many events scheduled at a given time and many clients start reserving immediate resources just before that time. As there is little traffic, more reservations will be allowed. When all events start, the network will be over-allocated with poor service as a result.

In our agent scheme, we use only parameter-based admission control for a number of reasons. First, measurement-based admission control can be efficiently implemented only in routers or switches (where traffic measurement can be done directly). When admission control is in the agent, the signaling overhead for traffic measurements may be prohibitive. Second, traffic measurements provide little information about the future and are therefore not suitable for advance reservations [4]. Third, we want to provide a service model where reservations give a trusted level of service for people that are prepared to pay for it. The problems of lower utilization for reserved traffic are mitigated in a packet-switched network as unused resources always can be used by best-effort traffic. Thus the overall network utilization does not suffer very much from parameter-based admission control.

3.2. Service model

To support advance reservations, information about time must be added to the reservation requests. Therefore, advance reservation requests include a *duration* interval: $I = [t_s, t_e]$, where t_s is the starting time of the requested service and t_e is the finishing time of that service. This information could be automatically collected from a site where scheduled events are advertised, or generated by the client making a reservation. This is indifferent to our architecture.

Immediate reservations are open ended and do not specify a duration. A consequence of allowing open ended immediate reservations together with advance reservations is that there may be cases where resources are over-allocated. This is resolved by *preempting* immediate reservations to free resources for advance reservations. Preemption is used to ensure that there are always resources to honor advance reservations. Flows that lose reserved resources through preemption will be given best-effort service. We have chosen *not* to include a service for immediate requests that are *guaranteed* never to be preempted. Instead, we aim at keeping the preemption probability low, so that the user need not worry about preemption when asking for admission. A fundamental reason for this decision is that indefinite non-preemptible service would block resources indefinitely and no advance reservations of those resources could, therefore, be made. Furthermore, resources that are reserved in advance for some time in the far future cannot be given to indefinite sessions in the meantime unless we have the option of preempting those sessions when resources are needed. Thus, we have the following service model:

- Durations must be specified for advance reservations. A duration includes the start time and the finish time for the requested service.
- Immediate reservations do not specify durations. Such requests are serviced immediately (if possible) and can be preempted later by flows that reserved in advance.
- Preemption means that a flow loses its service quality; however, it will still be serviced at a best-effort level.
- In order to manage resources safely there may be limits imposed on minimum and maximum bookahead times for advance reservations.
- The admission control for advance reservations may support finding the earliest point for which a requested service can be granted.

There are other proposals for advance reservations that do allow indefinite (open ended) resource reservations. To provide indefinite sessions without having preemption, resources could be set aside for this service exclusively. In the Tenet suite [5], resources are partitioned to support advance reservations and non-preemptible immediate sessions simultaneously. With that solution there will be some fragmentation due to it being hard to maintain the partitioning such that the demand for the two different services are met at each instant. In our architecture we can support indefinite sessions by partitioning resources. However, in this paper, we will investigate whether or not we can avoid some problems with resource fragmentation by sharing resources between immediate and advance reservations and instead allow for statistically rare cases of preemption.

There are also proposals relying on statistical guarantees of immediate and advance reservations where both immediate and advance reservations are preemptible. In [18] a model is presented where the call with the shortest bookahead time is preempted in case of congestion. The model relies on the theory of effective bandwidth to calculate the chance of a call being active. In our model we deal with aggregated bulk reservations rather than per-call reservations. For our model we have chosen to use deterministic guarantees in the admission control for advance reservations. Our belief is that advance reservations, once accepted, should not be preempted. Preemption of advance reservations is acceptable only when there are failures in the basic infrastructure. This consideration is made in a context where there is a lot of unplanned network traffic that use immediate reservations and an even larger amount of bursty traffic that uses no reservation at all. In this context resources will be well utilized.

3.3. Bookahead and lookahead

When immediate reservations and advance reservations share resources and the duration of each immediate reservation is unknown, there is a risk of over-allocation that can be resolved with preemption. The probability of

preemption can be controlled by either monitoring immediate requests when admitting advance reservations, or vice versa, by monitoring advance reservations when admitting immediate requests.

In practice, however, the risk of preemption for immediate requests can not be controlled very well by advance admission control. When advance reservations are requested far in advance, the information obtained from monitoring present traffic or present immediate reservations will not be useful [4]. Most of the active flows would have finished at the crucial time anyway. It is only when advance reservations are requested with very short notice that it would be valuable to know about present traffic. Since advance reservations are provided to meet the demand for scheduled sessions primarily, we believe that most reservations will be requested on a notice much longer than the average session length. We denote the time from making an advance reservation until the resources are to be available as the *bookahead time*. The bookahead time is set by the clients asking for admission and varies from request to request. In our admission control we may impose a minimum bookahead time. Then, advance admission control need not consider the immediate reservations state at all.

The risk of preemption can be effectively controlled by immediate admission control. By monitoring advance reservations for the near future when admitting immediate requests, we can control the risk of preemption. If there are insufficient resources available in the near future, the request could be rejected instead of serviced. It is up to the provider to decide whether rejection or preemption is preferable in these cases. Rejecting immediate requests because the resources are reserved in the near future means that the overall utilization will go down to make room for the advance reservations. We denote the pre-allocation time, i.e., the time for starting to set aside resources for advance reservations as the *lookahead time* into the advance reservation state. Setting the lookahead time is an internal issue to the admission control scheme. As shown in simulations (Section 4), the lookahead time can be tuned to obtain reasonable utilization and preemption probability.

3.4. Admission control criteria

Figure 5 is a snap-shot of admitted flows in a time/bandwidth diagram. Flows a, b and c are immediate reservations. For immediate reservations we add up the aggregate bandwidth. When a new admission request α arrives, admission is granted if the new reservation would not cause the aggregate bandwidth to exceed the link bandwidth. For advance reservations, the admission conditions consider other advance reservations that overlap with the new flow (d, e, g, h). The conditions are checked at all points where new flows begin (t_s , t_x and t_y). For immediate requests, the admission conditions consider the current aggregate of immediate reservations and the advance reservations that are active within the lookahead time (d, e, f). The following notation is used in describing our admission criteria:




Fig. 5. Snap-shot of reservations.

R	Aggregate bandwidth reserved by current immediate reservations,
$A(t)$	Aggregate bandwidth reserved in advance at time t ,
$look$	The lookahead time used at immediate admission,
μ	The link bandwidth.

Immediate reservations. When a client requests an immediate reservation α with bandwidth r^α , the admission control algorithm performs the following checks:

- Determine if the bandwidth usage, after adding the new load r^α , will exceed the available link bandwidth μ during the lookahead time

$$\mu > \max_{t \in look} (r^\alpha + R + A(t)). \quad (1)$$

Advance reservations. When a client requests an advance reservation α with bandwidth and duration (r^α, I^α) , the admission control algorithm performs the following checks:

- Check whether the starting time is acceptable as a bookahead time. The minimum bookahead time is configurable to any value larger than the lookahead time.
- Check whether the finishing time is acceptable. It should be larger than the starting time. Admission control may also set a limit on how far in the future reservations can be granted.
- Determine whether the aggregate bandwidth of advance reservations exceeds the link bandwidth at some point during the requested duration:

$$\mu > \max_{t \in I^\alpha} (r^\alpha + A(t)). \quad (2)$$

3.5. Reservation state

As described in Section 2.2, our model has mechanisms to minimize the amount of reservation state. Nevertheless, it is important to have data structures that manage reservations efficiently and allow quick admission control. For immediate admission control, we need to keep information about the current reserved aggregate bandwidth. We also need access to the advance reservation state within the lookahead time. For advance admission control, we need to keep the aggregated bandwidth reserved in advance over time. For this we have two alternative data structures. Our first data structure keeps exact starting times and finishing times for each admission request. The amount of state is proportional to the number of reservations maintained. The other data structure uses time slots where each reservation is extended to occupy a number of time slots. For each slot, the aggregate reserved bandwidth can be found. This data structure has fixed size and depends only on the maximum bookahead time allowed and the slot length.

Also, we must be able to find a victim for preemption among the immediate reservations. For this we need information about individual immediate reservations. However, potential victims can be chosen when they are admitted. It is enough to record a small number of potential victims corresponding to the amount of resources to preempt.

4. Simulations

The objective of our simulations is to explore the effects of allowing reservations in advance. The performance measures are *total utilization*, *rejection probability*, and *preemption probability*. We vary the *fraction* of the total link resources used for *advance reservations* and the *lookahead time* into advance reservation state at admission control for immediate requests. Admission decisions are parameter-based, i.e., based on traffic specifications and static link resources. For simplicity we use the requested bandwidth as the only traffic parameter. This study is for admission control over a single link only.

4.1. Simulation parameters

- The duration of immediate requests, i.e., the call hold time, is exponentially distributed with a mean of 300 time units.
- The duration of advance requests, is also exponentially distributed with a mean of 300 time units.
- The bookahead time for advance requests, i.e., how far in advance the advance requests are made, is exponentially distributed with a mean of 20 000 time units, where the tail is cut at 60 000 units (i.e., there is a maximum bookahead time).
- The total link capacity is 45 000 resource units.
- The resources asked for in each request are uniformly distributed in [1 ... 3000] resource units.
- The inter-arrival time of admission requests is exponentially distributed with a mean of 11.76 time units. This number was selected to obtain a reasonable rejection probability, where around 6% of the requests are rejected. The average offered load is 85% of total capacity.
- When preemption is performed, the immediate request with the largest requested capacity is picked as a victim (which would encourage making requests in advance when large capacity is needed).

Note that the same parameters for duration and resource requirements are used for both immediate and advance reservations. This is to isolate the impact of reserving in advance versus reserving immediately. Should we have different parameters for the two classes, other effects would dominate the outcome in terms of utilization, rejection probability and preemption probability (the impact of having other parameters is discussed in Section 4.6). The bookahead time is chosen to be sufficiently large so it has little impact on the currently active immediate reservations.

The simulations are run many times for each parameter set, with different seeds for the random variables. The 90% confidential intervals are computed and presented as error bars in the plots (although the confidence is often so tight that the error bars appear as dots).

4.2. Advance reservation requests

We vary the fraction of the admission requests that are made in advance between zero and one. The upper bound on resources available for advance reservations equals the total link capacity, resulting in advance reservations being seldom rejected when a small fraction of the requests are made in advance (Fig. 6).

Figure 7 shows that the average duration of the granted requests goes down as there is a higher fraction of advance reservations. When resources are scarce, advance reservations will result in some fragmentation in time, and there is a greater chance that requests for short durations are accepted. This effect does not occur for immediate requests since admission is made without considering durations (in fact, durations are unknown).

