UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DENTAL IMAGING TECHNOLOGIES CORPORATION

Petitioner

v.

3SHAPE A/S

Patent Owner

U.S. Patent No. 11,368,667 to Fisker *et al.*
Issue Date: June 21, 2022
Title: INTRAORAL SCANNING APPARATUS

DECLARATION OF JAMES L. MULLINS, Ph.D.

Exhibit 1010 page 1 of 78
DENTAL IMAGING

TABLE OF CONTENTS

I. INTRODUCTION 3

II. BACKGROUND AND QUALIFICATIONS 3

III. PRELIMINARIES 5

A. Library Records 8

B. Periodicals 10

C. Cataloging/Indexing 10

IV. OPINION REGARDING INDIVIDUAL DOCUMENT 12

A. Zhang 12

1. Authentication 12

2. Public Accessibility 14

3. Conclusion Regarding Zhang 18

V. CONCLUSION 18

I. INTRODUCTION

1. I, James L. Mullins, Ph.D., hereby declare as follows:

2. I have been retained by Foley & Lardner LLP on behalf of Dental Imaging Technologies Corporation to authenticate and establish the date of public accessibility of a document for use in an *inter partes* review proceeding. For this service, I am being paid my usual hourly fee of $275/hour. My compensation in no way depends on the substance of my testimony or the outcome of the proceeding.

3. I have personal knowledge of the facts and opinions set forth in this declaration. I believe them to be true, and if called upon to do so, I would testify competently to them. I have been warned that willful false statements and the like are punishable by fine or imprisonment, or both.

II. BACKGROUND AND QUALIFICATIONS

4. I am a retired academic librarian working as the Founder and Owner of the firm Prior Art Documentation Librarian Services, LLC at 205 Saint Cuthbert, Williamsburg, VA 23188. Attached as Appendix B is a true and correct copy of my Curriculum Vitae ("CV") describing my background and experience. Further information about my firm, Prior Art Documentation Librarian Services, LLC (PADLS), is available at www.priorartdoclib.com.

5. In forming the opinions expressed in this Declaration, I have considered and relied upon my background and experience. In addition, I have

reviewed and relied upon additional materials in preparation of this Declaration, as described below and provided in the Appendixes to this Declaration.

6. I am presently Dean Emeritus of Libraries and Esther Ellis Norton Professor Emeritus, Purdue University, 2018 – present.

7. I was previously employed as follows:

- Dean of Libraries and Professor & Esther Ellis Norton Professor, Purdue University, West Lafayette, IN, 2004-2017.
- Assistant/Associate Director for Administration, Massachusetts Institute of Technology (MIT), Cambridge, MA, 2000-2004.
- University Librarian and Director, Falvey Memorial Library, Villanova University, Villanova, PA, 1996-2000.
- Director of Library Services, Indiana University South Bend, South Bend, IN, 1978-1996.
- Part-time instructor, School of Library and Information Science, Indiana University, Bloomington, IN, 1979-1996.
- Associate Law Librarian, and associated titles, Indiana University School of Law, Bloomington, IN, 1974-1978.
- Catalog Librarian, Assistant Professor, Georgia Southern College (now University), Statesboro, GA, 1973-1974.

8. Over the course of my career as a librarian, instructor of library science, author of scholarly publications, and presenter at national and international conferences, I have had experience with catalog records and online library management systems built around Machine-Readable Cataloging (MARC) standards.

9. In the course of more than forty-four years as an academic librarian and scholar, I have been an active researcher. In my years as a librarian, I have facilitated the research of faculty colleagues either directly or through the provision of and access to the requisite print and/or digital materials and services at the universities I worked. I have kept current on the professional library science literature and served on the editorial board of the most prominent library journal, *College and Research Libraries*. This followed service as the chair of the Research Committee of the Association of College and Research Libraries (ACRL), a division of the American Library Association (ALA). As an academic library administrator, I have had responsibility to ensure that students were educated to identify, locate, assess and integrate information garnered from library resources.

III. PRELIMINARIES

10. I am not a lawyer, and I am not rendering an opinion on the legal question of whether a particular document is, or is not, a "printed publication" under the law.

11. I am, however, rendering my expert opinion on the authenticity of the documents referenced herein and on when and how these documents were disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located the documents by at least June 17, 2008.

12. I am informed by counsel that an item is considered authentic if there is sufficient evidence to support a finding that the item is what it is claimed to be. I am also informed that authenticity can be established based on the contents of the documents themselves, such as the appearance, content, substance, internal patterns, or other distinctive characteristics of the item, taken together with all of the circumstances.

13. I am informed by counsel that a given reference is publicly accessible upon a satisfactory showing that such a document has been disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence, can locate it. I have also been informed by counsel that materials available in a library constitute printed publications if they are cataloged and indexed according to general library practices that make the references available to members of the interested public.

14. In forming the opinions expressed in this declaration, I have reviewed the documents and appendices referenced herein. These materials are records

created in the ordinary course of business by publishers, libraries, indexing services, and others. From my years of experience, I am familiar with the process for creating many of these records, and I know these records are created by people with knowledge of the information in the record. Further, these records are created with the expectation that researchers and other members of the public will use them. All materials cited in this declaration and its appendices are of a type that experts in my field would reasonably rely upon and refer to in forming their opinions.

15. I am told by counsel that the subject matter of this proceeding relates to a scanner for scanning three-dimensional objects.

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

17. I am told by counsel that a person of ordinary skill in the art ("POSA") would have had a bachelor's degree in Computer Engineering, Computer Science, Electrical Engineering, Physics, or an equivalent field, as well as at least one or two years of design experience, or alternatively, at least five years of comparable industry experience. In particular, I am informed that a POSA

would have had experience with and knowledge of three-dimensional ("3D") imaging systems and 3D modeling techniques.

18. It is my opinion that such a person would have been engaged in academic research, learning through study and practice in the field and possibly through formal instruction the bibliographic resources relevant to that person's research. Such a person would have had access to a vast array of print resources in 3D scanner design.

A. Library Records

19. Some background on MARC formatted records, OCLC, *WorldCat*, and OCLC's Connexion is needed to understand the library catalog records discussed in this declaration.

20. Libraries world-wide use the MARC format for catalog records; this machine-readable format was developed by the Library of Congress in the 1960s.

21. MARC formatted records provide a variety of subject access points based on the content of the document being cataloged. All subject headings may be found in the MARC fields 6XX. For example, MARC Field 600 identifies personal names used as subjects and the MARC Field 650 identifies topical terms. A researcher might discover material relevant to his or her topic by a search using the terms employed in the MARC Fields 6XX.

22. The MARC Field 040, subfield a, identifies the library or other entity

that created the original catalog record for a given document and transcribed it into machine readable form. The MARC Field 008 identifies the date when this first catalog record was entered on the file. This date persists in all subsequent uses of the first catalog record, although newly-created records for the same document, separate from the original record will show a new date. It is not unusual to find multiple catalog records for the same document.

23. *WorldCat* is the world's largest public online catalog, maintained by the Online Computer Library Center, Inc., or OCLC, and built with the records created by the thousands of libraries that are members of OCLC. *WorldCat* provides a user-friendly interface for the public to use MARC records; it requires no knowledge of MARC tags and codes. *WorldCat* is easily accessible through the World Wide Web to all who wish to search it; there are no restrictions to be a member of a particular community, etc. The date a given catalog record was created (corresponding to the MARC Field 008) appears in some detailed *WorldCat* records as the Date of Entry but not necessarily all.

24. Whereas *WorldCat* records are widely available, the availability of MARC formatted records varies from library to library.

25. When an OCLC participating institution acquires a document for which it finds no previously created record in OCLC, or when the institution chooses not to use an existing record, it creates a record for the document using

OCLC's Connexion, the bibliographic system used by catalogers to create MARC records. Connexion automatically supplies the date of record creation in the MARC Field 008.

26. Once the MARC record is created by a cataloger at an OCLC participating member institution, it becomes available to other OCLC participating members in Connexion and also in *WorldCat*, where persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, can locate it.

B. Periodicals

27. A library typically creates a catalog record for a periodical publication when the library receives its first issue. When the institution receives subsequent issues/volumes of the periodical, the issues/volumes are checked in (often using a date stamp), added to the institution's holding records, and made available very soon thereafter – normally within a few days of receipt or (at most) within a few weeks of receipt.

28. The initial periodicals record will sometimes not reflect all subsequent changes in publication details (including minor variations in title, frequency, etc.).

C. Cataloging/Indexing

29. A researcher may discover material relevant to their topic in a variety of ways. One common means of discovery is to search for relevant information in

an index of periodicals and other publications. Having found relevant material, the researcher will then normally obtain it online, look for it in libraries, or purchase it from the publisher, a bookstore, a document delivery service, or other provider. Sometimes, the date of a document's public accessibility will involve both indexing and library date information. Date information for indexing entries is, however, often unavailable. This is especially true for online indices.

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

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

32. ***ACM Digital Library*** - Association for Computing Machinery. The ACM was founded in 1947 and is a scientific and educational computing society. The portal has two main sections: the ACM Digital Library and the ACM Guide to Computing Literature. The ACM Digital Library provides access to the journals, newsletters and conference proceedings of the organization. The Guide to

Computing Literature is a bibliography of computing literature beginning in the 1950s with more than a million entries. https://dl.acm.org.

33. All metadata in the ACM database are freely available on the web. This includes abstracts, linked references, citing work, and usage statistics. Full text of articles are available with subscription.

34. ***Wisconsin TechSearch (WTS) -*** WTS offers an array of article delivery and research services to individuals and organizations not affiliated with the University of Wisconsin or who need the specialized skills of WTS staff in locating and retrieving information. It is the policy of WTS to keep confidential all material submitted by clients with respect to searches as well as search results. https://wts.wisc.edu/.

IV. OPINION REGARDING INDIVIDUAL DOCUMENT

A. Zhang

35. I have been asked to opine on an article authored by: L. Zhang *et al.*, titled: "Projection Defocus Analysis for Scene Capture and Image Display;" published in ACM SIGGRAPH 2006, July 2006, on pages 907-915. ("Zhang").

1. Authentication

36. I have evaluated the Zhang reference three ways: (1) by assessing the copy of Zhang (Ex. 1006 to the Petition for *Inter Partes* Review) provided to me by counsel; (2) by assessing scans of a copy of Zhang owned by the Library of

Congress, Appendix A-1; and (3) by accessing and assessing a download of Zhang I made on September 7, 2022.

37. Appendix A-1 is a scan of the print document held by the Library of Congress. It was provided to me at my request from the Wisconsin TechSearch from the print copy of Zhang owned by the Library of Congress on September 1, 2022. Appendix A-1 includes: bound front cover, spine and back cover with a stamp that reads "T 385.A28 Set 1", as shown below; published cover; publication information page; title page; verso of the title page (copyright) with a stamp that reads "Library of Congress AUG 15 2006", as shown below; Table of Contents; and Zhang article on pages 907-915.







38. All identifying characteristics, such as stamps and notations, on Appendix A-1, are consistent with library practice and procedure that I have observed during my career as a professional librarian, specifically with those items held by the Library of Congress. I have no cause for concern about the authenticity

or accuracy of these identifying attributes. In addition, Zhang was found within the custody of a library, the Library of Congress, one of the most likely locations for an authentic publication to be located.

39. Evidence of the ownership of Zhang by the Library of Congress will further be explored and detailed in the section below on public access.

40. Appendix A-2 is a download I made on September 8, 2022. I accessed and downloaded Appendix A-2 from the *ACM Digital Library* database by my emeritus faculty status through Purdue University Libraries at this URL: https://dl-acm-org.ezproxy.lib.purdue.edu/doi/pdf/10.1145/1141911.1141974.
It is available for a fee at this URL:
https://dl.acm.org/doi/abs/10.1145/1179352.1141974.

41. After comparison of Zhang between Appendix A-1 and corresponding pages of Appendix A-2, I found no difference between the two. Therefore, upon finding Zhang in the Library of Congress and also in the *ACM Digital Library* database, I have determined that Ex. 1006, Zhang, is an authentic document.

42. I conclude and affirm that Zhang, Ex. 1006, is an authentic document.

2. Public Accessibility

43. Appendix A-3 is a true and correct copy of the *WorldCat* entry for Zhang. I obtained Appendix A-3 by completing a search on *WorldCat* on September 10, 2022 by searching geographical location, Washington, D.C.

(Library of Congress). When I searched *WorldCat* for holdings of Zhang in Washington, D.C., the Library of Congress was third on the list among the 459 libraries shown as holding Zhang worldwide.

44. Appendix A-3 shows that Zhang is the document associated with this *WorldCat* entry, as verified by title and authors: *Proceedings. ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, SIGGRAPH, Association of Computing Machinery;* publication date: c2005-, which is an open entry indicating that it is an annual and each year would be added as received.

45. Zhang could have been located by searching for the author and title: *Proceedings. ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games;* and/or by searching the subject headings: *Computer graphics; Computer graphics - Congresses; Three-dimensional display systems;* and/or *Three-dimensional display systems - Congresses.*

46. The search discussed above could have been performed anywhere in the world by anyone who accessed *WorldCat* and its predecessor database through an OCLC member library, which have been available from the 1990s through today.

47. Appendix A-4 is a download I made from the Library of Congress OPAC (online catalog) on September 1, 2022. The document cataloged in this record is Zhang as verified by the OPAC record fields including meeting name:

ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games; main title: *Proceedings;* publisher and publication date; Association of Computing Machinery from 2005-.

48. Zhang could have been located in the Library of Congress OPAC by searching for meeting name and title: *Proceedings, ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*; or by searching the subject headings: *Computer graphics – Congresses; Three-dimensional – Congresses; Computer graphics;* and/or *Three dimensional display systems.*

49. The Library of Congress OPAC record documents that the *2006 Proceedings, ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* is held by the Library of Congress is listed in "Request in Jefferson or Adams Reading Rooms" Latest receipts 2009, 2008, 2007, **2006**, 2005, as shown below:

Request in	**Jefferson or Adams Building Reading Rooms**
Latest receipts	**2009** **2008** **2007** **2006** **2005**

50. The date stamped on Appendix A-1 of August 15, 2006, as shown below, indicates receipt of Zhang by the Library of Congress, after which it would have been sent for processing and labeling. From my extensive professional experience and my knowledge of the Library of Congress processing, it is my

opinion that Zhang would have been available at the Library of Congress within a week to ten days, therefore, Zhang would have been available no later than August 25, 2006.



51. Once Zhang was entered into the general collection of the Library of Congress, members of the public could access the book by having it brought to either the Jefferson or Adams Reading Rooms. The collections of the Library of Congress are searchable by subject matter, author, or title such that a skilled researcher could find works in which they were interested. For example, a member of the public could have located a copy of Zhang by searching for the meeting name and title: *Proceedings, ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*; or by searching the subject headings: *Computer graphics – Congresses; Three-dimensional – Congresses; Computer graphics;* and/or *Three dimensional display systems.*

52. Members of the public could read, study, and make notes about a selected work in the Library of Congress Reading Rooms. Further, members of the public were permitted to make photocopies of portions of the work while in the Reading Rooms. Accordingly, a copy of Zhang was accessible to the general

public when it was available at the Library of Congress.

3. Conclusion Regarding Zhang

53. I conclude that Zhang, Ex. 1006, is an authentic document and would have been publicly accessible through the Library of Congress no later than August 25, 2006.