A similar effect can be observed concerning the average bandwidth given to the granted requests. If the offered load is increased, the average bandwidth for granted requests will go down as the rejection probability goes up. For immediate reservations clients see a considerable decrease in the chance of succeeding with large reservations as the offered load is increased. However, for advance reservations the rejection probability depends on the bookahead




Fig. 6. Rejection probability for advance requests.




Fig. 7. Average duration for accepted advance requests.

time. Therefore, clients can increase their chance of succeeding with large requests by reserving further in advance. This effect is clearly seen in our simulations. However, as the bookahead time is chosen by clients it is not easy to draw any conclusions on *how* early one should make an admission request. The appropriate bookahead time depends on when other clients choose to book ahead. The only effect that can be claimed for sure is that the chance of having a request granted for a given duration increases strictly with the bookahead time.

4.3. Total utilization

We define the *total utilization* as being the ratio of the total reserved bandwidth versus the total bandwidth of the link. This metric considers how large part of the total resources that are reserved. The metric does not include increased link utilization due to best-effort traffic, nor drops in link utilization due to reserved resources not being utilized.

Figure 8 shows that the total utilization goes down as the fraction of advance reservations increases. There are two reasons for this. First, we start allocating resources ahead of the starting point of the durations for advance reservations. The interval during which resources are allocated ahead of time is denoted the *lookahead* time. Technically speaking the lookahead time is used by admission control for *immediate* requests to judge whether there are any advance reservations needing the resources within that lookahead time. If so, the immediate request will be rejected. Thus, having a longer lookahead time results in more immediate requests being rejected and thereby a lower utilization. Having a shorter lookahead time means that the network will more often be over-allocated, resulting in more immediate requests being preempted (as we will show in Section 4.4).

The second reason for lower utilization is that fragmentation in time is introduced when advance reservations are allowed. There is little chance that other advance requests ask for the small fragments that may be available between already-granted advance reservations. Consequently, a decrease in utilization can be observed as the rate of advance reservations is increased, even though the lookahead time is zero. When the lookahead is zero, admission control for immediate requests does not consider any resources reserved in advance.

This may sound like bad news for providing advance reservations, but in fact similar drops in total utilization would be observed as soon as there are scheduled events introduced in the network, even though users were to make only immediate reservations. Users would still be likely to request resources before events begin to make sure that they have resources allocated in time. The only difference is that the control is with the clients instead of with the network. The degree of over-allocation in time that clients choose would be a consequence of social behavior in an environment with scarce resources. In some situations we can expect that the utilization drops in an uncontrolled way. When there are scheduled events, advance reservations can improve the service and the utilization by ensuring that resources are allocated at an appropriate time.

Measurement-based immediate admission control could alleviate the problem of clients over-allocating in time by measuring traffic on the link. However, we cannot get any useful measurements for scheduled events until the traffic starts. Thus, we must choose between staying with the traffic specification or gradually losing the resource allocation as the measurement procedure finds that there is no traffic. When there are several events scheduled at the same time, there is a risk that allocated resources are insufficient to give acceptable service.

4.4. Rejection and preemption probability

Figure 9 shows that the rejection probability for immediate reservations increases as the fraction of advance reservations is increased. However, the change in rejection probability for immediate requests is primarily an effect of picking the victims among a smaller set of immediate requests as the rate of advance reservations increases. If we instead monitor the total rejection probability (Fig. 10) over both immediate and advance requests, we get a much lower risk of rejection (note that the scales are different). Perhaps this is a better reflection of overall user satisfaction.

We can see that the rejection probability is about the same when we have only immediate requests as it is when we have only advance requests, i.e., around 6%. The decrease in total utilization (Fig. 8) is explained by the fact that the average duration for accepted requests goes down (due to fragmentation in time) with increasing demand for advance reservations (Fig. 7).




Fig. 8. Total utilization.




Fig. 9. Rejection probability for immediate requests.




Fig. 10. Total rejection probability (for immediate and advance requests).




Fig. 11. Preemption probability for immediate requests.




Fig. 12. Total preemption probability (for immediate and advance requests).

Figure 11 shows that the preemption probability increases as the fraction of advance reservations is increased. Again, part of the explanation is that all preemption victims are picked within the immediate requests and as the fraction of immediate requests goes down we get a higher risk of preemption for this service class. In Fig. 12, we show the preemption probability over the total number of served reservations.

It is clear that increasing the lookahead time makes admission control more conservative, resulting in lower preemption probability at the cost of lower utilization and higher rejection probability. By decreasing the lookahead time, we allow the system to be more over-booked resulting in higher preemption probability. The fact that the system is allowed to be over-booked explains why the total rejection probability (Fig. 10) goes down when we use zero lookahead and there is a suitable fraction of advance reservations.

4.5. Lookahead time

Given a desirable target for preemption probability, the lookahead time can be conservatively computed by considering the bandwidth that must be allocated for each advance reservation and the amount of bandwidth that is returned per time-unit by immediate reservations ending. However, at reasonable load, it is likely that there are some unallocated resources that should be considered as well. In our simulations, between 75 and 79% of the resources are allocated on average, which explains why it is possible to allocate resources for advance reservations with a total preemption probability below 5%, even when using zero lookahead.

In previous examples we kept the lookahead times constant as the fraction of advance reservations was increased. We have also tried to use exponentially distributed lookahead times but we found a strict performance degradation compared with constant lookahead time. If there are problems in meeting the requirements for high utilization and low preemption probability, we may be forced to have a limit imposed on resources available for reservations in advance. As an example, we now fix the advance request rate to 40% of the total number of requests (i.e., the average is bounded and there is no firm bound) and we plot the performance measures as functions of the lookahead time.

Figures 13–15 show the effects. It is clear that the preemption probability falls steeply as the lookahead time is increased. In most systems, like cellular phone systems, preemption is considered bad and should be avoided. In a




Fig. 13. Rejection probability for immediate requests, at 40% advance request ratio.




Fig. 14. Preemption probability for immediate requests, at 40% advance request ratio.




Fig. 15. Utilization for immediate requests, at 40% advance request ratio.

connectionless packet switched network, preemption is not as bad since generally the packets are still served best-effort. We believe, however, that the relative weight of having low preemption is still high. Provided that we have the average duration for resource reservations (assuming exponentially distributed session lengths), the fraction of advance reservations, and the individual weights (set by the provider) for utilization, blocking and preemption, we can find a suitable lookahead time.

With a 40% advance reservation request rate, we can, for example, use a lookahead of 50 time units, resulting in 77.6% of the resources being allocated under a 1.1% immediate request preemption probability (i.e., 0.44% total preemption probability). By using a lookahead of 100 time units, we get a utilization of 76.9% at a preemption probability of 0.35% for immediate requests. This is to be compared with 79% of the resources being allocated when there are no advance reservations. The offered load is in both cases 85% of the total capacity and the total rejection probability is around 7%.

If there were reservation agents without supporting lookahead in the immediate admission control algorithm, the results would be comparable with having zero lookahead (the network must support preemption and non-blocking resource allocation that can be used by the advance reservation agent). In that case, we have 5% preemption probability for immediate requests when 40% of the admission requests are made in advance.

4.6. Evaluation

Our simulations show the cost of advance reservations when they use the same holding-times and have the same resource demands as immediate reservations. We do this to isolate the impact of reserving in advance versus reserving immediately. By using different parameters for advance and immediate reservations, utilization increases when there are more advance reservations. Given a certain offered load, it is easy to show that there is a benefit from making large reservations in advance. It is generally easier to obtain a large utilization if we start packing the largest requests into the time-bandwidth diagram shown in Fig. 5. In reality, we expect that users needing large resources will tend to make reservations in advance to avoid being blocked.

5. Conclusion

In this paper we claim that resources can be more accurately allocated over time by supporting advance reservations for scheduled events. The price of supporting advance admission control is a small but controlled decrease in network utilization. This is compared with users making immediate reservations for their scheduled events. When resources are scarce, users tend to ask for admission ahead of time to make sure that they have resources granted in due time for their own scheduled event. Users make their decision with little information about other scheduled events and with little concern about the overall utilization of the network. In some situations, the effective utilization and the probability of having reservation requests granted can drop without control.

Our simulations show the effects of providing advance reservations. We assume that we allocate resources at some point for each reservation and the utilization is measured in terms of how much we can allocate. We do not consider drops in link utilization due to reserved resources not being used.

In our model, resources are shared between advance and immediate admission control, i.e., there is no partitioning of resources. The most important parameter in the admission control algorithm is the so-called *lookahead* time, i.e., the point at which we actually start making the resources in the routers available for approaching advance reservations. We show how a suitable lookahead time can be found through simulations.

The effects of advance reservations can be expressed in terms of rejection probability, preemption probability, and total utilization. In the simulations presented, we can allocate 79% of the link resources when there are no advance reservations and around 77% of the resources when 40% of the admission requests are made in advance, and there is a 0.4% preemption probability for immediate requests. The offered load is in both cases 85% of the total capacity and the total rejection probability is around 7%.

Acknowledgements

This work was supported by a grant from the Center for Distance Spanning Technology (CDT), Luleå, Sweden, and Ericsson Radio Systems, Sweden.

References

- [1] A. Banerjee, D. Ferrari, B.A. Mah and M. Moran, The tenet real-time protocol suite: design, implementation, and experiences, Technical Report, TR-94-059, University of California at Berkeley, Berkeley, California, November, 1994.
- [2] D.D. Clark and W. Fang, Explicit allocation of best effort packet delivery service, Technical Report, MIT Lab for Computer Science, Boston, Massachusetts, November, 1997.
- [3] D.D. Clark, S. Shenker and L. Zhang, Supporting real-time applications in an integrated services packet network: architecture and mechanism, in: *Proc. of SIGCOMM*, Baltimore, Maryland, 1992, pp. 14–26; *ACM, Computer Communication Review* 22(4).
- [4] M. Degermark, T. Köhler, S. Pink and O. Schelén, Advance reservations for predictive service, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 3–14.