V. CONCLUSION

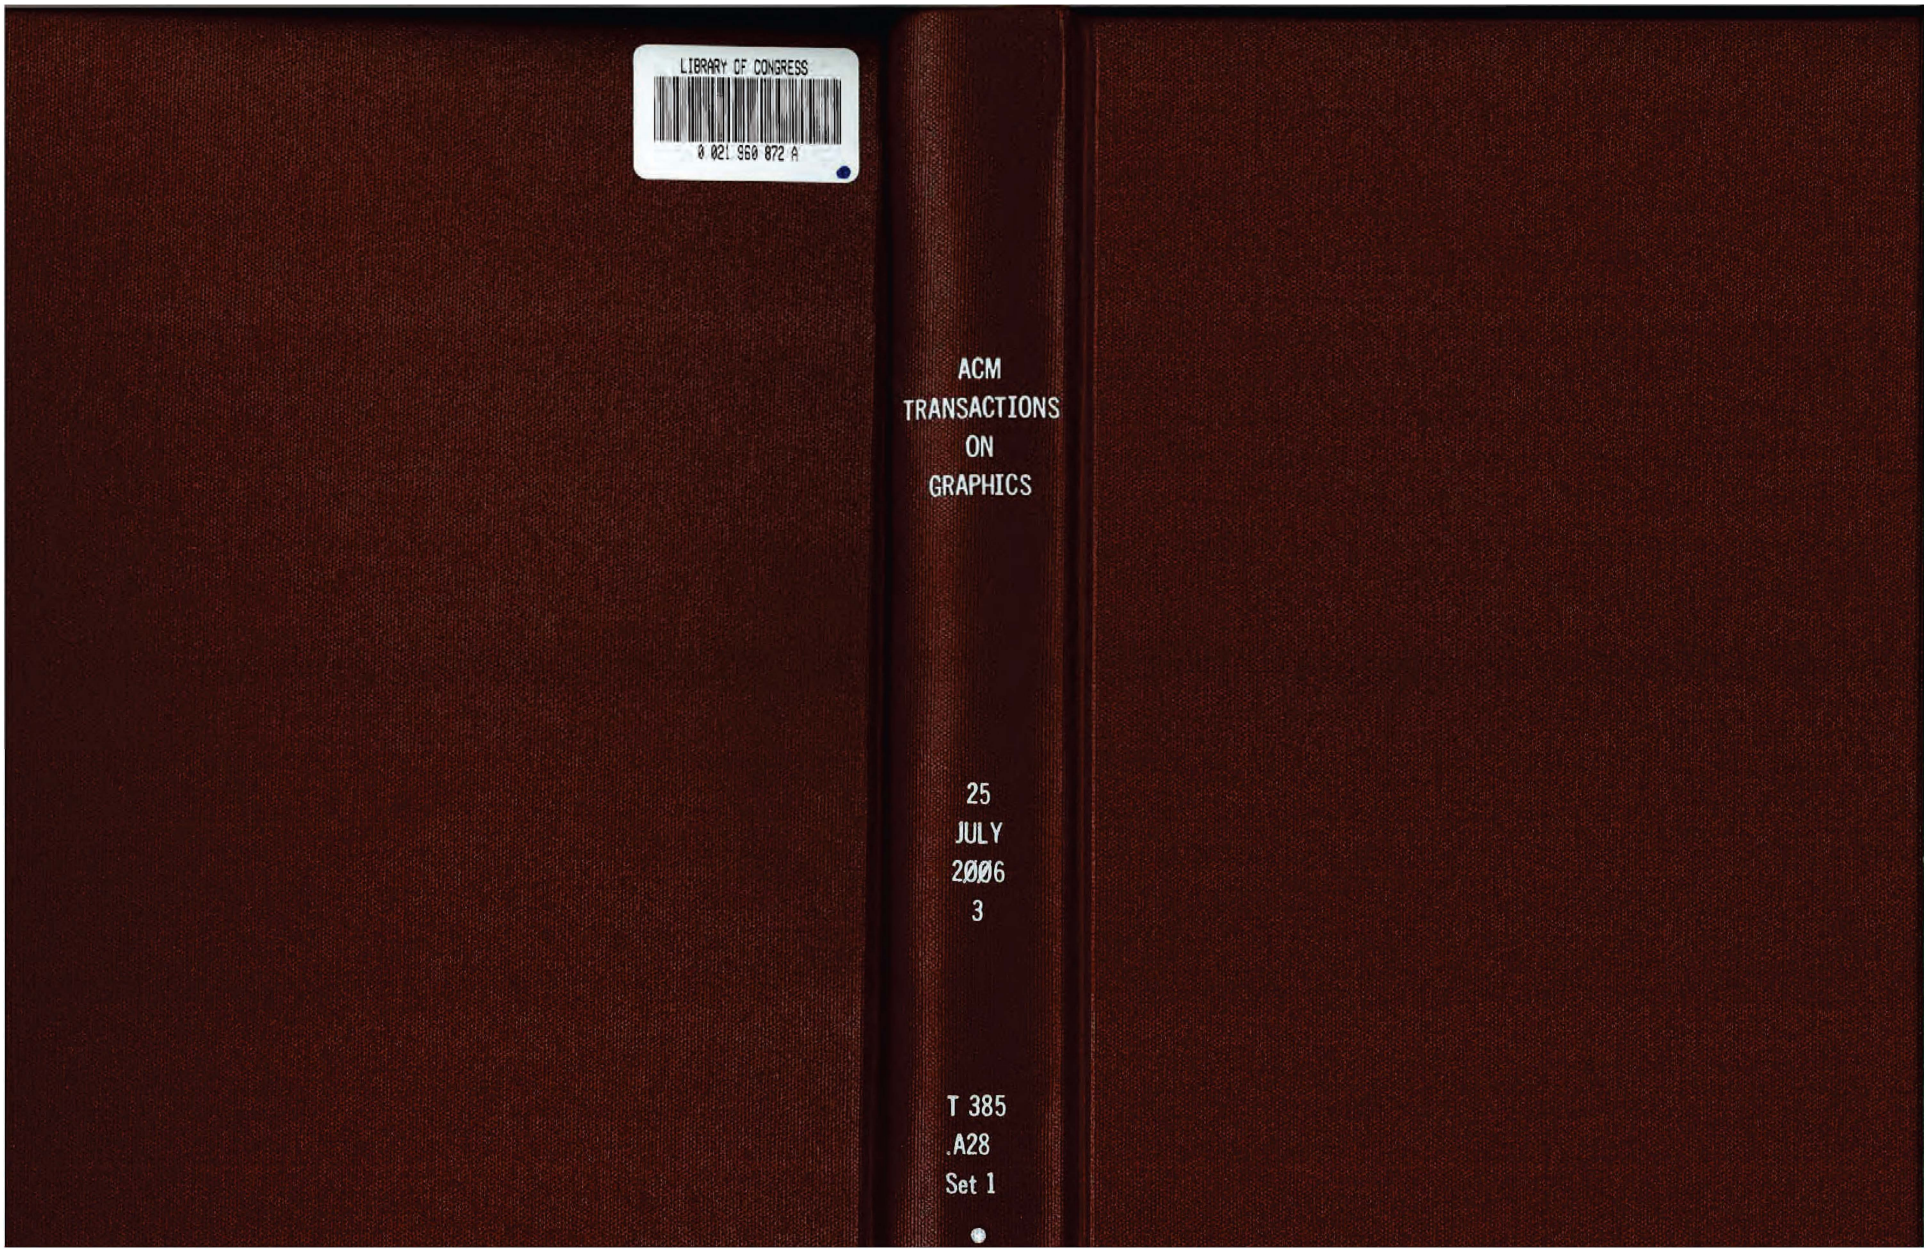
54. I reserve the right to supplement my opinions in the future to respond to any arguments that Patent Owner or its expert(s) may raise and to take-into account new information as it becomes available to me.

I hereby declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct, and that all statements made of my own knowledge are true and that all statements made on information and belief are believed to be true. I understand that willful false statements and the like are punishable by fine or imprisonment, or both (18 U.S.C. § 1001).

Executed on December 12, 2022 in Williamsburg, Virginia.

JAMES L. MULLINS, Ph.D

APPENDIX A-1


LIBRARY OF CONGRESS
0 021 960 872 A
ACM
TRANSACTIONS
ON
GRAPHICS
25
JULY
2006
3
T 385
.A28
Set 1

acm Transactions on Graphics

July 2006
Volume 25 Number 3

Proceedings of ACM SIGGRAPH 2006, Boston, MA

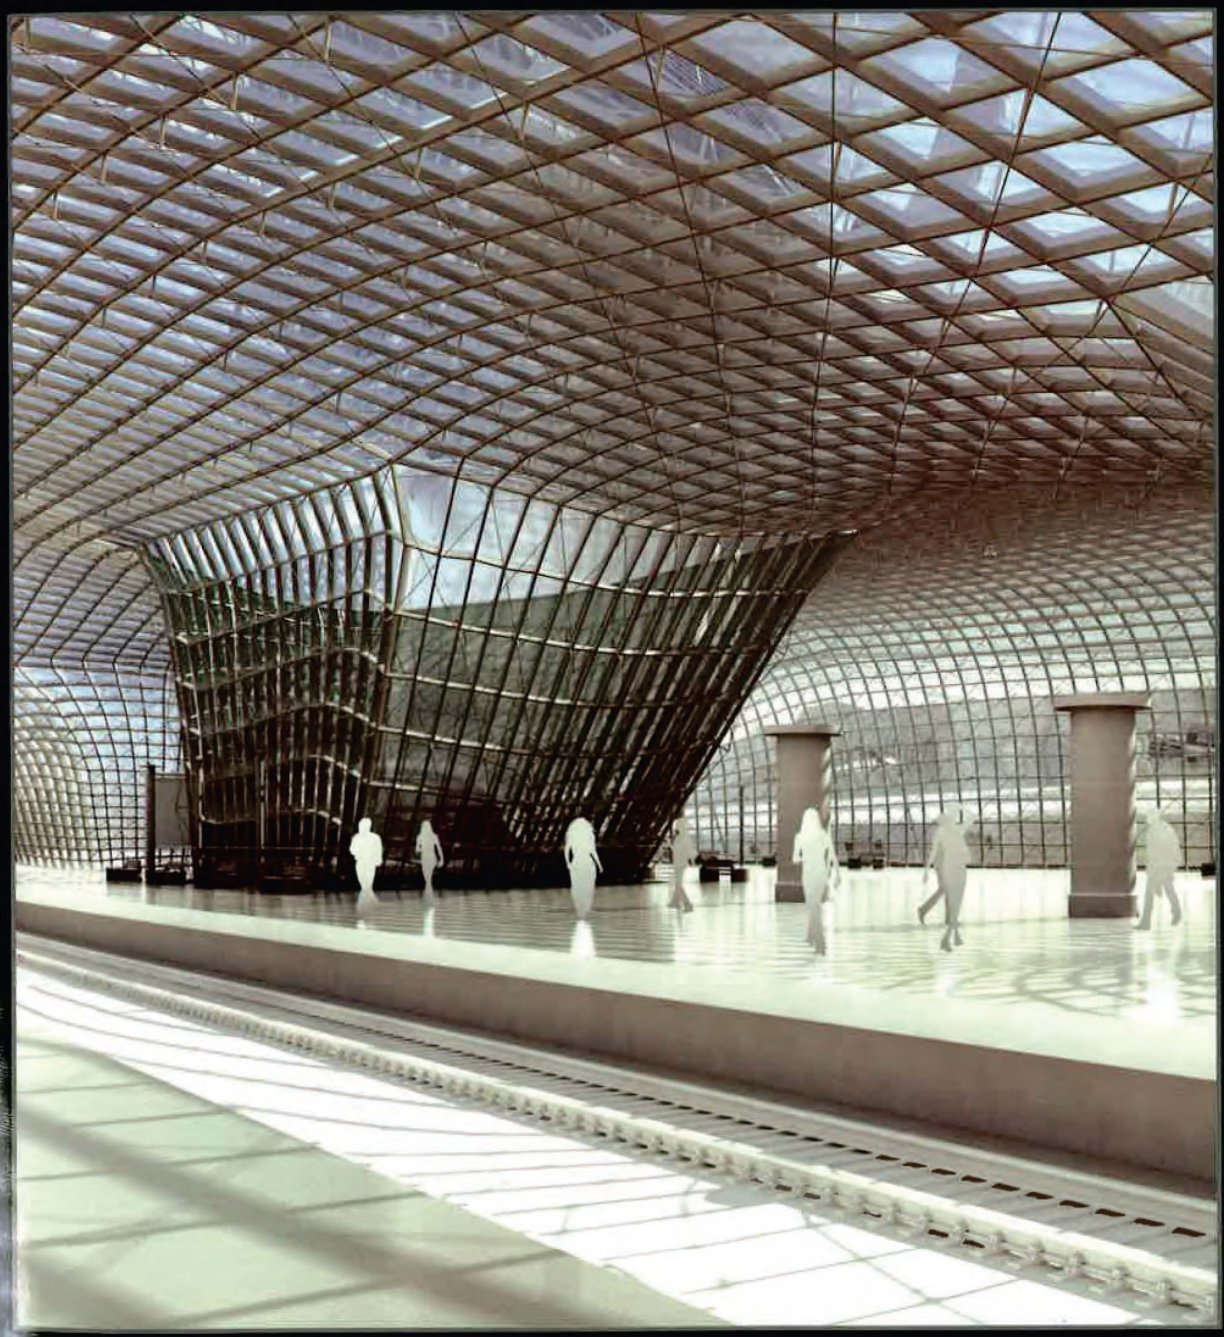


Exhibit 1010 page 21 of 78
DENTAL IMAGING

ACM Transactions on Graphics

ACM
1515 Broadway
New York, NY 10036
Tel.: (212) 869-7440
Fax: (212) 869-0481

TOG Home Page: http://www.acm.org/pubs/tog

Editor-in-Chief
John C. Hart — Department of Computer Science / University of Illinois at Urbana-Champaign / Urbana, IL 61801 / email: jch@cs.uiuc.edu

Guest Editor
Julie Dorsey — Computer Science Department / Yale University / New Haven, CT 06511 / email: julie.dorsey@yale.edu

Online Editor
Eric Haines — Autodesk, Inc. / Ithaca, NY 14850 / email: erich@acm.org

Associate Editors
Nina Amenta — Department of Computer Science / University of California / Davis, CA 95616 / email: amenta@cs.ucdavis.edu
Chandrajit Bajaj — Department of Computer Science and TICAM / University of Texas / Austin, TX 78712 / email: bajaj@cs.utexas.edu
Mathieu Desbrun — Computer Science Department / University of Southern California / Los Angeles, CA 90089 / email: desbrun@usc.edu
Frédo Durand — Computer Graphic Group, CSAIL / Massachusetts Institute of Technology / Cambridge, MA 02139 / email: fredo@mit.edu
Michael Gleicher — Department of Computer Sciences / University of Wisconsin–Madison / Madison, WI 53706 / email: gleicher@cs.wisc.edu
Cindy Grimm — Department of Computer Science / Washington University / St. Louis, MO 63130 / email: cmg@cs.wustl.edu
Radek Grzeszczuk — Intel Corporation / Santa Clara, CA 95052 / email: radek.grzeszczuk@intel.com
Christoph Hoffmann — Computer Sciences Department / Purdue University / West Lafayette, IN 47907-1398 / email: cmh@cs.purdue.edu
Hugues Hoppe — Microsoft Research / Redmond, WA 98052 / email: hhoppe@microsoft.com
Leonard McMillan — Department of Computer Science / University of North Carolina / Chapel Hill, NC 27599 / email: mcmillan@cs.unc.edu
Jörg Peters — Department of CISE / University of Florida / Gainesville, FL 32611-6120 / email: jorg@cise.ufl.edu
Frits H. Post — Faculty of Technical Informatics / Delft Technical University / Delft, The Netherlands / email: frits.post@twi.tudelft.nl
Przemyslaw Prusinkiewicz — Department of Computer Science / University of Calgary / Calgary, Alberta T2N 1N4, Canada / email: pwp@cpsc.ucalgary.ca
Thomas W. Sederberg — Department of Computer Science / Brigham Young University / Provo, UT 84602 / email: tom@cs.byu.edu
Peter Shirley — Computer Science Department / University of Utah / Salt Lake City, UT 84112 / email: shirley@cs.utah.edu
François Sillion — iMAGIS-GRAVIR, ZIRST / 38330 Montbonnot, France / email: francois.sillion@imag.fr
Jack Tumblin — Computer Science Department / Northwestern University / Evanston, IL 60201 / email: jet@cs.northwestern.edu
Michiel van de Panne — Department of Computer Science / University of British Columbia / Vancouver BC V6T 1Z4 Canada / email: van@cs.ubc.ca

Headquarters Staff
Mark Mandelbaum — Director of Publications
Jono Hardjowirogo — Publisher, Associate Director of Publications
Roma Simon — Managing Editor

ACM Transactions on Graphics (ISSN: 0730-0301) is published quarterly in January, April, July, and October by the Association for Computing Machinery (ACM), 1515 Broadway, New York, NY 10036. Periodicals class postage paid at New York, NY 10011, and at additional mailing offices. Printed in the U.S.A. POSTMASTER: Send address changes to ACM Transactions on Graphics, ACM, 1515 Broadway, New York, NY 10036.

For manuscript submissions, subscription, and change of address information, see inside backcover.