- [5] D. Ferrari, A. Gupta and G. Ventre, Distributed advance reservation of real-time connections, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 15–26.
- [6] V. Firoiu, J. Kurose and D. Towsley, Efficient admission control for edf schedulers, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1996.
- [7] R. Guerin, C. Partridge and S. Shenker, Specification of guaranteed quality of service, Request for Comments (Proposed Standard) 2212, Internet Engineering Task Force, October, 1997.
- [8] A. Gupta, W. Howe, M. Moran and Q. Nguyen, Resource sharing for multi-party real-time communication, in: *Proc. of IEEE Infocom*, Boston, Massachusetts, 1995.
- [9] S. Herzog, Rsvp extensions for policy control, Internet Draft, Internet Engineering Task Force, November, 1996, Work in progress.
- [10] S. Jamin, P.B. Danzig, S. Shenker and L. Zhang, A measurement-based admission control algorithm for integrated services packet networks, in: *Proc. of SIGCOMM*, Cambridge, Massachusetts, 1995.
- [11] D.J. Mitzel, D. Estrin, S. Shenker and L. Zhang, A study of reservation dynamics in integrated services packet networks, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1996.
- [12] K. Nichols, V. Jacobson and L. Zhang, A two-bit differentiated services architecture for the Internet, Internet draft, Bay Networks, LBNL and UCLA, November, 1997.
- [13] V. Paxson, End-to-end routing behavior in the Internet, in: *Proc. of SIGCOMM*, Palo Alto, California, 1996, ACM, pp. 25–38.
- [14] W. Reinhardt, Advance reservation of network resources for multimedia applications, in: *2nd International Workshop on Advanced Tele-services and High Speed Communication Architectures (IWACA '94)*, Heidelberg, Germany, 1994.
- [15] O. Schelén and S. Pink, An agent-based architecture for advance reservations, in: *IEEE 22nd Annual Conference on Computer Networks (LCN '97)*, Minneapolis, Minnesota, 1997.
- [16] O. Schelén and S. Pink, Sharing resources through advance reservation agents, in: *Proc. of IFIP Fifth International Workshop on Quality of Service (IWQoS '97)*, New York, 1997.
- [17] O. Schelén and S. Pink, Aggregating resource reservations over multiple routin domains, in: *Proc. of IFIP Sixth International Workshop on Quality of Service (IWQoS '98)*, Napa Valley, California, 1998.
- [18] D. Wischnik and A. Greenberg, Admission control for booking ahead shared resources, in: *Proc. of IEEE Infocom*, San Fransisco, California, 1998.
- [19] L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller and H. Wittig, Issues of reserving resources in advance, in: *Proc. of NOSSDAV, Lecture Notes in Computer Science*, Durham, New Hampshire, Springer, 1995, pp. 27–37.
- [20] J. Wroclawski, Specification of the controlled-load network element service, Request for Comments (Proposed Standard) 2211, Internet Engineering Task Force, October, 1997.
- [21] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, RSVP: a new resource ReSerVation protocol, *IEEE Network Magazine* 7(5) (1993), 8–18.

APPENDIX 1007D



JOURNAL

Journal of high speed networks

PEER REVIEWED

Available at Northern Regional Library Facility NRLF (UCB) (TK5105.5 J687) >

Available Online >

TOP

VIEW ONLINE

GET IT

SEARCH INSIDE

DETAILS

SEND TO

View Online

EBSCOhost Academic Search Complete

Available from 06/01/1996.

Most recent 9 month(s) not available.

SHOW LICENSE



EBSCOhost Business Source Complete

Available from 06/01/1996.

Most recent 9 month(s) not available.

SHOW LICENSE



IOS Press Journals

Available from 1995 volume: 1 issue: 1 until 2010 volume: 18 issue: 4.



view full text

IOS Press online journals UCB Access.



Get It

Please sign in to check if there are any request options.  Sign in

 BACK TO LOCATIONS

LOCATION ITEMS

Northern Regional Library Facility

Available , NRLF (UCB) ; TK5105.5 J687



LOCATION ITEMS

Northern Regional Library Facility
Available , NRLF (UCB) ; TK5105.5 J687
Holdings: v.4(1995)-18:3(2011/12)



TOP

VIEW ONLINE

GET IT

SEARCH INSIDE

LINKS

DETAILS

SEND TO

VIRTUAL BROWSE

Item in place (0 requests) v.4 (1995)
Loanable

Item in place (0 requests) v.5 (1996)
Loanable

Item in place (0 requests) v.6 (1997)
Loanable

Item in place (0 requests) v.7 (1998)
Loanable

Additional call number: TK5105.5 .J687
Location: Northern Regional Library Facility NRLF (UCB) TK5105.5 J687
Barcode: B 5 318 143
Copy: 1

Item in place (0 requests) v.8 (1999)
Loanable

Item in place (0 requests) v.9 (2000)
Loanable

Item in place (0 requests) v.10 (2001)
Loanable

Item in place (0 requests) v.11-12 (2002)
Loanable

Item in place (0 requests) v.13 (2004)
Loanable

Item in place (0 requests) v.14 (2005)
Loanable

TOP

VIEW ONLINE

GET IT

SEARCH INSIDE

LINKS

DETAILS

SEND TO

VIRTUAL BROWSE

Display MARC Record 

Details

Title

Other title

Format

Creation Date

Publisher

Frequency

General

Additional forms

Subject

Related titles

Identifier

Source

Journal of high speed networks

High speed networks

JHSN

v. : ill. ; 24 cm

c1992-

Amsterdam ; Washington, D.C. : IOS Press

Vol. 1, no. 1 (1992)-

Quarterly

Title from cover

Latest issue consulted: Vol. 13, no. 1 (2004)

Also issued online

Telecommunication systems 

Computer networks 

Available in other form: Online version: Journal of high speed networks (Online) 

ISSN : 0926-6801

Library Catalog

Send to

 PRINT

 EMAIL

 CITATION

 REFWORKS

 ENDNOTE

 EASYBIB

 EXPORT RIS

 EXPORT BIBTEX

 DIRECT LINK

https://search.library.berkeley.edu/permalink/01UCS_BER/1thf9n/alma991061649799706532

 COPY THE DIRECT LINK TO CLIPBOARD

leader 01824cas a2200481 a 4500
001 9912545834406531
005 20220618131719.0
008 930129c19929999ne qir p 0 0eng d
010 ##\$asn 96023015
012 ##\$i9604
022 ##\$a0926-6801
030 ##\$aJHSNEB
035 ##\$a(CU-S)b32251269-0lucs_sdi
035 ##\$a(OCOLC)27344595
037 ##\$BIOS Press, Van Diemenstraat 94, 1013 CN Amsterdam, Netherlands
040 ##\$ANED \$cNED \$dCAS \$dEYM \$dOCL \$dNLM \$dMYG \$dHMK \$dOCL
042 ##\$alcd
090 ##\$aTK5101.A1 \$bJ68
090 ##\$aTK5101.A1 \$bJ68 \$9LOCAL
245 00\$aJournal of high speed networks
246 30\$aHigh speed networks
246 13\$aJHSN
260 ##\$aAmsterdam ; \$aWashington, [D.C.] : \$bIOS Press, \$cc1992-
300 ##\$av. : \$bill. ; \$c24 cm
310 ##\$aQuarterly
362 0\$aVol. 1, no. 1 (1992)-
500 ##\$aTitle from cover
500 ##\$aLatest issue consulted: Vol. 13, no. 1 (2004)
515 ##\$aSome issues called also: Special issue
515 ##\$aNone published 2003
530 ##\$aAlso issued online
650 #0\$aTelecommunication systems \$vPeriodicals
650 #0\$aComputer networks \$vPeriodicals
776 18\$iOnline version: \$tJournal of high speed networks (Online) \$w(DLC) 2008233956 \$w(OCOLC)41949723
953 ##\$aTier 3 \$9LOCAL
955 ##\$a20160601 \$bstp \$cMT \$9LOCAL
956 41\$iIOS Press online journals UCB Access. \$uhttp://content.iospress.com/journals/journal-of-high-speed-networks \$9LOCAL
966 ##\$gOldMigrData: Quarterly \$9LOCAL
973 0\$aIOS Press online journals UCB access. \$5CU \$9LOCAL
993 ##\$a20201204 migrate this bib; has wb order syep \$9LOCAL
993 ##\$a20210216 Migration order records with 856s. sydr \$9LOCAL
996 ##\$eGLADN84216102 \$9LOCAL
996 ##\$ab140756267 \$bUpdated: 02-28-21 \$cCreated: 03-28-09 \$dCatDate: 02-15-11 \$9LOCAL

PRIVACY | ACCESSIBILITY | NONDISCRIMINATION

Berkeley

24/7 chat

APPENDIX 1032A



CONFERENCE
ON LOCAL
COMPUTER
NETWORKS
PROCEEDINGS

TK 5105
.5
.C66a
22nd
Nov 2-5
1997

PROCEEDINGS

22nd CONFERENCE ON LOCAL COMPUTER NETWORKS

MINNEAPOLIS, MINNESOTA

NOVEMBER 2-5, 1997

LCN '97

sponsored by



IEEE COMPUTER SOCIETY
IEEE COMPUTER SOCIETY TECHNICAL COMMITTEE ON COMPUTER COMMUNICATIONS






Published by the IEEE Computer Society
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

IEEE Computer Society Press Order Number PR08141
IEEE Catalog Number 97TB100179
ISBN 0-8186-8141-1
ISSN 0742-1303

ISBN 0-8186-8141-1



9 780818 681417

Conference on Local Computer Networks
Proceedings

**22nd Annual Conference on
Local Computer Networks**

LCN'97

November 2-5, 1997

Minneapolis, Minnesota

TX 4-577-503



•TX0004577503•

Sponsored by
**IEEE Computer Society Technical Committee on
Computer Communications**



Los Alamitos, California

Washington

•

Brussels

•

Tokyo

Copyright © 1997 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved

TK 5105
.5
.C66a
22nd Nov. 2-5, 1997

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number PR08141

ISBN 0-8186-8141-1

ISBN 0-8186-8142-X (case)

ISBN 0-8186-8143-8 (microfiche)

IEEE Order Plan Catalog Number 97TB100179

ISSN 0742-1303

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-908-981-1393
Fax: + 1-908-981-9667
mis.custserv@computer.org

IEEE Computer Society
13, Avenue de l'Aquilon
B-1200 Brussels
BELGIUM
Tel: + 32-2-770-2198
Fax: + 32-2-770-8505
euro.ofc@computer.org

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Editorial production by Ian Torwick

Cover design by Joseph Daigle/Studio Productions

Printed in the United States of America by Technical Communication Services



IEEE
COMPUTER
SOCIETY



Table of Contents

LCN'97 22nd Annual Conference on Computer Networks

Message from the General Chair	x
Message from the Program Chairs	xi
Conference Committee	xii

Keynote Address:

Electronic Commerce on the Internet	2
<i>Tim Wild, Tandem Computers</i>	

Expert Panel Session 1: Electronic Commerce

Chair: Professor Agnes Hui Chan, Northeastern University

Expert Panel Session2: Various Approaches to IP over ATM

Chairs: Dr. Rose Tsang, Sandia Nat'l Labs and Tim Salo, MSC

Session I: Network Analysis I

Chair: P. Krishna, Digital Equipment Corporation

Counting the Cycles: A Comparative Study of NFS Performance over High Speed Networks	8
<i>J. Hall, R. Sabatino, S. Crosby, I. Leslie, and R. Black</i>	
Performance of Asynchronous, Variable Bit Rate Communications over ATM: A Comparison with FDDI	20
<i>S. Dempsey, E. Fryland, B. Hazzard, and D. Cyganski</i>	
Performance Analysis of ATM Switches with Multistage Multipath Packet Switching Interconnection Networks	28
<i>A. Tentov</i>	
A Numerical Method for Performance Analysis of ATM Multiplexers	36
<i>M. Dowlatshahi and S.K. De</i>	

Session II: Network / Protocol Design

Chair: Burkhard Stiller,

Institut fuer Technische Informatik und Kommunikationsnetze, Switzerland

Improving Communication Network Topologies Using Tabu Search	44
<i>S. Pierre and A. Elgibaoui</i>	
LAN-ATM Internetworking over a Priority-Based Slotted-Ring Network	54
<i>H. Çam</i>	
Design and Implementation of Bridge between CORBA and DCE	64
<i>D.J. Kim, H. Namgoong, and Y.C. Lew</i>	
Distributed Shared Memory Layer for Cooperative Work Applications	72
<i>H. Guyennet, J.C. Lapayre, and M. Tréhel</i>	

Session III: Routing/ Switching

Chair: Peter Martini, University of Bonn, Germany

Performance of Integrated Telephony Service over Standalone LAN and ATM Internetworked LANs	80
<i>V. Mirchandani and D. Everitt</i>	
Fuzzy Optimization Model for Routing in B-ISDN	89
<i>E. Aboelela and C. Douligeris</i>	
A Routing Algorithm for Distributed Communication Networks.....	99
<i>R. Beaubrun and S. Pierre</i>	
Next Hop Resolution Using Classical IP over ATM	106
<i>C. Park, H.S. Choi, J.O. Kim, and J.H. Lee</i>	

Session IV: PCS and Wireless Networks

Chair: Bikram Bakshi, Intel Corp.

Preserving Cell Sequence of Handoff Connections in Wireless ATM Networks	112
<i>H.T. Lin and H.D. Hughes</i>	
Multiple Access for Third-Generation Personal Communication Systems.....	123
<i>B. Bing and R. Subramanian</i>	
An Adaptive Location Management Algorithm for Mobile Computing	133
<i>R.V.J. Chintalapati, V. Kumar, and A. Datta</i>	

Session V: Network Analysis II

Chair: Parmesh Ramanathan, University of Wisconsin-Madison

Efficient Estimation of Cell Loss and Cell Delay of Nonblocking ATM Switches	142
<i>J. Wang, K.B. Letaief, M. Hamdi, and X.R. Cao</i>	
Efficient Fault-tolerant Reliable Broadcast in an Extended LAN	149
<i>D. Venkatesulu and T.A. Gonsalves</i>	
The Optimal Virtual Path Design of ATM Networks.....	159
<i>Z. Qin, T. Lo, and F.F. Wu</i>	

Session VI: Network Protocol Design II

Chair: Ellis Nolley, Digi International

Video Distribution to the Home Using Advanced Reservation	170
<i>J.J. Harms</i>	
Protocol Implementation in a Vertically Structured Operating System	179
<i>R. Black, P. Barham, A. Donnelly, and N. Stratford</i>	
An Optimal Cache Replacement Algorithm for Internet Systems	189
<i>K.H. Yeung and K.W. Ng</i>	
A New Explicit Rate-Based Congestion Control Scheme for ABR Services	195
<i>K.T. Ma, R. Sankar, and K.J. Christensen</i>	

Session VII: High Performance Protocols

Chair: Tim Strayer, BBN

A Native ATM Protocol Architecture Design and Performance Evaluation.....	204
<i>R. Karabeck</i>	

Scaling CSMA/CD to 1Gb/s with Frame Bursting.....	211
<i>M. Molle, M. Kalkunte, and J. Kadambi</i>	
A Fast Optimal Slot Reuse Scheme for CRMA High Speed Networks	220
<i>P.R. Sheu and C.M. Lin</i>	
Session VIII: Wireless Network Security	
<i>Chair: Gary Kessler, Hill Associates</i>	
Performance Monitoring of a Wireless Campus Area Network.....	232
<i>A. Messier, J. Robinson, and K. Pahlavan</i>	
MC ² Platform for Mobile and Wireless Computing	239
<i>H. Shigeno, Y. Fujimaki, H. Omori, K. Okada, and Y. Matsushita</i>	
Security and Efficiency in Authentication Protocols Resistant to Password Guessing Attacks	245
<i>T. Kwon and J. Song</i>	
Security Services and Public Key Infrastructure for ATM Networks.....	253
<i>V. Varadharajan, R. Shankaran, and M. Hitchens</i>	
Session IX: Quality of Service I	
<i>Chair: Cathy Fulton, Schlumberger APC-Research</i>	
The Effectiveness of Cell-Level FEC for Packet Delivery in ATM Networks.....	264
<i>A. Guha and T.S. Chang</i>	
The Design of an Application Programming Interface for QoS-based Multimedia Middleware	274
<i>C. Conrad and B. Stiller</i>	
Selective Discarding with FEC for ATM-based Internetworks	284
<i>S. Chatterjee and W. Xiao</i>	
DBP-M: A Technique for Meeting End-to-End (m, k)—Firm Guarantee Requirements in Point-to-Point Networks.....	294
<i>W. Lindsay and P. Ramanathan</i>	
Session X: Network/Protocol Design III	
<i>Chair: Joe Bumblis, Tandem</i>	
Early Selective Packet Discard for Alternating Resource Access of TCP over ATM-UBR	306
<i>K. Cheon and S.S. Panwar</i>	
Reducing the Effects of Large Propagation Delays on High Speed IEEE 802.3 CSMA/CD Networks Using Collision Truncation.....	317
<i>M.L. Molle and B. Yeger</i>	
An Enhanced Communication Protocol for Flexible Interconnection	327
<i>H. Yamashita, T. Suguri, and S. Kinoshita</i>	
Practical Experiences with a Transport Layer Extension for End-to-End Bandwidth Regulation	337
<i>M. Frank and P. Martini</i>	

Session XI: Multicast Communications

Chair: Torsten Braun, IBM European Networking Center

An Efficient Scheduling Mechanism Using Multiple Thresholds for Multimedia Traffic in ATM Switching Nodes	348
<i>T.I. Jeong, J.H. Jeong, and S.J. Kim</i>	
A Centralized Generic Protocol for Multipoint Connections	357
<i>P.K. McKinley and Y. Huang</i>	
Reliable Multicast Protocols: a Practical View	369
<i>C.K. Miller</i>	
Reliable Multicast Under Error Conditions	379
<i>G.S. Johnstone and G.D. Williams</i>	

Session XII: Multimedia Communications

Chair: Kanti Prasad, University of Massachusetts, Lowell

Rate Control of Real-time MPEG-2 over ATM DBR Service with Bandwidth Re-negotiation	388
<i>K. Fukuda, N. Wakamiya, M. Murata, and H. Miyahara</i>	
Adaptive Streaming of MPEG Video over IP Networks	398
<i>R.S. Ramanujan, J.A. Newhouse, M.N. Kaddoura, A. Ahamad, E.R. Chartier, and K.J. Thurber</i>	
A Network Manager on Ethernets for Distributed Multimedia Systems	410
<i>V. Balachandran, G. Bhat, S. Chakravarty, and A. Zhang</i>	
An Approach for Imprecise Transmission of TIFF Image Files through Congested Real-Time ATM Networks	420
<i>C. Wong and A.M.K. Cheng</i>	

Session XIII: Quality of Service II

Chair: Cathy Fulton, Schlumberger APC-Research

An Ultra Light Weight CAC Mechanism (UL-CAC) for Soft Guarantee in ATM Networks	432
<i>T. Yokotani, T. Ichihashi, and M. Yabe</i>	
Design of a Network Shared Memory Based Workstation Cluster	441
<i>J. Bonney, R. Ramanujan, and K. Thurber</i>	
An Agent-based Architecture for Advance Reservations	451
<i>O. Schelén and S. Pink</i>	
A Dynamic Routing Procedure for Connections with Quality of Service Requirements	460
<i>M. Nour, A. Hafid, and M. Gendreau</i>	

Session XIV: Optical Networks

Chair: Allalaghata Pavan, Honeywell Technology Center

Reliability Analysis of Partitioned Optical Passive Stars Networks	470
<i>F. Cao</i>	
On the Number of Fiber Connections and Star Couplers in Multi-Star Single-Hop Networks	476
<i>P.J. Wan</i>	

Performance of a Symmetric Robust WDM Network when the Channel Access Pattern at Nodes is Known	482
<i>S.A. Abd-Elmalak, C. Vaishnav, and A.P. Jayasumana</i>	
Session XV: ATM Networks	
<i>Chair: Ken Christensen, University of South Florida</i>	
Design and Evaluation of a Control Plane Architecture for the ATM Access Network	494
<i>E.K. Fragouloupoulos, K.I. Mavrommatis, G.T. Kolyvas, and I.S. Venieris</i>	
Algorithm-Based Error Detection in ATM Cell Schedulers	504
<i>Y.H. Choi</i>	
Reduction of Self-Similarity by Application-Level Traffic Shaping	511
<i>K.J. Christensen and V. Ballingam</i>	
An Adaptive and Selective Cell Drop Policy with Dynamic Data Partitioning for Best Effort Video over ATM	519
<i>A. Mehaoua, R. Boutaba, and G. Pujolle</i>	
Session XVI: Enterprise Networking	
<i>Chair: Bhumi Khasnabish, GTE Labs</i>	
An Experimental Study of Logical Office Service over International ATM Networks	530
<i>H. Fujii, S. Tanimoto, T. Hariu, H. Miyashiro, and L. Galasso</i>	
Interconnecting Fieldbuses through ATM	538
<i>O. Kunert and M. Zitterbart</i>	
A Comparative Study of Simple Mail Transfer Protocol (SMTP) Post Office Protocol (POP) and X.400 Electronic Mail Protocols	545
<i>P. Tzerefos, C. Smythe, I. Stergiou, and S. Cvetkovic</i>	
Dial-in Virtual Private Networks Using Layer 3 Tunneling	555
<i>G.S. Malkin</i>	
Author Index	562

An Agent-based Architecture for Advance Reservations

Olov Schelén and Stephen Pink

Computer Science and Electrical Engineering

Luleå University of Technology

S – 971 87 Luleå, Sweden

{olov, steve}@cdt.luth.se

Tel: +46 920 91000 Fax: +46 920 72821

Keywords: Advance Reservations, Reservation Agents, Quality of Service

Abstract

We propose an architecture where clients can make advance reservations through agents. For each routing domain in the network there will be an agent responsible for admission control on behalf of the routers in the domain. Requests involving several routing domains are forwarded for admission control with agents along the path for the requested service. Agents maintain hard reservation state using a reliable protocol for agent intercommunication.