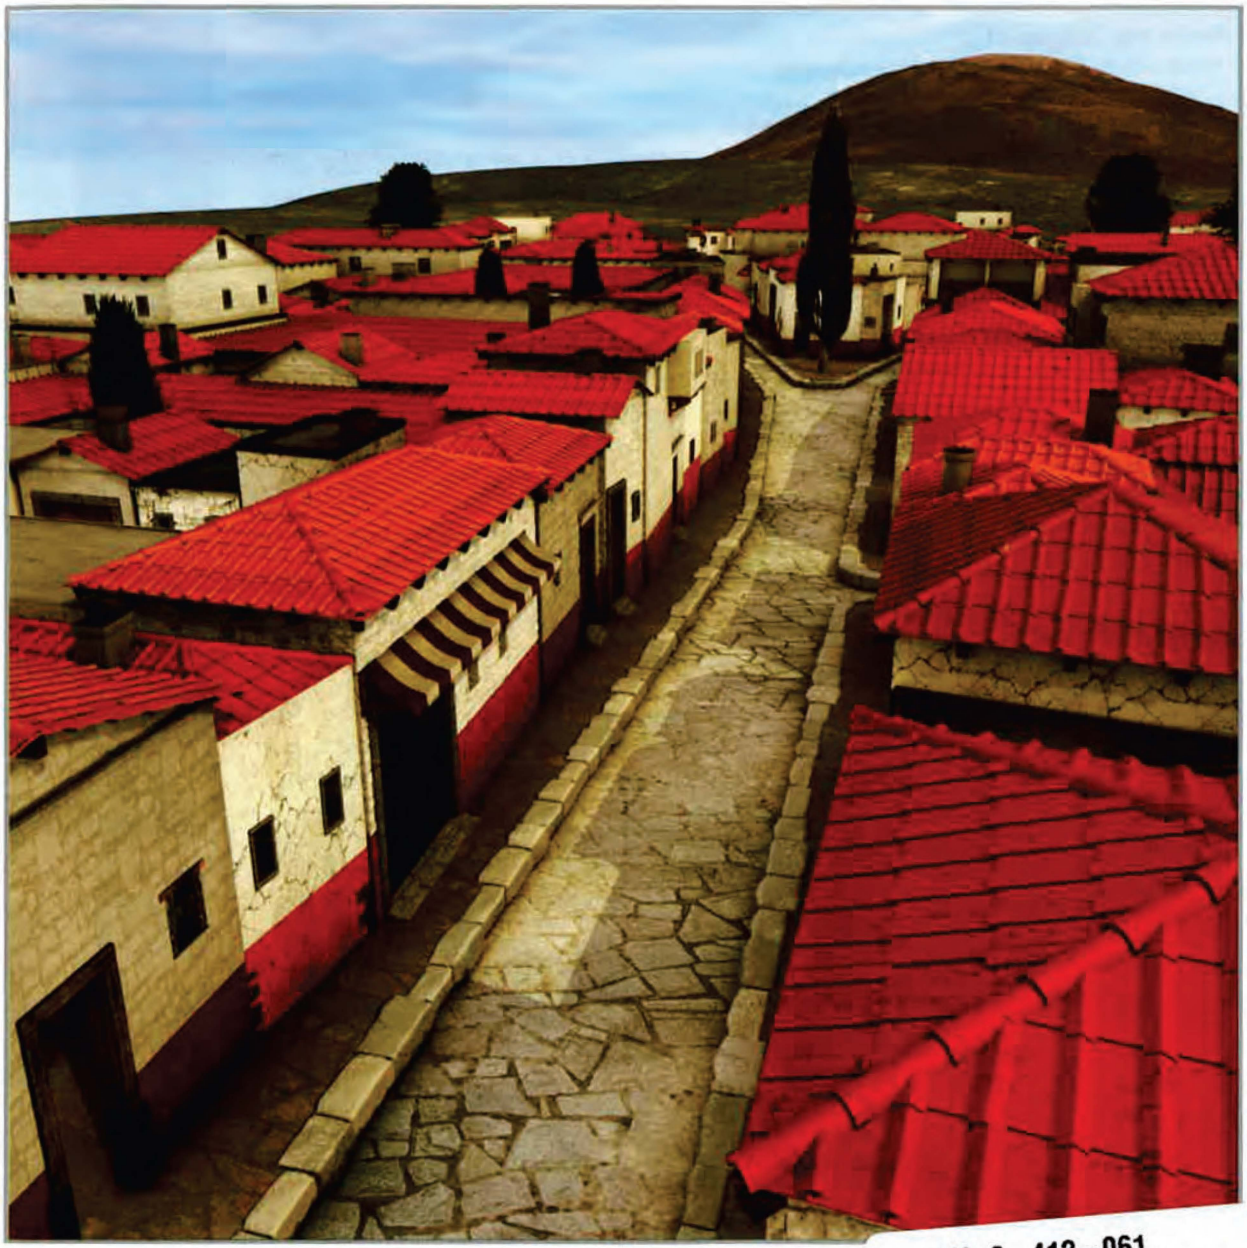
Copyright © 2006 by the Association for Computing Machinery (ACM). Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Department, ACM, Inc. Fax +1 212-869-0481 or email <permissions@acm.org>.

For other copying of articles that carry a code at the bottom of the first or last page or screen display, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

acm Transactions on Graphics

July 2006
Volume 25 Number 3

Proceedings of ACM SIGGRAPH 2006, Boston, MA



TX 6-412-061



Exhibit 1010 page 23 of 78
DENTAL IMAGING

The Association for Computing Machinery, Inc.
1515 Broadway
New York, New York 10036

Copyright © 2006 by the Association for Computing Machinery, Inc (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. **Copyright for components of this work owned by others than ACM must be honored.** Abstracting with credit is permitted.

To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Department, ACM, Inc. Fax +1-212-869-0481 or e-mail `permissions@acm.org`.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Notice to Past Authors of ACM-Published Articles
ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that was previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform `permissions@acm.org`, stating the title of the work, the author(s), and where and when published.

ACM ISSN 0730-0301
ACM Order Number 428060

Printed in the USA

Additional copies may be ordered from ACM.

ACM
1515 Broadway
New York, NY 10036
+1-212-869-7440
+1-212-869-0481 (fax)



Exhibit 1010 page 24 of 78
DENTAL IMAGING

Table of Contents

Preface 478
Editorial 480

2006 Computer Graphics Achievement Award 482
2006 Significant New Researcher Award 483
2006 Outstanding Service Award 484

Paper Sessions – Monday, 31 July 2006

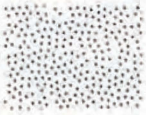
8:30 – 10:15 **Sampling and Ray Tracing**
Session Chair: Kavita Bala, Cornell University



Ray Tracing Animated Scenes using Coherent Grid Traversal 485
Ingo Wald, Thiago Ize, Andrew Kensler, Aaron Knoll, Steven G. Parker



Guided Visibility Sampling 494
Peter Wonka, Michael Wimmer, Kaichi Zhou, Stefan Maierhofer, Gerd Hesina, Alexander Reshetov



A Spatial Structure for Fast Poisson-Disk Sample Generation 503
Daniel Dunbar, Greg Humphreys



Recursive Wang Tiles for Real-Time Blue Noise 509
Johannes Kopf, Daniel Cohen-Or, Oliver Deussen, Dani Lischinski

8:30 – 10:15 **Image Processing**
Session Chair: Yizhou Yu, University of Illinois at Urbana-Champaign



Fast Median and Bilateral Filtering 519
Ben Weiss



Hybrid Images 527
Aude Oliva, Antonio Torralba, Philippe G. Schyns



Image Deformation Using Moving Least Squares 533
Scott Schaefer, Travis McPhail, Joe Warren



Appearance-Space Texture Synthesis 541
Sylvain Lefebvre, Hugues Hoppe

Exhibit 1010 page 25 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Monday, 31 July 2006

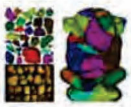
10:30 – 12:15 **Shape Matching and Symmetry**
Session Chair: Ioana Boier-Martin, IBM T.J. Watson Research Center



A Planar Reflective Symmetry Transform for 3D Shapes 549
Joshua Podolak, Philip Shilane, Aleksey Golovinskly, Szymon Rusinkiewicz, Thomas Funkhouser



Partial and Approximate Symmetry Detection for 3D Geometry 560
Niloy J. Mitra, Leonidas J. Guibas, Mark Pauly



Reassembling Fractured Objects by Geometric Matching 569
Qi-Xing Huang, Simon Flöry, Natasha Gelfand, Michael Hofer, Helmut Pottmann



Perfect Spatial Hashing 579
Sylvain Lefebvre, Hugues Hoppe

10:30 – 12:15 **Shape Modeling and Textures**
Session Chair: David Ebert, Purdue University



SmoothSketch: 3D Free-Form Shapes From Complex Sketches 589
Olga A. Karpenko, John F. Hughes



Image-based Plant Modeling 599
Long Quan, Ping Tan, Gang Zeng, Lu Yuan, Jingdong Wang, Sing Bing Kang



Interactive Decal Compositing With Discrete Exponential Maps 605
Ryan Schmidt, Cindy Grimm, Brian Wyvill



Procedural Modeling of Buildings 614
Pascal Müller, Peter Wonka, Simon Haegler, Andreas Ulmer, Luc Van Gool

Exhibit 1010 page 26 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Monday, 31 July 2006

3:45 – 6:00 **Image Manipulation**
Session Chair: Alexei Efros, Carnegie Mellon University


Color Harmonization 624
Daniel Cohen-Or, Olga Sorkine, Ran Gal, Tommer Leyvand, Ying-Qing Xu


Drag-and-Drop Pasting 631
Jiaya Jia, Jian Sun, Chi-Keung Tang, Heung-Yeung Shum


Two-scale Tone Management for Photographic Look 637
Soonmin Bae, Sylvain Paris, Frédo Durand


Interactive Local Adjustment of Tonal Values 646
Dani Lischinski, Zeev Farbman, Matt Uyttendaele, Rick Szeliski


Image-Based Material Editing 654
Erum Arif Khan, Erik Reinhard, Roland W. Fleming, Heinrich Bülthoff

Paper Sessions – Tuesday, 01 August 2006

8:30 – 10:15 **Surfaces**
Session Chair: Leif P. Kobbelt, RWTH Aachen University


Real-Time GPU Rendering of Piecewise Algebraic Surfaces 664
Charles Loop, Jim Blinn


Point-Sampled Cell Complexes 671
Anders Adamson, Marc Alexa


Geometric Modeling With Conical Meshes and Developable Surfaces 681
Yang Liu, Helmut Pottmann, Johannes Wallner, Yong-Liang Yang, Wenping Wang


Mesh Quilting for Geometric Texture Synthesis 690
Kun Zhou, Xin Huang, Xi Wang, Yiying Tong, Mathieu Desbrun, Baining Guo, Heung-Yeung Shum

Exhibit 1010 page 27 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Tuesday, 01 August 2006

10:30 – 12:15 **HDR and Systems**
Session Chair: Greg Ward, BrightSide Technologies



High Dynamic Range Texture Compression for Graphics Hardware 698
Jacob Munkberg, Petrik Clarberg, Jon Hasselgren, Tomas Akenine-Möller



High Dynamic Range Texture Compression 707
Kimmo Roimela, Tomi Aarnio, Joonas Itäranta

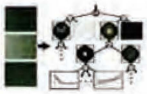


Backward Compatible High Dynamic Range MPEG Video Compression 713
Rafal Mantiuk, Alexander Efremov, Karol Myszkowski, Hans-Peter Seidel



The Direct3D 10 System 724
David Blythe

1:45 – 3:30 **Appearance Representation**
Session Chair: Holly Rushmeier, Yale University



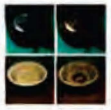
Inverse Shade Trees for Non-Parametric Material Representation and Editing 735
Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, Szymon Rusinkiewicz



A Compact Factored Representation of Heterogeneous Subsurface Scattering 746
Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, Philip Dutré



Appearance Manifolds for Modeling Time-Variant Appearance of Materials 754
Jiaping Wang, Xin Tong, Stephen Lin, Minghao Pan, Chao Wang, Hujun Bao, Baining Guo, Heung-Yeung Shum



Time-Varying Surface Appearance: Acquisition, Modeling and Rendering 762
Jinwei Gu, Chien-I Tu, Ravi Ramamoorthi, Peter Belhumeur, Wojciech Matusik, Shree Nayar

Exhibit 1010 page 28 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Tuesday, 01 August 2006

3:45 – 5:30 **Matting & Deblurring**
Session Chair: Paul Debevec, USC Centers for Creative Technologies


Flash Matting 772
Jian Sun, Yin Li, Sing Bing Kang, Heung-Yeung Shum


Natural Video Matting using Camera Arrays 779
Neel Joshi, Wojciech Matusik, Shai Avidan


Removing Camera Shake from a Single Photograph 787
Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T. Roweis, William T. Freeman


Coded Exposure Photography: Motion Deblurring using Fluttered Shutter 795
Ramesh Raskar, Amit Agrawal, Jack Tumblin

Paper Sessions – Wednesday, 02 August 2006

8:30 – 10:15 **Fluids**
Session Chair: Ming Lin, University of North Carolina at Chapel Hill


Efficient Simulation of Large Bodies of Water by Coupling
Two and Three Dimensional Techniques 805
Geoffrey Irving, Eran Guendelman, Frank Losasso, Ronald Fedkiw


Multiple Interacting Liquids 812
Frank Losasso, Tamar Shinar, Andrew Selle, Ronald Fedkiw


Fluid Animation With Dynamic Meshes 820
Bryan M. Klingner, Bryan E. Feldman, Nuttapong Chentanez, James F. O'Brien


Model Reduction for Real-Time Fluids 826
Adrien Treuille, Andrew Lewis, Zoran Popović

Exhibit 1010 page 29 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Wednesday, 02 August 2006

8:30 – 10:15 **Image Collections**
Session Chair: Wojciech Matusik, Mitsubishi Electric Research Laboratories (MERL)


Photo Tourism: Exploring Photo Collections in 3D 835
Noah Snavely, Steven M. Seitz, Richard Szeliski


AutoCollage 847
Carsten Rother, Lucas Bordeaux, Youssef Hamadi, Andrew Blake


Photographing Long Scenes with Multi-Viewpoint Panoramas 853
Aseem Agarwala, Maneesh Agrawala, Michael Cohen, David Salesin, Richard Szeliski


Schematic Storyboarding for Video Visualization and Editing 862
Dan B Goldman, Brian Curless, David Salesin, Steven M. Seitz

10:30 – 12:15 **Motion Capture**
Session Chair: Doug L. James, Carnegie Mellon University


Interaction Capture and Synthesis 872
Paul G. Kry, Dinesh K. Pai


Capturing and Animating Skin Deformation in Human Motion 881
Sang Il Park, Jessica K. Hodgins


Compression of Motion Capture Databases 890
Okan Arikan


Motion Patches: Building Blocks for Virtual Environments Annotated With Motion Data 898
Kang Hoon Lee, Myeong Geol Choi, Jehee Lee

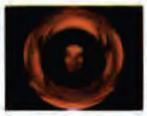
Exhibit 1010 page 30 of 78
DENTAL IMAGING

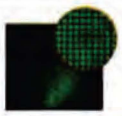
Table of Contents

Paper Sessions – Wednesday, 02 August 2006

10:30 – 12:15 **Image Capture**
Session Chair: Szymon Rusinkiewicz, Princeton University


Projection Defocus Analysis for Scene Capture and Image Display 907
Li Zhang, Shree Nayar


Multiview Radial Catadioptric Imaging for Scene Capture 916
Sujit Kuthirummal, Shree K. Nayar


Light Field Microscopy 924
Marc Levoy, Ren Ng, Andrew Adams, Matthew Footer, Mark Horowitz


Fast Separation of Direct and Global Components of a Scene
using High Frequency Illumination 935
Shree K. Nayar, Gurunandan Krishnan, Michael D. Grossberg, Ramesh Raskar

3:45 – 6:00 **Precomputed Transfer**
Session Chair: Fabio Pellacini, Dartmouth College


Real-Time BRDF Editing in Complex Lighting 945
Aner Ben-Artzi, Ryan Overbeck, Ravi Ramamoorthi


Generalized Wavelet Product Integral for Rendering Dynamic Glossy Objects 955
Weifeng Sun, Amar Mukherjee


All-Frequency Precomputed Radiance Transfer using Spherical Radial Basis Functions
and Clustered Tensor Approximation 967
Yu-Ting Tsai, Zen-Chung Shih


Real-Time Soft Shadows in Dynamic Scenes using Spherical Harmonic Exponentiation 977
Zhong Ren, Rui Wang, John Snyder, Kun Zhou, Xinguo Liu, Bo Sun,
Peter-Pike Sloan, Hujun Bao, Qunsheng Peng, Baining Guo


Precomputed Acoustic Transfer: Output-Sensitive, Accurate Sound Generation
for Geometrically Complex Vibration Sources 987
Doug L. James, Jernej Barbič, Dinesh K. Pai

Exhibit 1010 page 31 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Thursday, 03 August 2006

8:30 – 10:15 **Appearance Modeling**
Session Chair: Philip Dutré, Katholieke Universiteit Leuven



Photorealistic Rendering of Rain Streaks 996
Kshitiz Garg, Shree K. Nayar



Acquiring Scattering Properties of Participating Media by Dilution 1003
Srinivas G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, Henrik Wann Jensen



Analysis of Human Faces using a Measurement-Based Skin Reflectance Model 1013
Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jansen, Markus Gross



A Statistical Model for Synthesis of Detailed Facial Geometry 1025
Aleksey Golovinskiy, Wojciech Matusik, Hanspeter Pfister, Szymon Rusinkiewicz, Thomas Funkhouser

8:30 – 10:15 **Meshes**
Session Chair: Joe Warren, Rice University



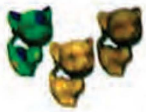
Modified Subdivision Surfaces with Continuous Curvature 1035
Adi Levin



Edge Subdivision Schemes and the Construction of Smooth Vector Fields 1041
Ke Wang, Weiwei, Yiying Tong, Mathieu Desbrun, Peter Schröder



Streaming Computation of Delaunay Triangulations 1049
Martin Isenberg, Yuanxin Liu, Jonathan Shewchuk, Jack Snoeyink



Spectral Surface Quadrangulation 1057
Shen Dong, Peer-Timo Bremer, Michael Garland, Valerio Pascucci, John C. Hart

Exhibit 1010 page 32 of 78
DENTAL IMAGING

Table of Contents

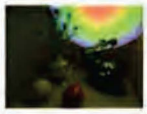
Paper Sessions – Thursday, 03 August 2006

10:30 – 12:15 **Light Transport**
Session Chair: Ravi Ramamoorthi, Columbia University


Simulating Multiple Scattering in Hair Using a Photon Mapping Approach 1067
Jonathan T. Moon, Stephen R. Marschner


Statistical Acceleration for Animated Global Illumination 1075
Mark Meyer, John Anderson


Multidimensional Lightcuts 1081
Bruce Walter, Adam Arbree, Kavita Bala, Donald P. Greenberg


Direct-to-Indirect Transfer for Cinematic Relighting 1089
Miloš Hašan, Fabio Pellacini, Kavita Bala

10:30 – 12:15 **Shape Deformation**
Session Chair: Marc Alexa, Technische Universität Berlin


Editing Arbitrarily Deforming Surface Animations 1098
Scott Kircher, Michael Garland


A Fast Multigrid Algorithm for Mesh Deformation 1108
Lin Shi, Yizhou Yu, Nathan Bell, Wei-Wen Feng


Vector Field Based Shape Deformations 1118
Wolfram von Funck, Holger Theisel, Hans-Peter Seidel


Subspace Gradient Domain Mesh Deformation 1126
Jin Huang, Xiaohan Shi, Xinguo Liu, Kun Zhou, Li-Yi Wei, Shang-Hua Teng, Hujun Bao, Baining Guo, Heung-Yeung Shum

Exhibit 1010 page 33 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Thursday, 03 August 2006

1:45 – 3:30 **Numerical and Geometric Algorithms and Crowds**
Session Chair: John Snyder, Microsoft Research



Locally Adapted Hierarchical Basis Preconditioning 1135
Richard Szeliski



Fast Proximity Computation among Deformable Models
Using Discrete Voronoi Diagrams 1144
Avneesh Sud, Naga Govindaraju, Russell Gayle, Ilknur Kabul, Dinesh Manocha



Resolving Surface Collisions through Intersection Contour Minimization 1154
Pascal Volino, Nadia Magnenat-Thalmann



Continuum Crowds 1160
Adrien Treuille, Seth Cooper, Zoran Popović

1:45 – 3:30 **Animation**
Session Chair: Dinesh K. Pai, Rutgers University and The University of British Columbia



The Cartoon Animation Filter 1169
Jue Wang, Steven M. Drucker, Maneesh Agrawala, Michael F. Cohen



Inverse Kinematics for Reduced Deformable Models 1174
Kevin G. Der, Robert W. Sumner, Jovan Popović



Super-Helices for Predicting the Dynamics of Natural Hair 1180
Florence Bertails, Basile Audoly, Marie-Paule Cani, Bernard Querleux, Frédéric Leroy, Jean-Luc Lévéque



Heads Up! Biomechanical Modeling and Neuromuscular Control of the Neck 1188
Sung-Hee Lee, Demetri Terzopoulos

Exhibit 1010 page 34 of 78
DENTAL IMAGING

Table of Contents

Paper Sessions – Thursday, 03 August 2006

3:45 – 5:30 **Non-Photorealistic Rendering**
Session Chair: Aaron Hertzmann, University of Toronto



Exaggerated Shading for Depicting Shape and Detail 1199
Szymon Rusinkiewicz, Michael Burns, Doug DeCarlo



Image Enhancement by Unsharp Masking the Depth Buffer 1206
Thomas Luft, Carsten Colditz, Oliver Deussen



Manga Colorization 1214
Yingge Qu, Tien-Tsin Wong, Pheng-Ann Heng



Real-Time Video Abstraction 1221
Holger Winnemöller, Sven C. Olsen, Bruce Gooch

Committees and Reviewers 1227
Exhibitors 1235
ACM SIGGRAPH Professional and Student Chapters 1236
Cover Image Credits 1239
Author Index 1240

Exhibit 1010 page 35 of 78
DENTAL IMAGING

This material may be protected by Copyright law (Title 17 U.S. Code)

Projection Defocus Analysis for Scene Capture and Image Display *

Li Zhang Shree Nayar
Columbia University

Abstract

In order to produce bright images, projectors have large apertures and hence narrow depths of field. In this paper, we present methods for robust scene capture and enhanced image display based on projection defocus analysis. We model a projector's defocus using a linear system. This model is used to develop a novel temporal defocus analysis method to recover depth at each camera pixel by estimating the parameters of its projection defocus kernel in frequency domain. Compared to most depth recovery methods, our approach is more accurate near depth discontinuities. Furthermore, by using a coaxial projector-camera system, we ensure that depth is computed at all camera pixels, without any missing parts. We show that the recovered scene geometry can be used for refocus synthesis and for depth-based image composition. Using the same projector defocus model and estimation technique, we also propose a defocus compensation method that filters a projection image in a spatially-varying, depth-dependent manner to minimize its defocus blur after it is projected onto the scene. This method effectively increases the depth of field of a projector without modifying its optics. Finally, we present an algorithm that exploits projector defocus to reduce the strong pixelation artifacts produced by digital projectors, while preserving the quality of the projected image. We have experimentally verified each of our methods using real scenes.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning, display algorithm; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Depth cues, range data, shape.

Keywords: projector defocus, temporal defocus analysis, depth recovery, multi-focal projection, projector depixelation, refocus synthesis, image composition.

1 Introduction

Digital projection technologies, such as Digital Light Processing (DLP) and Liquid Crystal Displays (LCD), are increasingly used in consumer, commercial and scientific applications. In computer graphics and vision, video projectors have recently been used as per-pixel controllable light sources for real-time shape acquisition [Huang et al. 2003; Zhang et al. 2004; Davis et al. 2005; Koninckx et al. 2005], for complex appearance capture [Levoy et al. 2004; Sen et al. 2005] and control [Raskar et al. 2001; Grossberg et al. 2004; Bimber et al. 2005]. All these applications require the projectors to be focused for best performance. In practice, projectors are built with large apertures to maximize their brightness. As a result, virtually all projectors have very narrow depths of field; they are designed to produce focused images on a single fronto-parallel screen. An analysis of the defocus properties of projectors is therefore beneficial as it could lead to new methods that take advantage of, as well as compensate for, projection defocus.

In this paper, we provide the first systematic analysis of projector defocus and demonstrate its applications to robust scene capture and enhanced image display. We first present a simple linear model for projector defocus. Based on this model, we present a frequency-domain method for estimating the spatially-varying defocus kernel of a projector. The kernel estimated at each scene point is used to recover the 3D geometry of the scene. Based on the estimated kernel, we also present a technique that computationally manipulates an input image to minimize its defocus blur when it is projected onto a non-planar scene. Finally, we demonstrate that defocusing can be put to good use. A slight amount of defocusing, in conjunction with a compensation algorithm, can be used to reduce pixelation, a strong artifact produced by all digital projectors. Specifically, this paper makes the following three contributions:

Scene Geometry using Temporal Defocus Analysis: We propose a method called temporal defocus analysis that estimates depth at each camera pixel, independently, without using information from neighboring pixels (Section 3). Our method, compared to existing methods, such as stereo or even traditional depth from focus/defocus algorithms, is more accurate near depth discontinuities. Since the method is not based on triangulation, we can use a coaxial projector-camera system and compute depth at all camera pixels, without any missing parts. These advantages make the method uniquely suited to computer graphics applications like refocus synthesis and image composition, both of which we demonstrate.

Focused Projection at Multiple Depths: We present an iterative, spatially-varying filtering algorithm that compensates for defocus blur based on scene geometry (Section 4). This technique effectively increases the depth of field of a projector without modifying its optics. As a result, we are able to use a single projector to project well-focused images on multiple planes that are at different depths as well as on curved projection surfaces. We believe this capability addresses an important limitation of current projectors and widens their applicability in the real world.

Depixelation by Defocusing: Finally, we generalize our defocus compensation algorithm to reduce the strong pixelation artifacts produced by all digital projectors (Section 5). The key idea is to slightly defocus the projector so as to attenuate the high frequencies produced by pixelation and use the defocus compensation algorithm to make up for the induced projector defocus. This method is especially suited for projecting high resolution images using a low resolution projector, and we show several examples of results to illustrate its effects.

2 Previous Work

Many methods have been proposed in computer vision to recover 3D shape from images, including multi-view triangulation methods, single-view photometric methods, and camera focus/defocus methods. Triangulation-based methods [Faugeras 1993], e.g., structure from motion and stereo, require a point to be visible in at least two views to be reconstructed. Scenes with complex occlusions remain a challenging problem for these methods. Photometric methods [Horn and Brooks 1989], e.g., photometric stereo and shape from shading, estimate surface normals instead of surface depth. Converting normals to depths is an ill-posed problem for scenes with depth discontinuities.

*This work was conducted at the Computer Vision Laboratory at Columbia University. It was supported by an ITR grant from the National Science Foundation (No. IIS-00-85864). The authors thank Gurunandan Krishnan for his help with the experiments and the video and Anne Fleming for her help with the audio recording and the proofreading. The authors also thank www.dpreview.com for permitting us using their images as some of the examples shown in the paper.

Copyright © 2006 by the Association for Computing Machinery, Inc.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.
© 2006 ACM 0730-0301/06/0700-0907 $5.00

Exhibit 1010 page 36 of 78
DENTAL IMAGING

Depth recovery methods based on camera focus and defocus, e.g. [Pentland 1987; Nayar and Nakagawa 1994; Nayar et al. 1996; Schechner et al. 2000; Favaro and Soatto 2005], have the potential to recover depth at every pixel, regardless of the scene complexity and occlusions. To resolve the focus ambiguity of textureless surfaces, patterns can be projected to force scene texture [Girod and Scherock 1989]. However, camera defocus kernels depend on local surface geometry. To simplify the kernel analysis, most previous works assume that, within a small spatial window, the surface depth is constant, the so-called *equalfocal assumption*. This assumption smears shape details and is invalid across depth discontinuities. To alleviate this problem, Jin et al. [2002] and Rajagopalan and Chaudhuri [1997] estimate depths for all pixels simultaneously via a large scale energy minimization, which is computationally expensive and prone to local minima. Our key observation is that, unlike camera defocus, the kernel for projector defocus is *scene independent*, for most scene surfaces. Specifically, when a 3D scene point sees the entire projector aperture, its defocus kernel depends only on its distance to the projector lens and not on its neighboring surface geometry. This difference arises from the fact that projector defocus convolution happens on the projector's image plane while camera defocus convolution happens on the scene surface. Exploiting this scene-independent property, we project a shifting pattern over the scene and compute depth at each pixel using just its intensity variation over time. Without using the equalfocal assumption, our method works well at depth discontinuities. Furthermore, it is simple and not subject to local minima.

Our temporal per-pixel defocus analysis is inspired by previous structured light range finding methods. In particular, Kanade et al. [1991] and Curless and Levoy [1995] use temporal intensity variation to resolve correspondences between camera pixels and a sweeping laser stripe. Huang et al. [2003] have developed a real-time range finder by sweeping periodic sinusoidal stripes using a DLP projector. Zhang et al. [2004] and Davis et al. [2005] have generalized these ideas to space-time stereo. All these range finding techniques are based on the principle of triangulation, which cannot estimate depth for points that are visible to only the camera or the source but not both. Since our method is not based on triangulation, we are able to use a coaxial configuration where the camera and projector share the same optical center, and compute depth at all camera pixels (no missing parts). To our knowledge, the only other technique that is able to estimate a complete depth map with clean discontinuity boundaries is time-of-flight, e.g., [Gonzales-Banos and Davis 2004], which requires expensive specialized hardware. Raskar et al. [2004] proposed a method to detect depth discontinuities using multiple flash lights. The detected discontinuities can be used to enhance stereo matching algorithms. Compared to this work, our method directly generates reliable depth estimation for all pixels, including the ones near discontinuities.

In addition to scene capture, several methods have been proposed that use cameras to enable projectors to display "better" images. These include compensating distortions due to surface geometry [Raskar et al. 2003] and correcting brightness variations due to surface color and texture [Grossberg et al. 2004; Bimber et al. 2005]. In these works, the projectors are often used to display images onto scenes with considerable depth variation, where defocus blur is inevitable and most often spatially-varying. Complementary to these previous methods, our work seeks to compensate for defocus blur, and hence could be beneficial to the previous methods. In [Bimber and Emmerling 2006], a system is proposed that can project focused images on multiple planes at different depths. This system uses multiple projectors where each projector is focused on a single plane. In our case, we use a single projector to simultaneously project well-focused images at multiple depths. Another interesting related work is by Majumder and Greg [2001], in which multiple projectors at different focal settings are combined to generate depth of field effects at an interactive rate.

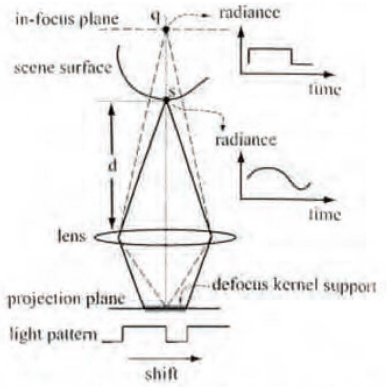


Figure 1: The principle of depth from projection defocus. Points at different distances to the projector lens exhibit different amounts of blur in their temporal radiance profile as a periodic illumination pattern is shifted across the scene.



(a) (b)

Figure 2: (a) A coaxial projector-camera system for depth from projection defocus. The system is made coaxial only to ensure that depth can be computed at all camera pixels. (b) The illumination pattern that is shifted across the scene to measure depth at each pixel, independently.

Deblurring is a well-studied topic in image processing where numerous techniques have been proposed, ranging from classical Wiener filtering and conjugate gradient optimization [Jain 1989] to more recent algorithms based on graph cuts [Raj and Zabih 2005] and belief propagation [Tappen et al. 2004]. Our technique is based on bound-constrained quadratic programming [Nocedal and Wright 1999], which is a variant of existing algorithms that best suits our formulation of the problem. As we show, our deblurring method can also reduce pixelation effects produced by digital projectors.

3 Depth from Projection Defocus

In this section, we formulate the problem of recovering 3D shape from projection defocus and present our solution.

3.1 Temporal Defocus Analysis

Consider a scene that is illuminated by structured light from a projector which is focused behind the scene, as shown in Figure 1. For a point **q** that is in focus, its irradiance comes from a single point on the projector's image plane. For a point **s** that is out of focus, its irradiance equals the convolution of its defocus kernel with the structured light pattern on the projector's image plane. Assuming that the surface is opaque, the radiance I of **s** along any given outgoing direction can be written as

$$I = \alpha f(\mathbf{x}; z) * P(\mathbf{x}) + \beta, \tag{1}$$

where $*$ denotes convolution, α is a factor depending on surface reflectance[1], β is the radiance due to the ambient light, $f(\mathbf{x}; z)$ is the

[1]To be precise, α takes into account all surface shading factors: BRDF, orientation with respect to the irradiance direction and the squared distance fall-off of the projector brightness.