Agents start allocating resources for advance reservations in the routers by setting up forwarding state shortly before resources are needed for packet forwarding. Resources are made available for advance reservations by means of rejecting further immediate requests and ultimately by preempting some immediate reservations. We have shown that the risk of preemption can be kept very low.

Thus, agents can set up packet classifiers and schedulers in their routers, allowing routers to get on with their main task, packet forwarding.

1 Introduction

Real-time applications such as audio and video conferencing may need resource reservations in the network to perform well. However, many current resource reservation proposals allow only immediate reservations, i.e., reservations at the time the session begins, while real-time events often are scheduled and advertised in advance. Support for advance reservations allows both the network and its users to more

effectively schedule resources when the time of the future event needing reservations is known in advance.

Current protocols designed for resource reservations such as RSVP [29] and the ATM signaling protocol Q.2931 [31], and its Internet analog ST-2 [25] are designed for immediate reservations. In [10], we explain how to extend the soft-state approach in RSVP for support of advance reservations. However, using these protocols for advance reservations would be difficult since senders and receivers would have to be in place and capable of making reservations and large amounts of state would have to be stored in the network for potentially long periods of time.

In this paper we present an architecture where advance reservations are provided by reservation agents. For each routing domain in the network, there will be a reservation agent responsible for admission control on behalf of the routers in the domain (figure 1). Users can make reservation requests by contacting any reservation agent. Reservation requests are forwarded to the agent responsible for the routing domain concerned by the request.

Requests involving several routing domains are forwarded for admission control to the agents along the path. Agents can avoid forwarding each request along the path by making long-term bulk reservations with other agents so that admission requests over several routing domains can be granted locally. Agents establish reservation state through reliable handshaking.

When the time to make resources available approaches, reservation agents reserve resources in routers by setting up forwarding state along the path. A router's primary task is packet forwarding, i.e., packet classification and scheduling. Agents are responsible for setting up appropriate forwarding state in the

⁰This work was supported by a grant from the Center for Distance Spanning Technology (CDT), Luleå, Sweden




Figure 1: Routing domains (AS1 - AS3) and reservation agents (A1 - A3)

routers when it is needed.

An advantage with an agent-based approach is that advance reservations can be provided with little support from routers and without having to maintain reservation state in routers through periods when reservations have no effect on packet forwarding. Furthermore, the agent-driven approach does not require reserving endpoints to be present until the session for which resources have been reserved is starting. Reservations can also be made for remote locations, thereby supporting nomadic computing, i.e., reservations for places where the host will be moving in the future. Since reservations are maintained by agents, reservation requests can be made by a third party, who may make reservations for several endpoints. Thus, neither senders nor receivers need to be in place when resource reservations are made.

2 Overview of the architecture

The Internet has relied on a more-or-less stateless architecture for providing best-effort service. Soft-state solutions such as RSVP have been proposed to provide immediate reservations in an otherwise stateless network. However, advance reservations do not fit into this model because of the amount of overhead that can be generated over long periods of time.

The basic principle of soft-state is that state is refreshed periodically by endpoints either by sending explicit state setup messages or implicitly by sending data packets. This approach works for state that is deployed only during periods when it is actually used for packet forwarding. Many policies such as immediate resource reservations have this property. The signaling overhead for maintaining state is marginal,

especially if policies don't change very frequently. The refresh frequency can be low as route changes will be detected locally and trigger immediate forwarding of messages for setting up state over new paths. For *immediate* reservations this scheme is efficient and flexible.

However, since we expect the book-ahead time for advance reservations to be longer than the actual duration of the session, the relative signaling overhead for maintaining advance reservation state may become unacceptable. We propose that *agents*, deployed in end systems, maintain state about reservations that will be needed for *future* packet forwarding and avoid deploying state in routers until there are packets in the network whose forwarding depends on the reservation.

A reservation agent is a host that is responsible for reserving resources on a part of the Internet. Typically there is one reservation agent responsible for each autonomous system (AS) in the Internet (an AS is an administrative domain containing a number of routers running a routing protocol internally). Such reservation agents advertise themselves periodically on a well known multicast address, giving their identity and which routing domain (AS) they are responsible for. By listening to this well known multicast address, each top level agent will know the existence of all other top level agents. Inside an AS there may be backup agents as well as a hierarchy of agents, e.g., different agents for different areas of the AS. Most agents are internal, i.e., not visible outside the AS, but if there is an agent failure, one of the backup agents is advertised instead.

Each reservation agent must know the topology of its routing domain to perform admission control on behalf of that domain, and also to detect topology changes that require re-computation of forwarding state setup in routers, or renegotiation with adjacent reservation agents. This topology information can be obtained through a link state routing protocol that provides each node, including the agent, with a topology database. It is in the nature of hierarchical routing that an agent cannot have knowledge about the internals of remote autonomous systems and therefore can only provide admission control for its own routing domain. Thus, we need at least one agent for each routing domain.

For agent-based resource reservations, static link performance measures such as link capacity must be known. In the present version of OSPF [20], there are link metrics for different type of services to allow for type-of-service (tos) routing, but there is no information on link capacity. In QOSPF [30], both static link resources and dynamic available resources are advertised. In our scheme, only static link performance is

necessary. This information *could* be included in link state advertisements, which would require an addition to the OSPF specification. Instead we suggest using SNMP for finding out specific properties of links. This solution does not require any changes to the link-state routing protocol, and allows agents to monitor any information made available through SNMP.

To avoid undue signaling overhead and scalability problems, agents do not measure traffic dynamics. Admission control in agents is parameter-based, i.e., considering static link resources and traffic descriptions given in admission requests. Agents may setup resources along shortest paths or perform QoS routing based on the parameters in the admission requests. QoS routes are installed in the forwarding state of the routers.

The architectural goal is to relieve the routers from admission control and other QoS related computing as far as possible. Routers must have an interface so that agents can set up packet forwarding directly. This can be accomplished through SNMP and a special management information base.

For reservation agents to find out which autonomous systems are on the path between source and destination, agents need access to routing information that resides in the border routers of their autonomous system. This can be accomplished either by getting routing information from a border router through SNMP, or by having border routers inject external routing information into their AS. The former may be suitable for stub ASes where no external routing information is injected, while the latter may be suitable for ASes with more than one border router. External routing information together with agent advertisements can be used for forwarding admission requests to other agents.

2.1 Mixing advance and immediate admission control

Agents also support immediate reservations. Network resources are available for both advance and immediate reservations. Advance requests specify a duration of the requested service while immediate requests do not specify durations. In case of conflicts due to unknown durations, immediate requests may be preempted.

In a related paper [24], we present simulation results on sharing network resources between advance and immediate admission control when the admission control algorithms are loosely coupled. Advance admission control is performed without considering current immediate reservations. Immediate admission control can be performed without considering advance

reservations, but the accuracy will be increased by considering resources to be allocated for advance reservations in the near future. Resources are made available for advance reservations by blocking immediate resource requests, or ultimately by means of preempting service for flows that made immediate reservations. In [24], we present results on resource utilization, rejection probability and preemption probability when resources are shared according to this model.

The results show that admission control agents can handle advance and immediate parameter-based admission control, or alternatively coexist with a separate mechanism for immediate admission control, e.g., the IETF int-serv model where immediate admission control (guaranteed service or controlled load service) is performed in routers.

3 Advance reservations over several domains

Agents can interoperate to support end-to-end advance reservation requests in the global network. A host requesting a reservation may consult any reservation agent, e.g., the local one or the favorite agent, just as one may use different travel agents for booking airline tickets. Some agents may already have records for managing accounting etc.

3.1 Unicast

A typical unicast advance request includes a duration, a traffic specification, a source, and a destination. There is no notion of sender or receiver associated with the node making the request. The source and destination descriptions may be the addresses of a sending host and a receiving host respectively. However, in many cases we don't know the addresses of endpoints at the time of making the reservation, but we may know the local area network where addresses temporarily acquired during a session. This may be the case for mobile or stationary computers that acquire addresses through a dynamic host configuration protocol (DHCP), or for computers that will be setup for a specific event, e.g., for an IETF conference. An address mask associated with each endpoint address will solve this problem. Address masks also provide a mechanism for making aggregate reservations, i.e., bulk reservations between different domains of the network.

3.2 Multicast

A multicast advance reservation consists of a duration, a multicast address and a list of senders and re-

ceivers respectively. Each sender and receiver is specified by an address and an address mask. A multicast reservation tree is identified by its multicast destination address and can be extended by making further reservations for the same multicast address. Overlapping requests, i.e., requests having links in common, is merged by the reservation agent forming one shared multicast tree.

3.2.1 Rendezvous agents

To support advance reservations with little overhead we have a *rendezvous* agent associated with each multicast address. The rendezvous agent has group membership information, i.e., the complete list of senders and receivers. When making advance multicast reservations, the rendezvous agent provides an updated list of participants so that multicast reservations can be made, starting from the agents responsible for receiver domains. Building the tree starting at receivers is more efficient as sessions often involve fewer senders than receivers, and adding a receiver may result in reservations over a few additional links for joining an already established multicast reservation tree.

Some current multicast routing protocols, e.g., sparse mode Protocol Independent Multicast [8] and Core Based Trees [1], use rendezvous points. They are designed for setting up multicast routing trees immediately, i.e., when receivers are in place and communicating. In PIM a rendezvous point only keeps track of senders of the group and each receiver (local router) is responsible for joining with the rendezvous point to obtain a branch of the multicast tree. PIM also allows the multicast tree to be rebuilt on demand from individual receivers (local routers) to use shortest paths. To support new receivers joining, the rendezvous point is always a member of the multicast tree (thus receiving data from all senders).

Our notion of rendezvous agents is different. First to support reservations for not yet present receivers, a rendezvous agent must record all members of the group, i.e., both senders and receivers. Second, the rendezvous agent is only used as a database for finding group membership. There is no obligation that sender data pass the rendezvous agent. Group membership information is used to build multicast trees that will finally be used so that resource reservations can be made on the right branches. It does not make sense to change the distribution tree on demand when there are advance reservations, unless there are sudden changes in topology making it impossible to meet committed service along the chosen paths. The decision about whether to build a shortest-path tree or a core-based

tree can be taken by the agents. We leave this issue open for now.

3.2.2 Hierarchical agents

To support scaling, we use a hierarchical approach, i.e., agents register border points of their routing domains rather than specific endpoints. The internal multicast tree and the local members of a routing domain are concealed by the reservation agent. Adding another member to a session where members are sharing resources can be done internally. Inside an AS there is a hierarchy of routing domains and there will be a corresponding hierarchy of reservation agents that aggregate reservations and register with the next highest level. At the top level there are several ASes and corresponding peer agents. These agents should only register involved border routers of their autonomous systems with the rendezvous agent. The task of global rendezvous agents is to maintain the list of member ASes of their multicast groups. Hierarchical approaches have been suggested before, in [15] a hierarchical approach for organizing PIM rendezvous points is suggested.