Exhibit 1010 page 37 of 78
DENTAL IMAGING

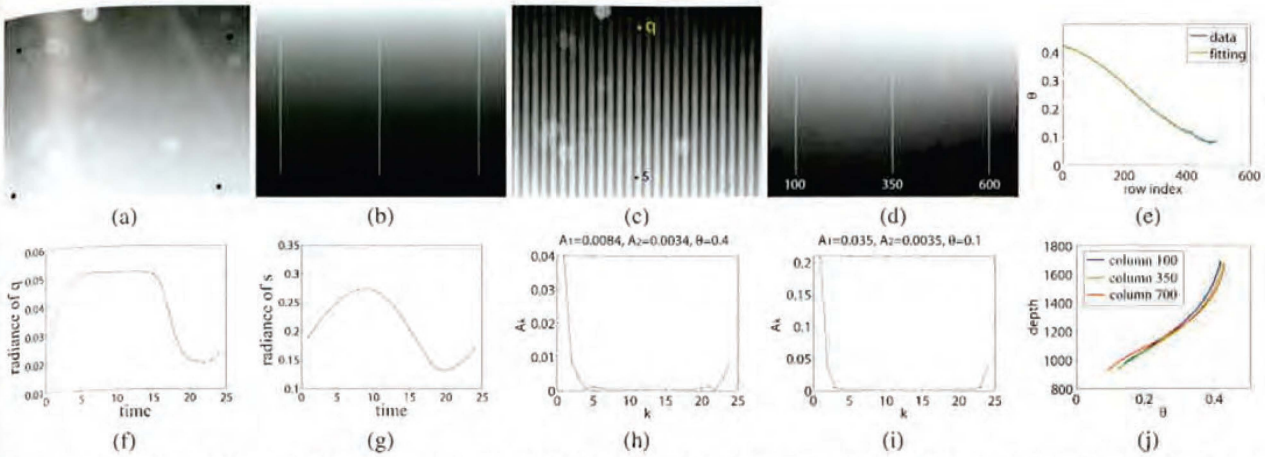


Figure 3: Illustration of the calibration procedure. (a) A white board with four markers tilted in front of the camera. (b) Depth map of the board computed using the markers. (c) The board under one of the stripe projection patterns. (d) θ map for the board. (e) θ values in column 350 of (d). (f,g) Temporal radiance profiles of the points $\mathbf{q}$ and $\mathbf{s}$, respectively. (h,i) Discrete-time Fourier series (DFS) of (f) and (g). (j) Mappings from θ to depth z for columns 100, 350 and 600, respectively.

defocus kernel, and $P(\mathbf{x})$ is the illumination pattern. The defocus kernel f depends on the depth z.[2] We now describe how to recover the depth z by estimating the kernel f from radiance measurements.

3.2 Depth Estimation

To estimate the kernel f, we note that Eq. (1) defines a linear system in which the projection pattern and the scene radiance are the input and the output, respectively. Estimating the kernel of a linear system is a classical problem in system identification [Ljung 1998] and we take a frequency-domain approach to solve it in our setting. Our basic idea is to shift an illumination pattern with a wide range of frequencies within it across the scene. The radiance of a surface point over time is then the response of its defocus kernel to the excitation by the illumination pattern. As the pattern is shifted, points at different distances to the projector exhibit different amounts of blur in their *temporal* radiance profile, as illustrated in Figure 1. We use this temporal blur for depth recovery.

Given the temporal radiance sequence, I_l, $l = 0,\ldots,L-1$, for a point, we quantify its blur by decomposing it into a discrete-time Fourier series (DFS) [Oppenheim and Willsky 1997] as

$$I_l = A_0 + \sum_{k=1}^{L-1} A_k \cos(\omega_k l - \phi_k), \tag{2}$$

where $\omega_k = \frac{2k\pi}{L}$, $A_k = (B_k^2 + C_k^2)^{\frac{1}{2}}$, $\phi_k = \arctan(B_k, C_k)$, $B_k = \frac{1}{L}\sum_{k=0}^{L-1} I_l \sin(\omega_k l)$ and $C_k = \frac{1}{L}\sum_{k=0}^{L-1} I_l \cos(\omega_k l)$. Since the kernel f is a low-pass filter, how quickly the coefficients A_k diminish with k is a measure of the amount of defocus, which in turn yields depth. Note that A_0 cannot be used to estimate depth because it depends on the ambient light β. Although both A_1 and A_2 are scaled by the albedo α, their ratio can be used to determine how severely the defocus kernel attenuates the second-order harmonic with respect to the first-order one. Therefore, we use the following ratio

$$\theta = \frac{A_2}{A_1} \tag{3}$$

as a measure of depth. In Eq. (3), $A_1 > A_2 > 0$ and $\theta \in [0,1]$ because f is a low-pass filter. Deriving the analytic mapping between θ and z is tedious but possible if we know precisely the optical design of the specific projector. However, analytically deriving a projector-specific mapping is not worth the effort as it would not apply to other projectors — each projector tends to have a unique optical design. Therefore, in the next section, we present a general data-driven approach that calibrates the $\theta - z$ mapping once and for all, for any given optical setting of any given projector-camera system.

[2] Due to lens aberration, f generally also varies across different pixels.

3.3 System Setup and Illumination Pattern

We have built a prototype camera-projector system to implement our depth from projection defocus method. Our system consists of an NEC LT260K DLP projector and a Basler A311f monochrome camera. To avoid shadows and occlusions, we approximately align the optical centers of the projector and the camera with a beam splitter (Edmund Optics stock #NT39-493), as shown in Figure 2(a)[3]. The projector is always focused on a plane behind the working volume to avoid a two-way defocus ambiguity. As the projector is quite bright[4], we stop-down the aperture of the camera to F11 so that it works approximately as a pinhole camera – any defocus introduced by the camera is negligible compared to that of the projector.

There are many choices of input sequences (excitation signals) in system identification theory. We have chosen to use a simple one – a binary periodic sequence 011011011011... with period 3, which is one type of M-sequence [Ljung 1998]. We encode this sequence as a stripe pattern in which each bit corresponds to an 8-pixel wide stripe, as shown in Figure 2(b). We shift this pattern, one pixel at a time, and take a total of $L = 24$ images for each experiment[5].

3.4 Calibration

We calibrate the mapping from the θ in Eq. (3) to the depth z in three steps. **Step 1**: We compute the correspondence between projector and camera pixels. This is achieved by projecting shifted sinusoids in both horizontal and vertical directions. The details of this procedure can be found in [Scharstein and Szeliski 2003]. **Step 2**:

[3] We align the camera and the projector by projecting an image onto a white board with a fence in front of it and adjusting the camera position until it does not see any shadows cast by the projector.

[4] During depth estimation, we have used the projector as a grayscale one by removing its color wheel from the light path to boost its brightness.

[5] In our experiments, we have found that the projector produces several undesirable effects: it vibrates at a high frequency, possibly due to the rotation of its fan; its brightness is not stable over time; the imprecise synchronization between the camera and projector causes measured radiances to fluctuate slightly over time. To resolve these issues, we repeat the shifting of the pattern about 85 times and compute a mean image for each shift. The camera runs at 60Hz and the total acquisition time is within a minute.

Exhibit 1010 page 38 of 78
DENTAL IMAGING

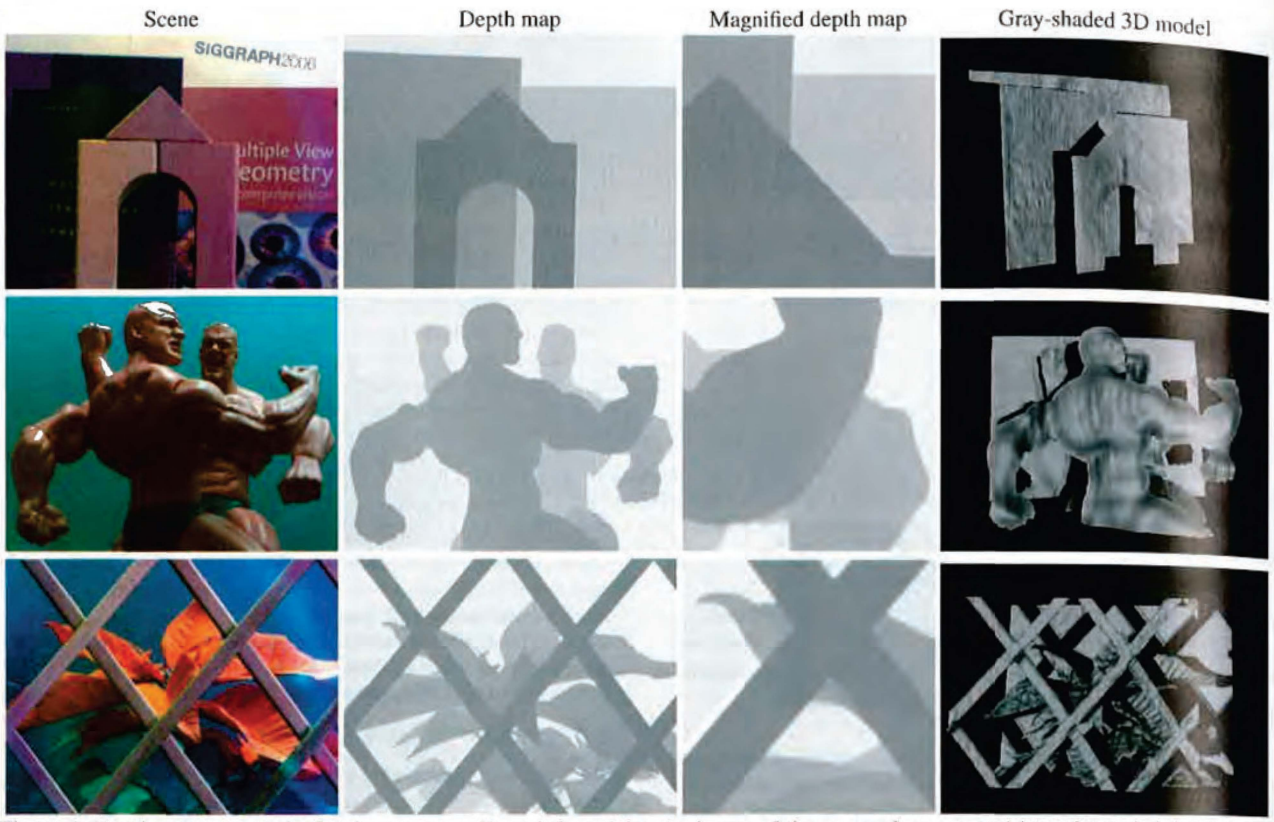


Figure 4: Depth recovery results for three scenes. From left to right: An image of the scene, the computed intensity-coded depth map, a close-up view of the depth map, a gray-shaded 3D model of the scene, as seen from a novel viewpoint. Notice the quality of the computed depth around scene discontinuities. **(Please see the companion video.)**

We tilt a foam board in front of the system and compute depth for each of its points. This is done by taking an image of the board with a few reference markers, as shown in Figure 3(a),[6] and computing the homography from the board to the projection plane. As detailed in [Zhang 2000], this homography allows us to estimate the position and orientation of the board[7], from which the depth of every point of the board can be easily computed. Figure 3(b) shows the intensity-coded depth map of the board. **Step 3**: We compute the θ values for all points on the board by shifting the stripe pattern and computing A_1 and A_2 for each pixel.

Figure 3(c) shows an image of the board under one of the shifted patterns, and (f) and (g) show the temporal radiance profiles of two points, **q** and **s**, on the board. As **s** is closer to the projector and more defocused than **q**, its temporal radiance profile is more blurred. Figures 3(h) and (i) show plots of the DFS coefficients A_k for the profiles in (f) and (g), respectively. Figure 3(d) shows the intensity-coded θ map and (e) shows a plot of the θ values (blue curve) for the column 350 in (d). The raw θ values include a small mount of noise and hence we fit a smooth curve which is shown as the green curve in (e). Using the depth estimation in Step 2, we can tabulate the mapping from θ to z. In our current implementation, we build a lookup table for each column, assuming that the defocus kernel is vertically invariant but has some horizontal variation. Figure 3(j) shows three θ–z mapping curves for the columns 100, 350 and 600. An even more comprehensive calibration would involve translating the board across the whole working volume and generating a lookup table for each pixel. Such a calibration would account for higher-order projector lens aberrations as well.

3.5 Depth Recovery Results

Once we know the mapping between θ and z, the depth recovery is straightforward. We take 24 images while shifting the illumination pattern across the scene. From these images, we compute θ for each camera pixel using Eqs. (2,3) and then transform the θ image to a depth image using the pre-computed lookup table. Figure 4 shows several depth recovery results. In the first row, we show results for a scene with books and wooden blocks that has simple boundaries. Our method recovers the sharp depth discontinuities at the boundaries. In the second row, we show results for a scene with two toy wrestlers with similar skin color. Again, our method can separate the two figures with clean boundaries, even though the surfaces have a specular component. In the third row, we show results for a scene with leaves behind a fence. This result shows that the performance of our method is independent of occlusion complexity. In fact, it has been shown [Scharstein and Szeliski 2003] that it is very difficult to obtain a depth map that is complete in any camera view for a scene like the one with leaves using triangulation-based methods, even if multiple cameras and projectors are used. All the results are computed within a minute using Matlab.

We have also computed noise statistics for our depth estimation. Specifically, we chose 20 planar patches from the experimented scenes that are located at different depth within the working volume. We then computed the mean and standard deviation of the

[6] The glare is produced by a number of dust particles on the half-mirror. Since the half-mirror is very close to the projector, these particles are lit with roughly 100 times the intensity compared to the scene, thereby producing bright spots in the image that have the shape of the camera aperture.

[7] Zhang [2000] takes several images of a board at different orientations to estimate a camera's intrinsic parameters. In our work, we approximately estimate the projector's intrinsics from the frustum specification in its manual. Therefore, we only take one image of the board to compute its pose.

Exhibit 1010 page 39 of 78
DENTAL IMAGING

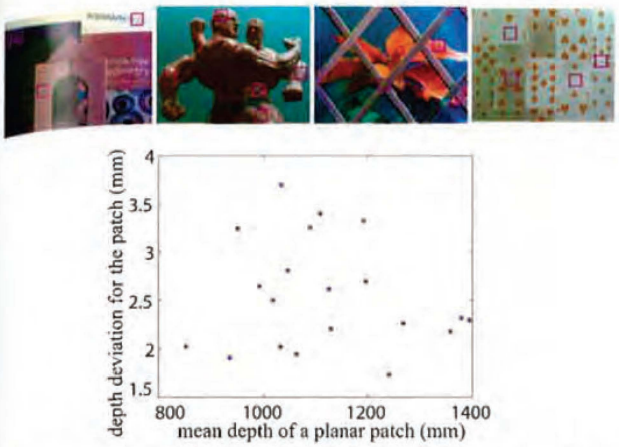
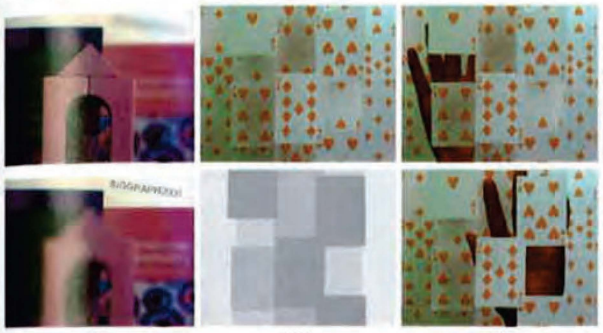


Figure 5: Noise statistics in the depth estimation. Twenty planar patches are chosen from the experimented scenes, shown on the top, and the mean and standard deviation of the depth for those patches are plotted at the bottom. Within a 600mm working volume, our depth estimation noise is around 4mm.



(a) (b) (c)

Figure 6: (a) Refocused images of the book scene in Figure 4. The defocus effects are synthesized using the recovered depth map. (b) A collection of playing cards placed at four depth layers and their estimated depth map. (c) A moving hand from another video is inserted into the card scene with all the desired occlusion effects. **(Please see the companion video.)**

depths inside these patches, as shown in Figure 5. From this figure, we can see that our depth estimation noise is about 4mm for a working volume of 600mm.

3.6 Applications: Refocusing & Video Composition

A distinctive feature of our depth recovery method is that it estimates depth at every pixel in the camera's view and the estimation is reliable near depth discontinuities. This makes our approach particularly well-suited for a variety of image-based applications in computer graphics. Here, we show two examples, namely, refocusing and object insertion to compose a new image or video.

Figure 6(a) shows two refocused images of the book scene in Figure 4: the top one is focused on the foreground wooden blocks and the bottom one is focused on the background letters. These two images were generated with the "lens blur" tool in Adobe PhotoShop.[8] This tool takes an image and its depth map as inputs to generate new images that are focused at any desired depth.