The scalability problem could be approached by getting rid of rendezvous agents completely. In dense-mode schemes for building immediate multicast trees a flood and prune approach is often used [9] [26]. However, for advance reservations there is no data available until the session starts and therefore the relative cost of flooding control information between agents and keeping prune state during the booking time would be high.

A related issue is how to find the rendezvous agent for a given multicast address. With a hierarchical approach, all agents know about all other agents at the same level (by multicasting), so the problem boils down to partitioning the address space between agents. We propose an approach based on dynamic prefix partitioning of multicast addresses. The range of multicast addresses can be advertised as an address prefix associated with each agent. Agents then need a distributed algorithm to partition a multicast address space between themselves. The solution requires that a range of consecutive multicast addresses is allocated for the agent architecture.

4 Packet classification

Deciding which packets that are to use against a particular reservation is known as *packet classification*. This is performed by the routers for each packet and therefore poses a scaling problem. There are basically two ways of doing packet classification.

First, a *filter specification* could be given by clients having reserved resources. Packet classification is performed on the headers of the packets. In RSVP filtering is performed primarily on addresses and destination and source ports. Second, a reservation key (a certificate) could be included in each packet that is to use against a reservation. Packet classification is done on this single field of the header. Depending on how many different service classes we support, the field could be one or many bits long.

The problem with filter specifications is efficiency. In the worst case of filtering on sender addresses, the routers need a filter for each sender to a multicast group. Multicast routing does not scale well even without packet classification, i.e., for best effort service. Current proposals for multicast routing scales at best as $O(G)$, where G is the number of multicast groups that are served by a router [1]. Thus, the size of the multicast routing table is proportional to the number of multicast addresses that are currently served by the router.

Our goal is to add multicast resource reservations while preserving the $O(G)$ scaling factor also for packet classification. This rules out making classification on specific source addresses. Instead we suggest that the rendezvous agent issue a key for the multicast address to other trusted agents, which in their turn may install the key in routers. The key may be added to packets by clients having reservations for sending to that group, or for better security, by the most local router under authority of a reservation agent. In the latter case only the local router must do classification and policing on packet headers to decide which packets that should be annotated with reservation keys. In IPv6, keys could be included in the hop-by-hop options header. Since agents can compute one aggregate reservation for each hop of the multicast tree, routers can make classification using only one key for each reservation.

This scheme provides scalable support for aggregate multicast reservations, i.e., where a set of senders share reserved resources. The amount of state needed for packet classification in routers is in $O(G)$, which is comparable with best-effort multicast routing. The RSVP model supports a stronger but unscalable model where it is possible to reserve resources in routers for individual senders (known as fixed filters) or for a set of senders (known as shared explicit filter). One motive for having this support is for receivers with insufficient network resources to be able to allocate resources only for the senders that it is interested in. However, in a typical multicast session the set of interesting senders changes dynamically, e.g., with the

current speaker of a meeting, so a receiver with insufficient network resources would have to make a lot of re-reservations during the session. Another motive is for misbehaving senders not to be able to steal bandwidth from reservations, i.e., some policing is done by packet classification.

In our model, applications must make sure to use multicast addresses so that independently of the number of sources there is a limit on the bandwidth that is used by the aggregate multicast flow. Audio and video servers, usually have a known bandwidth limit, e.g., they provide flows with different bandwidths that can be selected by receivers depending on the capacity of their access link. It is important to know which maximum bandwidth that is used by a flow before connecting. In multiparty audio sessions bandwidth is limited by senders taking turns speaking. Video applications can be designed in a similar way, e.g., so that only the most recent speakers and the upcoming speakers on the speakers list send motion video. We believe that it is a bad idea to introduce nonscalable network solutions in order to support applications that waste bandwidth.

The advantage of providing shared reservations only is that packet classification can be done using only one classification entry for each flow (multicast address). This together with the possibility for making bulk reservations between subdomains of the network by using address prefixes to specify leaves of the reservation tree should provide a scalable mechanism for quality of service.

4.1 Classification granularity

Agents may choose internally how many different service levels to support. The low end extreme is to have only one admission controlled service level shared between all flows. Then, packets only need to carry one bit to tell whether it is to use the primary service. Routers need only one queue for the primary service and another queue for best effort service. Several service levels can be supported by having a fixed number of queues in routers. The high end extreme would be to allow a queue for each flow, e.g., flow-wise weighted fair queuing. The scaling for this may be acceptable for multicast, as multicast routing already keeps per-flow state. For unicast, however, we suggest using a small number of admission controlled service levels to avoid per-flow state in the routers. If QoS routing is to be performed we may add service classes to distinguish between outgoing interfaces.

5 Security and policing

In our architecture, packet classification is performed by including a reservation key in each packet that is using against a reservation. Consequently, malicious senders may take currently used keys and steal bandwidth from resource reservations. However, the agents know which senders are entitled to use reserved resources and occasional policing can be done to catch violating senders. Receivers of a multicast group may even hint to their provider when policing of a group should be done. The multicast reservation scheme should not include policing for unwanted senders at every hop as default. Instead, agents should have a way to setup routers close to the edges with specific policing tasks. When there is a malicious sender identified, an agent should be able to tell the most local trusted router of the malicious sender to exclude the flow. Thus, per source filtering and policing is introduced at specific points close to the edges of the network when necessary. Per-source policing does not scale well and therefore local routers cannot be expected to have sufficient capacity for policing continuously for every local sender. In the general case, policing should be soft in the sense that only a few flows are checked at each time. The policing policy is driven by the reservation agents. An important issue is to develop policing and protection mechanisms that are fast and scalable enough to find and exclude violating flows accurately. The requirements depend on how common deliberate stealing of reserved resources will be.

6 Setting up forwarding state in routers

Resource set-up messages are distributed by agents to routers just before the resources are allocated. We use a soft state setup protocol where agents contact and refresh each involved router in their domain periodically. Such a protocol can be optimized by using multicast in cases where there are many routers that are to be updated with the same message. Since soft state is refreshed periodically, we do not need a reliable multicast protocol. Instead the refresh frequency can start high as new state is deployed and then gradually decrease as there is no change in state.

If the overhead for soft state turns out to be too high, a hard state protocol could be used instead. Agents could open a TCP connection (unicast) to each of their routers and setup state reliably. With this approach the routers must be able to clear up reservation state in case a reservation agent should go down and

forget about the state.

To allow reservation agents to query and set router state, the routers must expose an interface. For an initial implementation the most straight forward approach is to use SNMP, i.e., having agents setting router state by sending unicast SNMP messages to each of them. The main problem is to make necessary router state available through the management information base (MIB). The authentication problem will be solved by SNMP mechanisms.

7 Routing performed by agents

Agents can choose to follow best effort routes and establish forwarding state for reserved resources along the current path. However, sometimes routes change (or flap) although the old route is still available and has sufficient resources for the committed service.

We suggest that agents be responsible for selecting routes and that these routes can be included in the forwarding state set in routers by agents. This also makes quality of service (QoS) routing possible. For advance reservations there is time for calculating QoS routes. When allocating resources in routers, the agents provide the address/interface of the next hop as well. Thus, routing information is included in the packet classification state, so that flows using resources reserved through agents can take a different route than the best effort packets to the same destination.

It is the job of the agent to re-route at link or router failure. If there is a link failure, routers next to the failed link will detect the problem first, e.g., through the OSPF HELLO messages. This will trigger a link state advertisement to other nodes in the AS, including the reservation agent. As the reservation agent finds that a link is broken, it will process the reservations on that link by recomputing routes, checking if there is capacity left for advance reservations through alternative routes. This will work if the capacity is available and flows haven't yet started. If capacity on alternative routes is insufficient, the agent can just wait and retry, or choose to inform users who have made advance reservations about possible problems. If there are reservations for currently active flows on the broken link and there is capacity in an alternative path, the agent will immediately allocate resources on the alternative path. The old reservation state in the routers can be torn down by the agent or simply left to time out by not being refreshed.

Thus, when there are failures, routes will change only if there is an alternative path that can provide resources, otherwise the traffic is served best effort.

When routes change the agents take multi hop decisions over the involved reservations, so the scheme does not suffer from having an arbitrary outcome in terms of who gets what on which hops. When links between different administrative domains fail, several reservation agents need to react.

As rerouting doesn't happen instantaneously, the flow state may direct packets for a short time to an interfaces that is unavailable. When this happens the packets should be sent through best effort routing instead. That includes removing any reservation keys in the packet so that best-effort service is used all the way to the destination. This works fine for unicast, but for multicast there may not be any best effort multicast routing state available and then the packets will be dropped.

If all routers along the path support the agent architecture, all routes are loop-free. However, routing loops may occur when there are routers on the path that only support best-effort routing. To avoid this, the flow state routing must check so that packets arrive on the correct incoming interface. Should a packet arrive on an incoming interface different from the one negotiated during QoS route setup, it may be a packet that coming into the QoS path again after a loop.

8 Comparison with the RSVP policy architecture

The RSVP policy architecture of [16] [17] is an extension to the RSVP architecture [29]. Reservation requests goes through both resource based admission control at routers and through policy based admission control managed by local policy modules or by policy servers. Policies are bilateral agreements between neighbor providers and policy admission control is enforced at border nodes only. Reservation requests, carried via the normal RSVP soft state mechanisms, contain policy data objects that are passed to local policy modules at border nodes. Policy data objects may be added, removed, or rewritten by policy modules along the path. Routers have the choice to handle policies in their local policy module or to outsource policy decisions to a policy server. There is a protocol for communication between routers and policy servers using hard or soft state.

An important difference between our approach and the RSVP policy approach is the role of routers. In our architecture, agents handshake to make reservations and to compute the forwarding state for routers, while the routers primary task is to do packet forwarding, i.e., best effort service and service disciplines installed by the agents. Admission control and QoS routing

decisions are done by agents. Our goal has been to relieve the routers from computation concerned with QoS contracts and put it in agents instead. Thus admission control and service commitments are all managed by agents. In RSVP, routers play a central role in admission control and routing. Routing decisions are not part of RSVP but instead made by the best effort routing mechanisms in the routers. Consequently the RSVP mechanism can easily reestablish reservations along changing routes. However, when routes change (or flap) the outcome in terms of who gets reservations depend on the order reservations arrive on the new path. An annoying problem with RSVP is that users do not know if a reservation is granted end-to-end. Non-RSVP routers may be on the path and may disturb end-to-end quality. We believe that RSVP is suitable for providing loose service commitments but it cannot support guaranteed service very well. Our architecture may adapt more slowly to changing topology but it provides a better vehicle for guaranteed quality of service and certainly for advance reservations (which it was designed for from the very beginning). In RSVP senders must be present to make reservations, which might not always be possible far in advance.

Another problem with RSVP is poor support for aggregate reservations. Endpoints make reservations on a per-flow basis and although reservation messages are merged up the multicast tree, we still have reservations per flow. In addition, packet filtering is designed so there may be classification state for each sender. This is clearly not scalable. In our agent architecture aggregation of reservations is included from the beginning. Classification is based on keys (certificates) and the leaves of the multicast reservation tree may be routing domains (address prefixes). There is work in RSVP on introducing address prefix aggregation (aka CIDR prefixes), but alone this is not enough to handle sparsely located senders to the same multicast group. Packet filtering/classification may still have to be done on a per sender basis.

9 Conclusions

Advance and immediate reservations can be supported by agents responsible for admission control on parts of the Internet. Our architecture minimizes disturbing the primary task of the routers: packet forwarding. Soft state in the routers is deployed only during periods when it is needed for packet forwarding.

The agent-driven approach described in this paper does not require the communicating endpoints to be

present until the session for which resources have been reserved is starting. Reservations can also be made for remote locations, thereby supporting nomadic computing, i.e., reservations for places where the host will be moving in the future. Since reservations are maintained by agents, reservation requests can be made for a group of participants by a third party.

Reservation agents have access to routing databases for finding out the topology of the routing domain that it performs admission control for. Reservation agents use SNMP to find out the static link resources available in the network and for setting up forwarding state in the routers. Agents perform parameter-based admission control, using traffic descriptions provided by clients making reservation requests. Thus, agents work without signaling any traffic dynamics.

The agent architecture is suitable for parameter-based QoS routing of advance reservations. There is intrinsically more time available for admission control and QoS routing when requests are made in advance.

References

- [1] A. Ballardie, P.F. Francis, J. Crowcroft. *Core based trees*. In Proceedings of the ACM SIGCOMM, San Francisco, 1993.
- [2] A. Ballardie: *Core Based Trees (CBT) Multicast Routing Architecture*, Work in progress, May 1997.
- [3] A. Banerjee, B. Mah: *The Real-Time Channel Administration Protocol*. Proc. 2nd International Workshop on Network and Operating System Support for Digital Audio and Video, November 1991.
- [4] A. Banerjee, D. Ferrari, B. Mah, M. Moran, D. Verma, H. Zhang: *The Tenet Real-Time Protocol Suite: Design, Implementation, and Experiences*. Technical Report TR-94-059, ICSI, Berkeley, CA, November 1994.
- [5] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin: *Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification*. Work in progress, Nov 1996.
- [6] J. Case, K. McCloghrie, M. Rose, S. Waldbusser: *Introduction to version 2 of the Internet-standard Network Management Framework*. RFC 1441, April 1993.
- [7] D. Clark, S. Shenker, L. Zhang: *Supporting Real-Time Applications in an Integrated Packet Services Network: Architecture and Mechanism*. Proc. ACM SIGCOMM'92, 1992.
- [8] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C-G. Liu, L. Wei: *An Architecture for Wide-Area Multicast Routing*. Proc. ACM SIGCOMM'94, 1994.
- [9] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, A. Helmy, L. Wei: *Protocol Independent Multicast Version 2, Dense Mode Specification*, Work in progress, May, 1997.
- [10] M. Degermark, T. Köhler, S. Pink, O. Schelén: *Advance Reservations for Predictive Service*. Proc. 5th NOSSDAV, Durham, New Hampshire, April 1995.
- [11] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, L. Wei: *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*, RFC 2117, June 1997.
- [12] D. Ferrari: *Real-Time Communication in an Internetwork*. Journal of High Speed Networks, 1(1), 1992, pp. 79-103.
- [13] D. Ferrari, A. Gupta, G. Ventre: *Distributed advance reservation of real-time connections*. Proc. 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [14] A. Gupta: *Advance reservations in real-time communication services*. Sun Microsystems Laboratory. URL: <http://www.cs.berkeley.edu/~amit/research/Postscript/adv-issues.ps>
- [15] M. Handley, J. Crowcroft, I. Wakeman: *Hierarchical Protocol Independent Multicast (HPIM)*, Oct 1995. URL <ftp://cs.ucl.ac.uk/darpa/IDMR/hpim.ps>
- [16] S. Herzog: *Building Blocks for Accounting and Access Control in RSVP*. USC/ISI, November 1995.
- [17] S. Herzog: *RSVP Extensions for Policy Control*. Work in progress, draft-ietf-rsvp-policy-ext-02.txt, April 1997.
- [18] S. Jamin, P. Danzig, S. Shenker, L. Zhang: *A Measurement-based Admission Control Algorithm for Integrated Services Packet Networks*. Proc. ACM SIGCOMM'95, 1995.
- [19] D. J. Mitzel, D. Estrin, S. Shenker, L. Zhang: *A Study of Reservation Dynamics in Integrated Services Packet Networks*. Proc. IEEE INFOCOM'96, March 1996.
- [20] J. Moy: *OSPF Version 2*. RFC 1583, March 1994.
- [21] C. Partridge, S. Pink: *An Implementation of the Revised Internet Stream Protocol*. Internetworking Research and Experience, 3(1), March 1992, pp. 27-54.
- [22] W. Reinhardt: *Advance reservation of network resources for multimedia applications*. Proc. ICWA 94, Germany, October 1994.
- [23] Y. Rekhter, T. Li: *A Border Gateway Protocol 4 (BGP-4)*. RFC 1654, July 1994.
- [24] O. Schelén, S. Pink: *Sharing Resources through Advance Reservation Agents*. Proceedings from IWQoS'97, International Workshop on Quality of Service, New York, May 1997. URL: <http://www.cdt.luth.se/~olov/publications.html>
- [25] C. Toplocic: *Experimental Internet Stream Protocol, Version 2 (ST-II)*. RFC 1190, October 1990.
- [26] D. Waitzman, C. Partridge, S. Deering: *Distance Vector Multicast Routing Protocol*. RFC 1075, November 1988.

- [27] L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller, H. Wittig: *Issues of Reserving Resources in Advance*. Proc. 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [28] D. Zappala, B. Braden, D. Estrin, S. Shenker, D. Zappala: *Interdomain Multicast Routing Support for Integrated Services Networks*. November 1996.
- [29] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala: *RSVP: A New Resource ReSerVation Protocol*. IEEE Network Magazine, September 1993, pp. 8-18.
- [30] Zhang, Sanches, Salkewicz, Crawley: *Quality of Service Extensions to OSPF*. work in progress, June 1996.
- [31] "Broadband Integrated Service Digital Network (B-ISDN) Digital Subscriber Signaling System No.2 (DSS2) User Network Interface Layer 3 Specification for Basic Call/Connection Control", ITU-T Recommendation Q.2931, International Telecommunication Union: Geneva, 1994.

APPENDIX 1032B

LIBRARY

CATALOG

PERIODICAL OR NEWSPAPER

Proceedings

Full Record MARC Tags

Meeting name
Conference on Local Computer Networks.

Main title
Proceedings / Conference on Local Computer Networks.

Published/Created
Washington, D.C. : IEEE Computer Society Press, c1986-

Request this Item

 LC Find It

Links >

Links <https://ieeexplore.ieee.org/xpl/conhome/1000419/all-proceedings> 

More Information >

LCCN Permalink <https://lcn.loc.gov/98657356>

Publication history 11th (1986)-

Description v. : ill. ; 28 cm.

Current frequency	Annual
ISSN	2831-7742
LC classification	TK5105.5 .C66a
Continues	Conference on Local Computer Networks. Conference on Local Computer Networks : [proceedings] 0742-1303 (DLC) 83641144 (OCoLC)8309764
Variant title	Title on IEEE Computer Society Proceedings list: Annual IEEE International Conference on Local Computer Networks Additional title on IEEE Computer Society Proceedings list: LCN Proceedings of the IEEE ... Conference on Local Computer Networks 2007-
Portion of title	Local computer networks
Cover title	IEEE ... local computer networks 1986
Spine title	IEEE Conference on Local Computer Networks 1986
Serial key title	Proceedings (Conference on Local Computer Networks. Print)
Abbreviated title	Proceedings (Conf. Local Comput. Netw. Print)
Related names	IEEE Computer Society. Technical Committee on Computer Communications.
LC Subjects	Local area networks (Computer networks)--Congresses.
Other Subjects	Local area networks (Computer networks)
Form/Genre	Conference papers and proceedings. Congresses.
Browse by shelf order	TK5105.5
Notes	Some issues have also a distinctive title. Published: Los Alamitos, Calif., <1996-> Sponsored by: IEEE Computer Society, TC-Computer Communications.
Additional formats	Also issued online. Online version: Conference on Local Computer Networks. Proceedings 2832-1421 (DLC) 2022202008 (OCoLC)232333634

Online version: Conference on Local Computer Networks. Proceedings
(OCoLC)763170076

LCCN	98657356
Invalid LCCN	sn 89011888
Dewey class no.	004.68
Other system no.	(OCoLC)ocm17022071
Incorrect ISSN	0742-1303
Additional Links	https://ieeexplore.ieee.org/xpl/conhome/1000419/all-proceedings To access, click on LCN at http://csdl2.computer.org/persagen/ProceedingByFilter.jsp?filter=L
Type of material	Periodical or Newspaper
Source of description	Latest issue consulted: 30th (2005).

Item Availability



CALL NUMBER	Electronic Resource
Request in	.Electronic Journal
Links available	https://ieeexplore.ieee.org/xpl/conhome/1000419/all-proceedings

CALL NUMBER	TK5105.5 .C66a Set 1
Request in	Jefferson or Adams Building Reading Rooms
Status	Not Charged
Latest receipts	[no] 38 (2013 Oct.) [no] 34 (2009 Oct.) [no] 34 (2009 Oct.) [no] 31 (2006 Nov.) [no] 30 (2005 Oct.)

10/27/22, 8:23 PM

LC Catalog - Item Information (Full Record)

[no] 29 (2004 Oct.)
[no] 28 (2003 Oct.)
[no] 27 (2002 Nov.6/8)
[no] 25 (2000 Nov.)