[8] In theory, without having the full light field available, refocusing can not be simulated exactly. The PhotoShop implementation of lens blur is proprietary. We think a practical way to implement it would be to heuristically impaint the occluded layers a little bit to obtain the full light field.

Figure 6(b) shows a collection of playing cards placed at four depth layers. The repetitive textures on the cards and the complex occlusions make it very hard to segment the scene into layers automatically, or even manually. Using projection defocus, our method produces a clean depth map that segments the scene into layers. This segmentation makes it easy to insert objects into the scene with correct occlusion effects. Figure 6(c) shows a real hand from another video inserted into the scene with all the desired occlusion effects. **Please see the video for more results.**

We believe that when a complete depth map is available for an image, a variety of image editing applications become possible or easier to implement. Examples other than the ones we have shown are creating layered representations of complex scenes for view morphing, segmentation for matting, object replacement, and shadow removal. In the past, such applications have not been easy to implement as "clean" depth maps have not been easy to obtain. For most image editing tasks, the depth maps need not be highly precise (as with time-of-flight sensors), but they need to be complete and they need to be reliable at discontinuities, as humans are particularly sensitive to edge artifacts. We hope our depth recovery method, although not as accurate as laser scanners or structured light sensors, will inspire new types of image editing applications.

4 Focused Projection at Multiple Depths

Due to their optical design, standard projectors can only be focused on a single fronto-parallel plane. In some applications, it is desirable to project images onto non-planar structures; for example, multiple planes or a dome to create a virtual environment [Raskar et al. 1998]. In such cases, most parts of the image are blurred due to defocus. One way to solve this problem is to design sophisticated optics. Even if this is possible, the optics cannot be modified to accommodate changes in the structure that is being projected onto. Another approach is to use multiple projectors [Bimber and Emmerling 2006], where the number of depths for which the projected image can be in focus equals the number of projectors. In this section, complementary to the multi-projector approach, we propose a computational method that processes an input image in a scene-dependent way to minimize the defocus blur at all points in the projected output image. Our method requires just a camera and can be applied to any off-the-shelf projector.

4.1 Defocus Compensation Algorithm

Consider the scenario where an image I needs to be projected onto a surface with given depth variation. The radiance of a point on the surface due to illumination by the projector is governed by the projection defocus equation Eq. (1). As in previous work, e.g. [Raskar et al. 2003; Fujii et al. 2005], we use a camera as a proxy for the human eye and seek to make the scene radiance captured by the camera be the same as the input image I. To do so, we can project a *compensation image*, P^*, by solving the projection defocus equation as

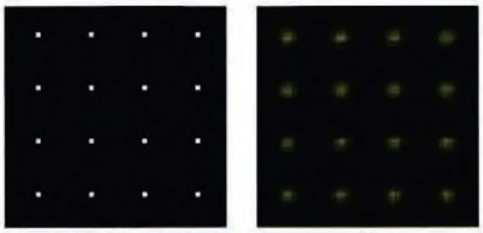
$$P^* = (\alpha f)^{-1} * (I - \beta), \tag{4}$$

where $(\alpha f)^{-1}$ is the inverse of the kernel αf. Theoretically, if P^* is projected, the scene brightness will be the same as the input image I. However, f is a low-pass filter and its inverse will have strong ringing effects. Therefore, Eq. (4) will not be feasible to implement unless the projector has an infinite dynamic range. Instead, we cast the problem of computing the compensation image as a constrained minimization problem, as follows:

$$P^* = \arg\min_P \{d(\alpha f * P + \beta, I) \mid \forall \mathbf{x},\ 0 \le P(\mathbf{x}) \le 255\}, \tag{5}$$

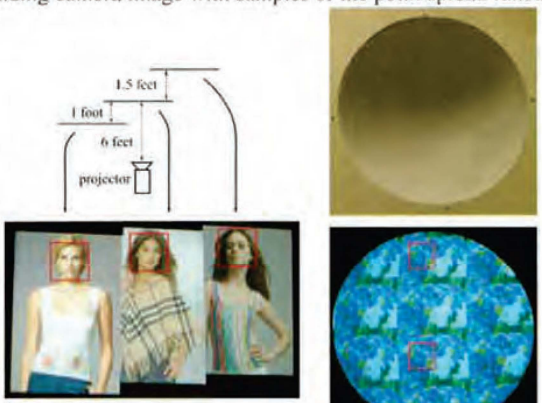
where $\mathbf{x}$ is the projector pixel coordinate and $d(\cdot,\cdot)$ is an image distance metric. Eq. (5) finds the compensation image P^*, with all its brightness values within the projector's dynamic range, that after defocus blurring most closely matches the input image I.

Exhibit 1010 page 40 of 78
DENTAL IMAGING



(a) (b)

Figure 7: Measuring the spatially-varying defocus kernel. (a) A portion of the dot pattern projected onto the scene. (b) The corresponding camera image with samples of the point spread function.



(a) (b)

Figure 8: Two test scenes for focused projection at multiple depths. (a) Projection onto three planes at different depths from the projector. (b) Projection onto a hemispherical dome.

In our implementation, we have used the sum-of-squared pixel differences for the image distance metric $d(\cdot,\cdot)$. The compensation image P^* is found by applying an iterative, constrained, steepest-descent algorithm [Nocedal and Wright 1999]. We represent the defocus convolution, $\alpha f * P$, as a matrix multiplication, FP, where each row of F is the defocus kernel of the corresponding pixel modulated by its albedo. The algorithm starts with $P_0 = I$ and iterates the following two steps:

$$\tilde{P}_{i+1} = P_i + \eta_i G_i \quad (6)$$

$$P_{i+1} = \texttt{CLAMP}(\tilde{P}_{i+1}; 0, 255) \quad (7)$$

where $G_i = F^{\mathsf{T}}(I - \beta - FP_i)$, $\eta_i = \frac{\|G_i\|^2}{\|FG_i\|^2}$, and CLAMP is a pixel-wise clamping operation. Notice that G_i is the gradient of the image distance $\|FP + \beta - I\|^2$ with respect to P. Evaluating G_i is straightforward – it involves two image filterings with the kernel matrices F and F^{T}, respectively. Note that these filterings are spatially-varying and scene-dependent, unlike the standard sharpening operations that are often built into projectors by their manufacturers. Without Eq. (7), iterating Eq. (6) alone is a standard steepest-descent algorithm, which converges to the solution of Eq. (4). Combining Eq. (7) and Eq. (6) minimizes the difference between the defocused compensation image and the original input image within the dynamic range of the projector.

4.2 Kernel Estimation

We now describe how we obtain β and αf for any given scene. We obtain the ambient term β by turning off the projector and taking an image. To obtain αf for each pixel, we could shift a dot pattern across the scene and the temporal radiance profile for each pixel is then its kernel. This method is effective for a scene with intricate geometry, such as the scenes shown in Section 3. However, the scenes used in our experiments here are at least piece-wise smooth. Therefore, we simply project a sparse binary dot pattern, like the one in Figure 7(a). In our experiments, the distance between neighboring dots is 12 pixels. The point spread patterns captured by the camera, shown in Figure 7(b), are approximately the kernels for the projector pixels that are '1'. From these kernels, we interpolate the kernels for other pixels that are '0' in a bilinear manner. In the end, we have a per-pixel kernel map that accounts for spatially-varying defocus effects. We repeat the same procedure to capture kernels for each of the three color channels of the projector. Calibrating such a kernel map also helps to compensate lens aberrations as well. For the above computations of β and αf, we always warp the image from the camera to the projector's coordinate frame, which makes it convenient to compare the defocused compensation image, $FP + \beta$, and the input image I. The correspondences between the camera's and the projector's pixels are determined by shifting sinusoidal patterns in both horizontal and vertical directions, as in Section 3.

4.3 Focused Projection on Multiple Planes

In our first experiment, our goal was to project an image of three fashion models onto three planes that are at different depths, as shown in Figure 8(a). The projector is focused on the middle plane. We show the focus compensation results for the three face regions in the left three columns of Figure 9. The first row of Figure 9 shows the original image regions for the three faces. The third row shows the defocused original image regions, captured by a camera, without any compensation. This row represents the best the projector (without compensation) can produce. Notice that even though the middle plane is in focus, the face still looks a little blurred for this plane due to projector artifacts.[9] As expected, the faces on the left and right planes are even more blurred due to defocus. The second row shows the compensation image regions resulting from Eqs. (6,7), which represent the new input to the projector. These image regions look like high-pass filtered versions of the original input ones. This is expected as our method boosts high-frequency components to compensate for the defocus induced by scene geometry. The fourth row shows the defocused compensation image regions captured by the same camera, which represent our results. Notice that these image regions are less blurry and more closely resemble the input image regions. Among these three defocused compensation image regions, the one on the middle plane (second column) is almost identical to the input. The ones on the left and right planes (first and third columns) are slightly blurred with respect to their inputs because defocus is more severe in these two cases, and certain high frequencies in the input image that are cut by the defocus cannot be fully compensated for due to the limited dynamic range of the projector.

4.4 Focused Projection on a Dome

In our second experiment, we project an image of a tiled flower texture onto the inside of a hemispherical dome, as shown in Figure 8(b). The projector is focused on the front plane of the dome. In the right two columns of Figure 9, we show the focus compensation results for two patches on the dome – one is from the top of the dome and the other is from the center. The top patch is less blurred than the center one because the former is closer to the focal plane of the projector. As with the three-plane experiment, the compensation results closely resemble the input images.

5 Depixelation

In this section, we show how the defocus compensation method discussed in Section 4 can be used to reduce projector pixelation.

[9] We found that even if we focus a projector on a fronto-parallel plane, when projecting a binary dot pattern, we see "light leakage" from one pixel into its neighbors. We observed this effect on different projector models.

Exhibit 1010 page 41 of 78
DENTAL IMAGING

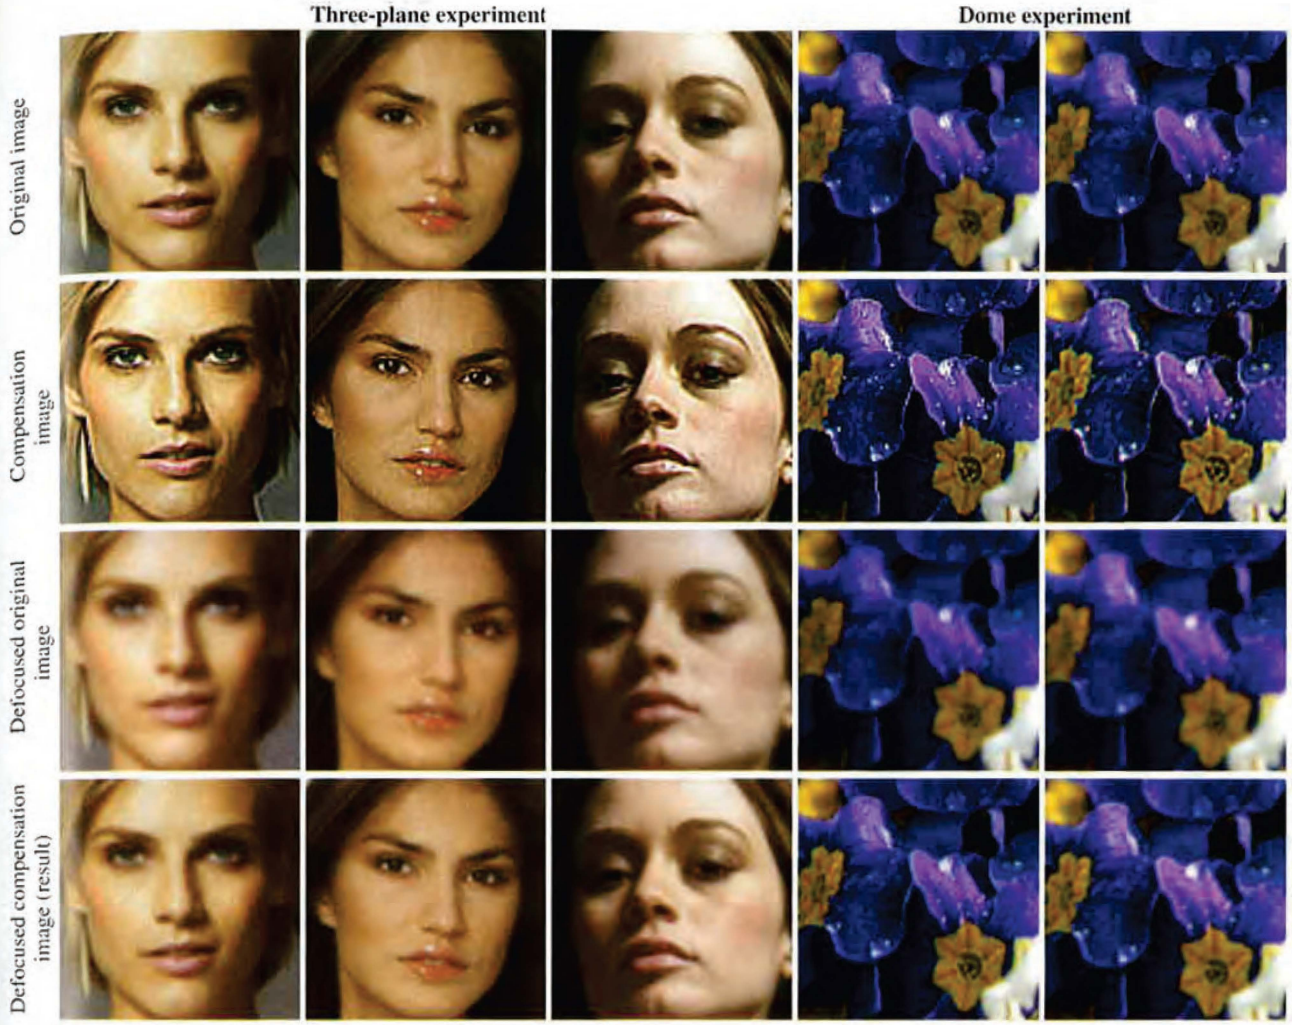


Figure 9: Results for defocus compensation. The three faces on the left were projected onto the three planar surfaces in Figure 8(a). The two flower textures on the right were projected onto the hemispherical dome in Figure 8(b). The defocus compensation method results in less blurry image regions (fourth row) than the uncompensated image regions (third row). The original image regions are shown in the first row and the compensation image regions computed by the method and used as input to the projector are shown in the second row. The compensation method is scene independent as it can handle a spatially-varying defocus kernel. **(Please see the companion video.)**

Pixelation is a clearly noticeable artifact produced by all digital projectors. It is caused by two factors. The first is the spatial digitization due to the finite resolution of the projector. The second is the gap (dead-zone) between adjacent pixels on the projector's physical image plane that arises because the pixel fill-factor is never 100 %. The digitization is known to create jaggy boundaries when the resolution of the projector is not high enough for the given application. This effect is observed particularly when high quality captured images are projected, as the resolution of most LCD and DLP projectors has remained at 1024x768 or less for the last seven years, while digital cameras have tripled in their resolution during the same time period. The dead-zone between pixels does not generate light and produces thin black lines on the projection screen, known as the *screen-door effect*. This effect makes pixelation more pronounced as it clearly marks out pixel boundaries on the screen.

In our implementation of depixelation, we assume that the images are projected onto a single fronto-parallel screen, i.e., the traditional projection scenario. Our basic approach is to focus the projector slightly in front of (or behind) the projection screen so that the image on the screen is slightly blurred. In this case, a slight amount of light is leaked into the black gaps and neighboring pixels, which reduces the screen-door effect as well as smoothes out the jaggy pixel boundaries. However, doing so also creates blurry images. Since this purposely induced blur is very slight, we can use the compensation method in Section 4 to process the input image so that the defocused projection looks very similar to the original input image. Therefore, the compensation method can be used in conjunction with the minor induced defocus to achieve depixelation.

Specifically, we use the compensation method to compute an optimal image P^* at the projector resolution whose defocused projection most closely resembles the input image, which may have higher resolution. Note that if the input image resolution is higher than the projector's native resolution, the number of rows in F is larger than the number of columns. Each row of F is a defocus kernel modulated by both surface albedo and the screen-door effect.

We have tested the depixelation method on several high resolution input images (each one with 3 times the projector resolution). The capture of β and F is done using the method in Section 4, except that we capture a higher resolution kernel map by zooming-in our camera so that 3×3 camera pixels see about one projector pixel.