Older receipts

11-24 (1986-1999), 26 (2001)

LIBRARY

CATALOG



PERIODICAL OR NEWSPAPER

Proceedings

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

000 03171cas a2200625 a 4500
001 11502468
005 20220512065707.0
008 871124c19869999cauar 1 a0eng c
010 __ |a 98657356 |z sn 89011888
022 0_ |a 2831-7742 |y 0742-1303 |2 1
035 __ |a (OCoLC)ocm17022071
040 __ |a MYG |b eng |c MYG |d NST |d COO |d NST |d OCL |d NST |d WAU |d MYG |d DLC |d MYG |d
OCLCQ |d CGU |d MYG |d OCLCQ |d OCLCF |d OCL |d OCLCO |d GILDS |d DLC
042 __ |a pcc |a nsdp
050 00 |a TK5105.5 |b .C66a
082 10 |a 004.68 |2 23/eng/20220413
111 2_ |a Conference on Local Computer Networks.
210 0_ |a Proceedings |b (Conf. Local Comput. Netw. Print)
222 _0 |a Proceedings |b (Conference on Local Computer Networks. Print)
245 10 |a Proceedings / |c Conference on Local Computer Networks.
246 30 |a Local computer networks
246 14 |a IEEE ... local computer networks |f 1986
246 18 |a IEEE Conference on Local Computer Networks |f 1986
246 1_ |i Title on IEEE Computer Society Proceedings list: |a Annual IEEE International Conference on
Local Computer Networks
246 1_ |i Additional title on IEEE Computer Society Proceedings list: |a LCN
246 1_ |a Proceedings of the IEEE ... Conference on Local Computer Networks |f 2007-
260 __ |a Washington, D.C. : |b IEEE Computer Society Press, |c c1986-

300 ___ |a v. : |b ill. ; |c 28 cm.

310 ___ |a Annual

362 0_ |a 11th (1986)-

500 ___ |a Some issues have also a distinctive title.

500 ___ |a Published: Los Alamitos, Calif., <1996->

530 ___ |a Also issued online.

550 ___ |a Sponsored by: IEEE Computer Society, TC-Computer Communications.

588 ___ |a Latest issue consulted: 30th (2005).

590 ___ |a SERBIB/SERLOC merged record

650 _0 |a Local area networks (Computer networks) |v Congresses.

650 _7 |a Local area networks (Computer networks) |2 fast |0 (OCoLC)fst01001200

655 _7 |a Conference papers and proceedings. |2 fast |0 (OCoLC)fst01423772

655 _4 |a Congresses.

710 2_ |a IEEE Computer Society. |b Technical Committee on Computer Communications.

776 08 |i Online version: |a Conference on Local Computer Networks. |t Proceedings |x 2832-1421 |w (DLC) 2022202008 |w (OCoLC)232333634

776 08 |i Online version: |a Conference on Local Computer Networks. |t Proceedings |w (OCoLC)763170076

780 00 |a Conference on Local Computer Networks. |t Conference on Local Computer Networks : [proceedings] |x 0742-1303 |w (DLC) 83641144 |w (OCoLC)8309764

850 ___ |a AzTeS |a CU-
SB |a CaAEU |a CoU |a DLC |a KMK |a MBU |a MCM |a MiDW |a MnU |a NBuU |a NIC |a NNC |a NSyU |a NcD |a PU |a UU |a ViW |a WaU

859 41 |u <https://ieeexplore.ieee.org/xpl/conhome/1000419/all-proceedings>

859 41 |z To access, click on LCN at |u <http://csdl2.computer.org/persagen/ProceedingByFilter.jsp?filter=L>

890 ___ |a Proceedings (Conference on Local Computer Networks) |i 98-657356

906 ___ |a 7 |b cbc |c serials |d u |e ncip |f 19 |g n-oclcserc

920 ___ |a Keep 1

955 ___ |k xf15 2022-04-13 (telework)

984 ___ |a srvf |d 2001-08-28

985 ___ |e eserial 200406

991 ___ |b c-GenColl |h TK5105.5 |i .C66a |w SERIALS

991 ___ |b c-GenColl |h TK5105.5 |i .C66a |p 00057176234 |t Copy 1 |v 22nd * 1997 Nov 2-5 |w CCF

992 ___ |b B/L |w SERLOC

[Request this Item](#)
 [LC Find It](#)

Item Availability



CALL NUMBER Electronic Resource

Request in .Electronic Journal

Links available <https://ieeexplore.ieee.org/xpl/conhome/1000419/all-proceedings>

CALL NUMBER **TK5105.5 .C66a**
Set 1

Request in Jefferson or Adams Building Reading Rooms

Status Not Charged

Latest receipts
[no] 38 (2013 Oct.)
[no] 34 (2009 Oct.)
[no] 34 (2009 Oct.)
[no] 31 (2006 Nov.)
[no] 30 (2005 Oct.)
[no] 29 (2004 Oct.)
[no] 28 (2003 Oct.)
[no] 27 (2002 Nov.6/8)
[no] 25 (2000 Nov.)

Older receipts 11-24 (1986-1999), 26 (2001)

APPENDIX 1032C



The Library's Main Reading Room will close to researchers earlier than usual, at 5:00pm, on Thurs., Nov. 3. All other reading rooms will be open according to their daily schedules.

Try the [Copyright Public Records System \(CPRS\)](#) pilot with enhanced search features and filters.


[Help](#)[Search](#)[History](#)[Titles](#)[Start Over](#)

Public Catalog

Copyright Catalog (1978 to present)

Search Request: Left Anchored Title = Conference on local computer networks

Search Results: Displaying 8 of 22 entries

[Labeled View](#)

Conference on Local Computer Networks : proceedings.

Type of Work: Serial

Application Title: IEEE ... conference on local computer networks

Title: Conference on Local Computer Networks : proceedings.

Serial Publication Year: 1997

Serial Key Title: Conference on Local Computer Networks

Imprint: Silver Spring, MD : IEEE Computer Society Press.

ISSN: 0742-1303

Description: print material.


Frequency: Annually.

Description based on: No. 9.

Copyright Claimant: Institute of Electrical and Electronics Engineers, Inc.

Issues Registered: no. 22, 1997. Created 1997; Pub. 1997-11-02; Reg. 1997-12-02; TX0004577503

Names: [Institute of Electrical and Electronics Engineers, Inc.](#)



Save, Print and Email ([Help Page](#))

Select Download Format	Full Record	▼	Format for Print/Save
Enter your email address:			Email

[Help](#) [Search](#) [History](#) [Titles](#) [Start Over](#)

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

APPENDIX 1032D

Appendix 1032D
Selected Early Citations to Schelén-1997

Source Document

Schelén, O., & Pink, S. (1997, November). An agent-based architecture for advance reservations. In *Proceedings of 22nd Annual Conference on Local Computer Networks* (pp. 451-459). IEEE.

Selected Early Citations from 1999 (in chronological order)

Berson, S., Lindell, R., & Braden, B. (1998). An Architecture for Advance Reservations in the Internet.

Wong, J. S., & Mikler, A. R. (1999, July). Intelligent mobile agents in large distributed autonomous cooperative systems. *Journal of Systems and Software*, 47(2-3), 75-87.

Karsten, M., Berier, N., Wolf, L., & Steinmetz, R. (1999, September). A policy-based service specification for resource reservation in advance. In *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON COMPUTER COMMUNICATION* (No. 2, pp. I-1).

Google Scholar

Articles

Any time

Since 2022

Since 2021

Since 2018

Custom range...

Sort by relevance

Sort by date

Create alert

About 60 results (0.02 sec)

An agent-based architecture for advance reservations

Search within citing articles

Networks with advance reservations: The routing perspective

RA Guérin, A Orda - ... Nineteenth Annual Joint Conference of the ..., 2000 - ieexplore.ieee.org

This paper provides an initial look at how support for advance reservations affects the complexity of the path selection process in networks. Advance reservations are likely to ...

☆ Save

🔖 Cite

Cited by 269

Related articles

All 9 versions

🔗

Performance issues of bandwidth reservations for grid computing

LO Burchard, HU Heiss... - 15th Symposium on ..., 2003 - ieexplore.ieee.org

In general, two types of resource reservations in computer networks can be distinguished: immediate reservations which are made in a just-in-time manner and advance reservations ...

☆ Save

🔖 Cite

Cited by 88

Related articles

All 5 versions

Networks with advance reservations: Applications, architecture, and performance

LO Burchard - Journal of Network and Systems Management, 2005 - Springer

In general, two types of resource reservations in computer networks can be distinguished: immediate reservations which are made in a just-in-time manner and advance reservations ...

☆ Save

🔖 Cite

Cited by 54

Related articles

All 8 versions

Resource sharing in advance reservation agents

O Schelén, S Pink - Journal of High Speed Networks, 1998 - content.iospress.com

We present an architecture where clients can make end-to-end resource reservations through agents. For each domain in the network, there is an agent responsible for immediate ...

☆ Save

🔖 Cite

Cited by 65

Related articles

All 7 versions

Analysis of data structures for admission control of advance reservation requests

LO Burchard - IEEE Transactions on Knowledge and Data ..., 2005 - ieexplore.ieee.org

[PDF] technion.ac.il

[PDF] academia.edu

[PDF] diva-portal.org

[PDF] archive.org

Ex. 1039

CISCO SYSTEMS, INC. / Page 244 of 245

Any time

Since 2022

Since 2021

Since 2018

Custom range...

1998

—

1999

Search

Sort by relevance

Sort by date

✓

Create alert

An agent-based architecture for advance reservations

Search within citing articles

Resource sharing in advance reservation agents

O Schelén, S Pink - Journal of High Speed Networks, 1998 - content.iospress.com

We present an architecture where clients can make end-to-end resource reservations through agents. For each domain in the network, there is an agent responsible for immediate ...

☆ Save 77 Cite Cited by 65 Related articles All 7 versions

Intelligent mobile agents in large distributed autonomous cooperative systems

JSK Wong, AR Mikler - Journal of Systems and Software, 1999 - Elsevier

A large distributed autonomous cooperative system is a system that provides fundamental services for integrating high-level services in a large distributed system with local autonomy ...

☆ Save 77 Cite Cited by 60 Related articles All 7 versions

Aggregating resource reservations over multiple routing domains

O Schelén, S Pink - ... Sixth International Workshop on Quality of ..., 1998 - ieeeexplore.ieee.org

Presents an agent-based architecture for resource reservations. For each domain in the network, there is an agent that is responsible for admission control. The architecture ...

☆ Save 77 Cite Cited by 52 Related articles All 6 versions

[CITATION] An architecture for advance reservations in the internet

S Berson, R Lindell, R Braden - 1998 - Citeseer

☆ Save 77 Cite Cited by 54 Related articles All 2 versions

[CITATION] A policy-based service specification for resource reservation in advance

M Karsten, N Berier, L Wolf, R Steinmetz - PROCEEDINGS OF THE ..., 1999 - Citeseer

☆ Save 77 Cite Cited by 28 Related articles All 10 versions

[CITATION] Steven Berson Robert Lindell Robert Braden USC Information Sciences Institute July 16, 1998

R Lindell - 1998

☆ Save 77 Cite Cited by 28 Related articles All 10 versions

[CITATION] Olav Schelen

[PDF] diva-portal.org

[PDF] academia.edu

[PDF] diva-portal.org

[PDF] researchgate.net

Ex. 1039

CISCO SYSTEMS, INC. / Page 245 of 245