Exhibit 1010 page 42 of 78
DENTAL IMAGING

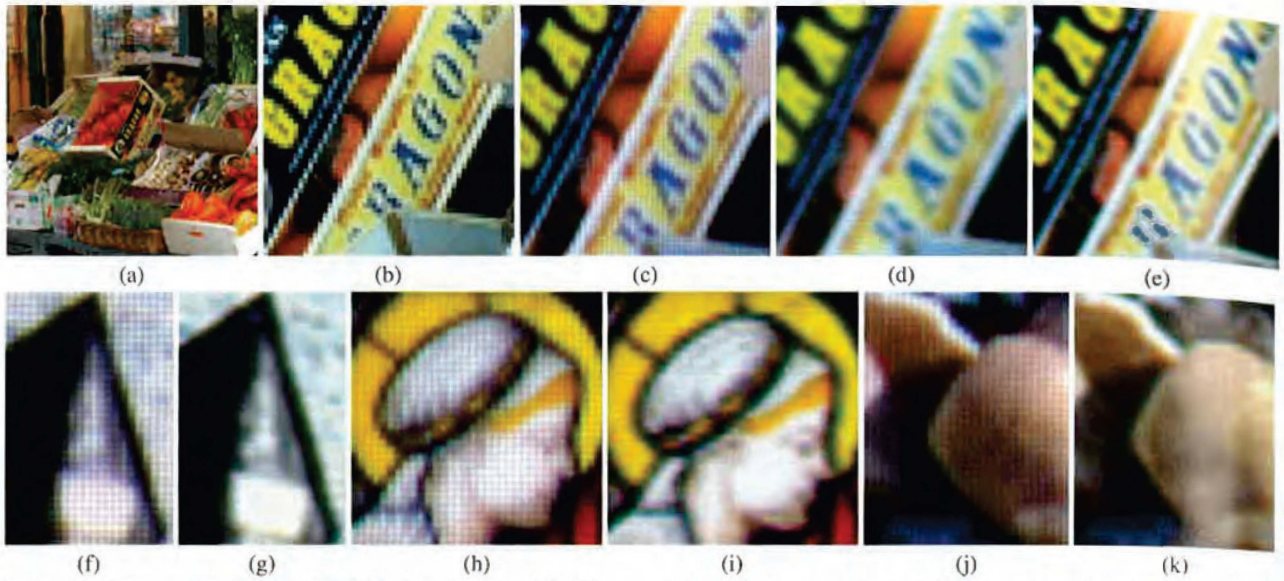


Figure 10: Examples of Projector Depixelation. (a) An original image of 3× projector resolution. (b) A patch in (a) of projector resolution. (c,d) Projection of (b) under focused and defocused settings, respectively. (e) Our depixelation result. (f-k) Three more examples of depixelation with (f,h,j) as focused projection and (g,i,k) as depixelated projection. In all examples, the compensation greatly reduces pixelation effects while preserving image quality. **(Please see the companion video.)**

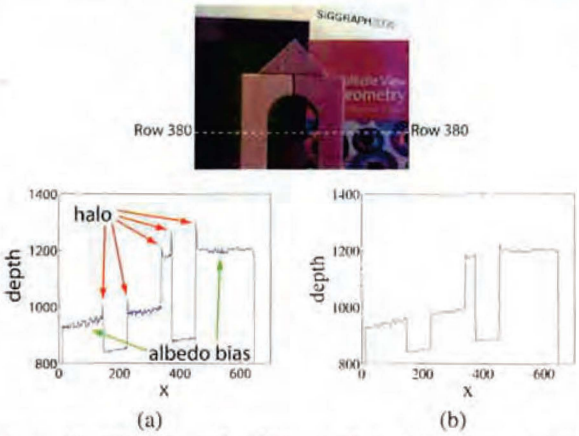


Figure 11: Albedo-correlated bias and halo effects in depth estimation. (a) The depth profile for row 380 of the scene. (b) The depth profile after removing noise and halo by a 7×7 median filter.

Figure 10 shows four examples of depixelation. In all the cases, spatial digitization and screen-door effects are greatly reduced and yet image sharpness is preserved. Our current implementation of the compensation algorithm uses Matlab. The algorithm takes about 10 iterations to converge and takes 3-5 minutes to produce the final compensation image. Since the compensation method only involves filtering with F and F^T, we believe it can be implemented on image processing chips or graphics hardware to achieve real-time (frame-rate) performance.

6 Discussion

In this paper, we have modelled projection defocus as a linear system and developed a temporal defocus method that recovers depth at all pixels in an image, regardless of the scene complexity. We also proposed a scene-dependent filtering algorithm that increases the depth of field of a projector without modifying its optics. Finally, we showed that defocus can be exploited to reduce pixelation artifacts produced by digital projectors. We now discuss the limitations of our work and open problems as topics for future research.

First, our temporal defocus method should, in principle, be invariant to surface albedo. In practice, we found that the surface albedo sometimes does affect the depth estimation, as shown in Figure 11. We believe that this is only an issue with our current implementation. Our beam-splitter not only reflects light from the scene to the camera, but also transmits light from a black backdrop placed behind it. Therefore, the defocus kernel f in Eq. (1) is actually a sum of the defocus kernels of both the scene and the backdrop. If the scene is much brighter than the black backdrop, Eq. (3) is dominated by the scene's defocus kernel. However, when the scene point is dark, due to low albedo or because its surface orientation is near-grazing for the projector, the light from the scene point to the camera is weak and comparable to the light from the backdrop. In this case, Eq. (3) will depend also on the defocus kernel of the backdrop, which causes errors in depth estimation. The backdrop kernel remains constant and can be calibrated to reduce depth estimation errors. Alternatively, the beam-splitter can be placed in a custom-designed chamber that absorbs all forms of stray light.

The projector defocus is for most cases independent of scene geometry, and this independence makes is superior to camera defocus for depth estimation. However, there is a special case that this independence does not hold. Specifically, when a closer object occludes a distant object, there will be a narrow band of points near the occluding contour on the distant object that "see" only part of the projector's aperture; so the projected pattern will be more infocus than it would be if the entire aperture were visible. This partial aperture effect results in a halo artifact: depth estimation biased toward the focal plane, as shown in Figure 11. However, even in this special case, projector defocus is still better than camera defocus. Camera defocus will cause estimation errors for both the foreground and background pixels near the occluding boundary. In contrast, projector defocus will only cause errors at the background pixels, while leaving the foreground pixels unaffected. Furthermore, in practice, because the solid angle subtended by the projector aperture is only about 1 degree with respect to a scene point, the number of affected background pixels is small. Therefore, as shown in Figure 11, the halo effect can be greatly reduced by a simple median filtering.

Exhibit 1010 page 43 of 78
DENTAL IMAGING

In practice, using projector defocus for depth estimation does not require moving detectors or changing optical settings as traditional camera defocus methods do [Pentland 1987; Nayar and Nakagawa 1994]. Therefore, our method can be more conveniently implemented. However, a coaxial implementation of depth from camera defocus, as in [Nayar et al. 1996], allows for capturing dynamic scenes, which our method currently is incapable of, because we use 24 images for depth estimation. We believe this number can be greatly reduced, which would enable us to develop a real-time sensor that can handle dynamic scenes. In Eq. (1), there are three unknowns: α, β, and z. If we can calibrate the defocus kernel f a priori, then, in theory, 3 images are enough to solve for these unknowns. Interestingly, this minimum number is the same for other shape acquisition methods, like the phase-shift method [Huang et al. 2003], the time-of-flight method [Gonzales-Banos and Davis 2004], and photometric stereo [Horn and Brooks 1989].

Our temporal defocus method can deal with specular surfaces, whenever the BRDF's are approximately constant within the solid angle subtended by the projector aperture from the scene point. Nevertheless, the extreme case of near-mirror reflection, as well as saturated highlight pixels, will cause errors. Although our method is developed under the direct illumination assumption, it's not very sensitive to the indirect illumination. This is because indirect illumination usually creates only very low frequency intensity fluctuation, which will mainly affect Fourier coefficient A_0. Our depth is however estimated via A_1 and A_2. Of course, for extremely specular reflection, the global illumination can change dramatically with the projected patterns and hence our method will break down.

One challenging open problem for defocus-based methods is handling translucent objects. A translucent material introduces additional blur due to subsurface scattering, which would cause a systematic bias in our methods. The amount of bias depends on the amount of translucency blur compared to defocus blur. If the defocus blur dominates the translucency blur, the error will be small. Otherwise, the error will be large. It would be interesting to extend our analysis to handle blur induced by translucency. In practice, the depth of mixture pixels at boundaries is often a blend of its neighboring depth values. It would be interesting to accurately model the mixture pixels, e.g., combining our approach with the defocus matting method [McGuire et al. 2005], to estimate mattes for scenes with continuous depth variations.

Our method for defocus compensation essentially extends the depth of field of the projector. However, when the blur is very severe, our method cannot fully compensate for the blur due to the limited dynamic range of the projector. In such a case, combining our method with multiple-projector solutions [Bimber and Emmerling 2006] could be beneficial. Finally, we plan to extend our work on depixelation to make it a practical and useful projector feature. How objectionable pixelation artifacts are, is really a subjective issue. For our technique to be incorporated into commercial projectors, a few preset options can be provided so that the user can decide how much focus quality they are willing to tradeoff for depixelation.

References

BIMBER, O., AND EMMERLING, A. 2006. Multi-focal projection. *IEEE Trans. on Visualization and Computer Graphics to appear*.

BIMBER, O., WETZSTEIN, G., EMMERLING, A., AND NITSCHKE, C. 2005. Enabling view-dependent stereoscopic projection in real environments. In *Proc. Int. Symp. on Mixed and Augmented Reality*, 14–23.

CURLESS, B., AND LEVOY, M. 1995. Better optical triangulation through spacetime analysis. In *Proc. Int. Conf. on Computer Vision*, 987–994.

DAVIS, J., NEHAB, D., RAMAMOOTHI, R., AND RUSINKIEWICZ, S. 2005. Spacetime stereo : A unifying framework for depth from triangulation. *IEEE Trans. on Pattern Analysis and Machine Intelligence 27*, 2, 296–302.

FAUGERAS, O. 1993. *Three-Dimensional Computer Vision*. MIT Press.

FAVARO, P., AND SOATTO, S. 2005. A geometric approach to shape from defocus. *IEEE Trans. on Pattern Analysis and Machine Intelligence* (in press).

FUJII, K., GROSSBERG, M., AND NAYAR, S. 2005. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 814–821.

GIROD, B., AND SCHEROCK, S. 1989. Depth from defocus of structured light. In *Proc. SPIE Conf. on Optics, Illumination, and Image Sensing for Machine Vision*.

GONZALES-BANOS, H., AND DAVIS, J. 2004. A method for computing depth under ambient illumination using multi-shuttered light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 234–241.

GROSSBERG, M., PERI, H., NAYAR, S., AND BELHUMEUR, P. 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. I, 452–459.

HORN, B., AND BROOKS, M. 1989. *Shape from Shading*. MIT Press.

HUANG, P. S., ZHANG, C. P., AND CHIANG, F. P. 2003. High speed 3-d shape measurement based on digital fringe projection. *Optical Engineering 42*, 1, 163–168.

JAIN, A. K. 1989. *Fundamentals of Digital Image Processing*. Prentice Hall.

JIN, H., AND FAVARO, P. 2002. A variational approach to shape from defocus. In *Proc. Eur. Conf. on Computer Vision*, 18–30.

KANADE, T., GRUSS, A., AND CARLEY, L. 1991. A very fast vlsi rangefinder. In *Proc. Int. Conf. on Robotics and Automation*, vol. 39, 1322–1329.

KONINCKX, T. P., PEERS, P., DUTR, P., AND GOOL, L. V. 2005. Scene-adapted structured light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 611–619.

LEVOY, M., CHEN, B., VAISH, V., HOROWITZ, M., MCDOWALL, I., AND BOLAS, M. 2004. Synthetic aperture confocal imaging. In *SIGGRAPH Conference Proceedings*, 825–834.

LJUNG, L. 1998. *System Identification: A Theory for the User*, 2 ed. Prentice Hall.

MAJUMDER, A., AND WELCH, G. 2001. Computer graphics optique: Optical superposition of projected computer graphics. In *Proc. Eurographics Workshop on Virtual Enviroment/Immersive Projection Technology*.

MCGUIRE, M., MATUSIK, W., PFISTER, H., HUGHES, J. F., AND DURAND, F. 2005. Defocus video matting. In *SIGGRAPH Conference Proceedings*, 567–576.

NAYAR, S. K., AND NAKAGAWA, Y. 1994. Shape from focus. *IEEE Trans. on Pattern Analysis and Machine Intelligence 16*, 8, 824–831.

NAYAR, S. K., WATANABE, M., AND NOGUCHI, M. 1996. Real-time focus range sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence 18*, 12, 1186–1198.

NOCEDAL, J., AND WRIGHT, S. J. 1999. *Numerical Optimization*. Springer.

OPPENHEIM, A. V., AND WILLSKY, A. S. 1997. *Signals and Systems*, 2 ed. Prentice Hall.

PENTLAND, A. 1987. A new sense for depth of field. *IEEE Trans. on Pattern Analysis and Machine Intelligence 9*, 4, 523–531.

RAJ, A., AND ZABIH, R. 2005. A graph cut algorithm for generalized image deconvolution. In *Proc. Int. Conf. on Computer Vision*.

RAJAGOPALAN, A. N., AND CHAUDHURI, S. 1997. A variational approach to recovering depth from defocused images. *IEEE Trans. on Pattern Analysis and Machine Intelligence 19*, 10, 1158–1164.

RASKAR, R., WELCH, G., CUTTS, M., LAKE, A., STESIN, L., AND FUCHS, H. 1998. The office of the future : A unified approach to image-based modeling and spatially immersive displays. In *SIGGRAPH Conference Proceedings*, 179–188.

RASKAR, R., WELCH, G., LOW, K., AND BANDYOPADHYAY, D. 2001. Shader lamps. In *Proc. Eurographics Workshop on Rendering*.

RASKAR, R., VAN BAAR, J., BEARDSLEY, P., WILLWACHER, T., RAO, S., AND FORLINES, C. 2003. ilamps: geometrically aware and self-configuring projectors. In *SIGGRAPH Conference Proceedings*, 809–818.

RASKAR, R., HAN TAN, K., FERIS, R., YU, J., AND TURK, M. 2004. Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. In *SIGGRAPH Conference Proceedings*, 679–688.

SCHARSTEIN, D., AND SZELISKI, R. 2003. High-accuracy stereo depth maps using structured light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 195–202.

SCHECHNER, Y. Y., KIRYATI, N., AND BASRI, R. 2000. Separation of transparent layers using focus. *Int. J. on Computer Vision 39*, 1, 25–39.

SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. In *SIGGRAPH Conference Proceedings*, 745–755.

TAPPEN, M. F., RUSSELL, B. C., AND FREEMAN, W. T. 2004. Efficient graphical models for processing images. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 2, 673–680.

ZHANG, L., SNAVELY, N., CURLESS, B., AND SEITZ, S. M. 2004. Spacetime faces: High-resolution capture for modeling and animation. In *ACM Annual Conference on Computer Graphics*, 548–558.

ZHANG, Z. 2000. A flexible new technique for camera calibration. *IEEE Trans. on Pattern Analysis and Machine Intelligence 22*, 11, 1330–1334.

Exhibit 1010 page 44 of 78
DENTAL IMAGING

APPENDIX A-2

Projection Defocus Analysis for Scene Capture and Image Display *

Li Zhang Shree Nayar
Columbia University

Abstract

In order to produce bright images, projectors have large apertures and hence narrow depths of field. In this paper, we present methods for robust scene capture and enhanced image display based on projection defocus analysis. We model a projector's defocus using a linear system. This model is used to develop a novel temporal defocus analysis method to recover depth at each camera pixel by estimating the parameters of its projection defocus kernel in frequency domain. Compared to most depth recovery methods, our approach is more accurate near depth discontinuities. Furthermore, by using a coaxial projector-camera system, we ensure that depth is computed at all camera pixels, without any missing parts. We show that the recovered scene geometry can be used for refocus synthesis and for depth-based image composition. Using the same projector defocus model and estimation technique, we also propose a defocus compensation method that filters a projection image in a spatially-varying, depth-dependent manner to minimize its defocus blur after it is projected onto the scene. This method effectively increases the depth of field of a projector without modifying its optics. Finally, we present an algorithm that exploits projector defocus to reduce the strong pixelation artifacts produced by digital projectors, while preserving the quality of the projected image. We have experimentally verified each of our methods using real scenes.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning, display algorithm; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Depth cues, range data, shape.

Keywords: projector defocus, temporal defocus analysis, depth recovery, multi-focal projection, projector depixelation, refocus synthesis, image composition.

1 Introduction

Digital projection technologies, such as Digital Light Processing (DLP) and Liquid Crystal Displays (LCD), are increasingly used in consumer, commercial and scientific applications. In computer graphics and vision, video projectors have recently been used as per-pixel controllable light sources for real-time shape acquisition [Huang et al. 2003; Zhang et al. 2004; Davis et al. 2005; Koninckx et al. 2005], for complex appearance capture [Levoy et al. 2004; Sen et al. 2005] and control [Raskar et al. 2001; Grossberg et al. 2004; Bimber et al. 2005]. All these applications require the projectors to be focused for best performance. In practice, projectors are built with large apertures to maximize their brightness. As a result, virtually all projectors have very narrow depths of field; they are designed to produce focused images on a single fronto-parallel screen. An analysis of the defocus properties of projectors is therefore beneficial as it could lead to new methods that take advantage of, as well as compensate for, projection defocus.

In this paper, we provide the first systematic analysis of projector defocus and demonstrate its applications to robust scene capture and enhanced image display. We first present a simple linear model for projector defocus. Based on this model, we present a frequency-domain method for estimating the spatially-varying defocus kernel of a projector. The kernel estimated at each scene point is used to recover the 3D geometry of the scene. Based on the estimated kernel, we also present a technique that computationally manipulates an input image to minimize its defocus blur when it is projected onto a non-planar scene. Finally, we demonstrate that defocusing can be put to good use. A slight amount of defocusing, in conjunction with a compensation algorithm, can be used to reduce pixelation, a strong artifact produced by all digital projectors. Specifically, this paper makes the following three contributions:

Scene Geometry using Temporal Defocus Analysis: We propose a method called temporal defocus analysis that estimates depth at each camera pixel, independently, without using information from neighboring pixels (Section 3). Our method, compared to existing methods, such as stereo or even traditional depth from focus/defocus algorithms, is more accurate near depth discontinuities. Since the method is not based on triangulation, we can use a coaxial projector-camera system and compute depth at all camera pixels, without any missing parts. These advantages make the method uniquely suited to computer graphics applications like refocus synthesis and image composition, both of which we demonstrate.

Focused Projection at Multiple Depths: We present an iterative, spatially-varying filtering algorithm that compensates for defocus blur based on scene geometry (Section 4). This technique effectively increases the depth of field of a projector without modifying its optics. As a result, we are able to use a single projector to project well-focused images on multiple planes that are at different depths as well as on curved projection surfaces. We believe this capability addresses an important limitation of current projectors and widens their applicability in the real world.

Depixelation by Defocusing: Finally, we generalize our defocus compensation algorithm to reduce the strong pixelation artifacts produced by all digital projectors (Section 5). The key idea is to slightly defocus the projector so as to attenuate the high frequencies produced by pixelation and use the defocus compensation algorithm to make up for the induced projector defocus. This method is especially suited for projecting high resolution images using a low resolution projector, and we show several examples of results to illustrate its effects.

2 Previous Work

Many methods have been proposed in computer vision to recover 3D shape from images, including multi-view triangulation methods, single-view photometric methods, and camera focus/defocus methods. Triangulation-based methods [Faugeras 1993], e.g., structure from motion and stereo, require a point to be visible in at least two views to be reconstructed. Scenes with complex occlusions remain a challenging problem for these methods. Photometric methods [Horn and Brooks 1989], e.g., photometric stereo and shape from shading, estimate surface normals instead of surface depth. Converting normals to depths is an ill-posed problem for scenes with depth discontinuities.

*This work was conducted at the Computer Vision Laboratory at Columbia University. It was supported by an ITR grant from the National Science Foundation (No. IIS-00-85864). The authors thank Gurunandan Krishnan for his help with the experiments and the video and Anne Fleming for her help with the audio recording and the proofreading. The authors also thank www.dpreview.com for permitting us using their images as some of the examples shown in the paper.

Copyright © 2006 by the Association for Computing Machinery, Inc.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.
© 2006 ACM 0730-0301/06/0700-0907 $5.00

Exhibit 1010 page 46 of 78
DENTAL IMAGING

Depth recovery methods based on camera focus and defocus, e.g. [Pentland 1987; Nayar and Nakagawa 1994; Nayar et al. 1996; Schechner et al. 2000; Favaro and Soatto 2005], have the potential to recover depth at every pixel, regardless of the scene complexity and occlusions. To resolve the focus ambiguity of textureless surfaces, patterns can be projected to force scene texture [Girod and Scherock 1989]. However, camera defocus kernels depend on local surface geometry. To simplify the kernel analysis, most previous works assume that, within a small spatial window, the surface depth is constant, the so-called *equalfocal assumption*. This assumption smears shape details and is invalid across depth discontinuities. To alleviate this problem, Jin et al. [2002] and Rajagopalan and Chaudhuri [1997] estimate depths for all pixels simultaneously via a large scale energy minimization, which is computationally expensive and prone to local minima. Our key observation is that, unlike camera defocus, the kernel for projector defocus is *scene independent*, for most scene surfaces. Specifically, when a 3D scene point sees the entire projector aperture, its defocus kernel depends only on its distance to the projector lens and not on its neighboring surface geometry. This difference arises from the fact that projector defocus convolution happens on the projector's image plane while camera defocus convolution happens on the scene surface. Exploiting this scene-independent property, we project a shifting pattern over the scene and compute depth at each pixel using just its intensity variation over time. Without using the equalfocal assumption, our method works well at depth discontinuities. Furthermore, it is simple and not subject to local minima.

Our temporal per-pixel defocus analysis is inspired by previous structured light range finding methods. In particular, Kanade et al. [1991] and Curless and Levoy [1995] use temporal intensity variation to resolve correspondences between camera pixels and a sweeping laser stripe. Huang et al. [2003] have developed a real-time range finder by sweeping periodic sinusoidal stripes using a DLP projector. Zhang et al. [2004] and Davis et al. [2005] have generalized these ideas to space-time stereo. All these range finding techniques are based on the principle of triangulation, which cannot estimate depth for points that are visible to only the camera or the source but not both. Since our method is not based on triangulation, we are able to use a coaxial configuration where the camera and projector share the same optical center, and compute depth at all camera pixels (no missing parts). To our knowledge, the only other technique that is able to estimate a complete depth map with clean discontinuity boundaries is time-of-flight, e.g., [Gonzales-Banos and Davis 2004], which requires expensive specialized hardware. Raskar et al. [2004] proposed a method to detect depth discontinuities using multiple flash lights. The detected discontinuities can be used to enhance stereo matching algorithms. Compared to this work, our method directly generates reliable depth estimation for all pixels, including the ones near discontinuities.

In addition to scene capture, several methods have been proposed that use cameras to enable projectors to display "better" images. These include compensating distortions due to surface geometry [Raskar et al. 2003] and correcting brightness variations due to surface color and texture [Grossberg et al. 2004; Bimber et al. 2005]. In these works, the projectors are often used to display images onto scenes with considerable depth variation, where defocus blur is inevitable and most often spatially-varying. Complementary to these previous methods, our work seeks to compensate for defocus blur, and hence could be beneficial to the previous methods. In [Bimber and Emmerling 2006], a system is proposed that can project focused images on multiple planes at different depths. This system uses multiple projectors where each projector is focused on a single plane. In our case, we use a single projector to simultaneously project well-focused images at multiple depths. Another interesting related work is by Majumder and Greg [2001], in which multiple projectors at different focal settings are combined to generate depth of field effects at an interactive rate.

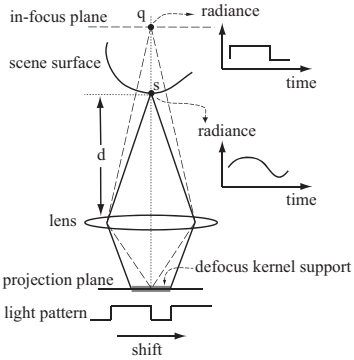
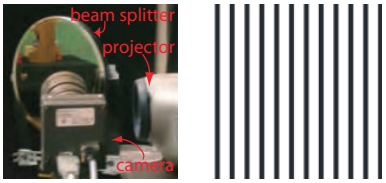


Figure 1: The principle of depth from projection defocus. Points at different distances to the projector lens exhibit different amounts of blur in their temporal radiance profile as a periodic illumination pattern is shifted across the scene.



(a) (b)

Figure 2: (a) A coaxial projector-camera system for depth from projection defocus. The system is made coaxial only to ensure that depth can be computed at all camera pixels. (b) The illumination pattern that is shifted across the scene to measure depth at each pixel, independently.

Deblurring is a well-studied topic in image processing where numerous techniques have been proposed, ranging from classical Wiener filtering and conjugate gradient optimization [Jain 1989] to more recent algorithms based on graph cuts [Raj and Zabih 2005] and belief propagation [Tappen et al. 2004]. Our technique is based on bound-constrained quadratic programming [Nocedal and Wright 1999], which is a variant of existing algorithms that best suits our formulation of the problem. As we show, our deblurring method can also reduce pixelation effects produced by digital projectors.

3 Depth from Projection Defocus

In this section, we formulate the problem of recovering 3D shape from projection defocus and present our solution.

3.1 Temporal Defocus Analysis

Consider a scene that is illuminated by structured light from a projector which is focused behind the scene, as shown in Figure 1. For a point **q** that is in focus, its irradiance comes from a single point on the projector's image plane. For a point **s** that is out of focus, its irradiance equals the convolution of its defocus kernel with the structured light pattern on the projector's image plane. Assuming that the surface is opaque, the radiance I of **s** along any given outgoing direction can be written as

$$I = \alpha f(\mathbf{x};z) * P(\mathbf{x}) + \beta, \quad (1)$$

where $*$ denotes convolution, α is a factor depending on surface reflectance[1], β is the radiance due to the ambient light, $f(\mathbf{x};z)$ is the

[1]To be precise, α takes into account all surface shading factors: BRDF, orientation with respect to the irradiance direction and the squared distance fall-off of the projector brightness.

Exhibit 1010 page 47 of 78
DENTAL IMAGING

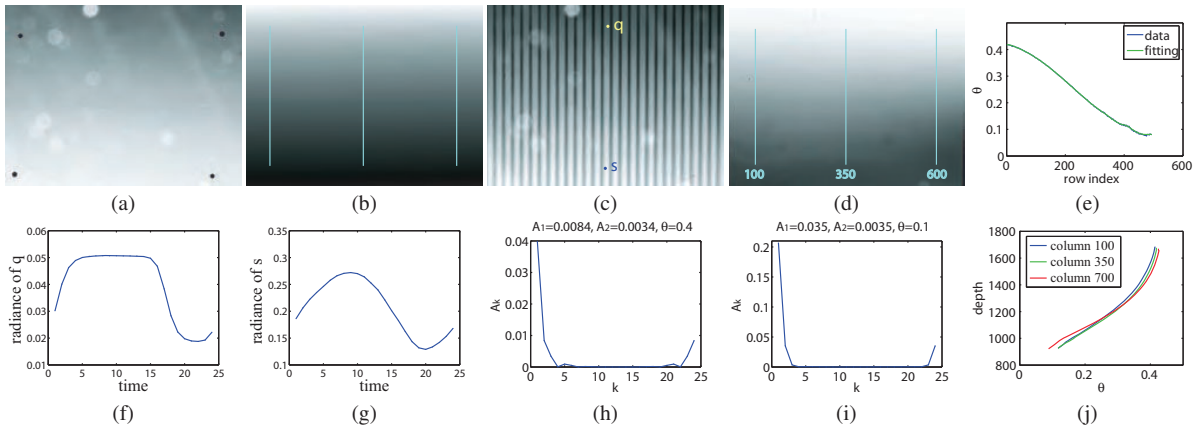


Figure 3: Illustration of the calibration procedure. (a) A white board with four markers tilted in front of the camera. (b) Depth map of the board computed using the markers. (c) The board under one of the stripe projection patterns. (d) θ map for the board. (e) θ values in column 350 of (d). (f,g) Temporal radiance profiles of the points **q** and **s**, respectively. (h,i) Discrete-time Fourier series (DFS) of (f) and (g). (j) Mappings from θ to depth z for columns 100, 350 and 600, respectively.

defocus kernel, and $P(\mathbf{x})$ is the illumination pattern. The defocus kernel f depends on the depth z.[2] We now describe how to recover the depth z by estimating the kernel f from radiance measurements.

3.2 Depth Estimation

To estimate the kernel f, we note that Eq. (1) defines a linear system in which the projection pattern and the scene radiance are the input and the output, respectively. Estimating the kernel of a linear system is a classical problem in system identification [Ljung 1998] and we take a frequency-domain approach to solve it in our setting. Our basic idea is to shift an illumination pattern with a wide range of frequencies within it across the scene. The radiance of a surface point over time is then the response of its defocus kernel to the excitation by the illumination pattern. As the pattern is shifted, points at different distances to the projector exhibit different amounts of blur in their *temporal* radiance profile, as illustrated in Figure 1. We use this temporal blur for depth recovery.

Given the temporal radiance sequence, I_l, $l = 0,\ldots,L-1$, for a point, we quantify its blur by decomposing it into a discrete-time Fourier series (DFS) [Oppenheim and Willsky 1997] as

$$I_l = A_0 + \sum_{k=1}^{L-1} A_k \cos(\omega_k l - \phi_k), \qquad (2)$$

where $\omega_k = \frac{2k\pi}{L}$, $A_k = (B_k^2 + C_k^2)^{\frac{1}{2}}$, $\phi_k = \arctan(B_k, C_k)$, $B_k = \frac{1}{L}\sum_{k=0}^{L-1} I_l \sin(\omega_k l)$ and $C_k = \frac{1}{L}\sum_{k=0}^{L-1} I_l \cos(\omega_k l)$. Since the kernel f is a low-pass filter, how quickly the coefficients A_k diminish with k is a measure of the amount of defocus, which in turn yields depth. Note that A_0 cannot be used to estimate depth because it depends on the ambient light β. Although both A_1 and A_2 are scaled by the albedo α, their ratio can be used to determine how severely the defocus kernel attenuates the second-order harmonic with respect to the first-order one. Therefore, we use the following ratio

$$\theta = \frac{A_2}{A_1} \qquad (3)$$

as a measure of depth. In Eq. (3), $A_1 > A_2 > 0$ and $\theta \in [0,1]$ because f is a low-pass filter. Deriving the analytic mapping between θ and z is tedious but possible if we know precisely the optical design of the specific projector. However, analytically deriving a projector-specific mapping is not worth the effort as it would not apply to other projectors — each projector tends to have a unique optical design. Therefore, in the next section, we present a general data-driven approach that calibrates the $\theta - z$ mapping once and for all, for any given optical setting of any given projector-camera system.

[2] Due to lens aberration, f generally also varies across different pixels.

3.3 System Setup and Illumination Pattern

We have built a prototype camera-projector system to implement our depth from projection defocus method. Our system consists of an NEC LT260K DLP projector and a Basler A311f monochrome camera. To avoid shadows and occlusions, we approximately align the optical centers of the projector and the camera with a beam splitter (Edmund Optics stock #NT39-493), as shown in Figure 2(a)[3]. The projector is always focused on a plane behind the working volume to avoid a two-way defocus ambiguity. As the projector is quite bright[4], we stop-down the aperture of the camera to F11 so that it works approximately as a pinhole camera – any defocus introduced by the camera is negligible compared to that of the projector.

There are many choices of input sequences (excitation signals) in system identification theory. We have chosen to use a simple one – a binary periodic sequence 011011011011... with period 3, which is one type of M-sequence [Ljung 1998]. We encode this sequence as a stripe pattern in which each bit corresponds to an 8-pixel wide stripe, as shown in Figure 2(b). We shift this pattern, one pixel at a time, and take a total of $L = 24$ images for each experiment[5].

3.4 Calibration

We calibrate the mapping from the θ in Eq. (3) to the depth z in three steps. **Step 1**: We compute the correspondence between projector and camera pixels. This is achieved by projecting shifted sinusoids in both horizontal and vertical directions. The details of this procedure can be found in [Scharstein and Szeliski 2003]. **Step 2**:

[3] We align the camera and the projector by projecting an image onto a white board with a fence in front of it and adjusting the camera position until it does not see any shadows cast by the projector.

[4] During depth estimation, we have used the projector as a grayscale one by removing its color wheel from the light path to boost its brightness.

[5] In our experiments, we have found that the projector produces several undesirable effects: it vibrates at a high frequency, possibly due to the rotation of its fan; its brightness is not stable over time; the imprecise synchronization between the camera and projector causes measured radiances to fluctuate slightly over time. To resolve these issues, we repeat the shifting of the pattern about 85 times and compute a mean image for each shift. The camera runs at 60Hz and the total acquisition time is within a minute.

Exhibit 1010 page 48 of 78
DENTAL IMAGING

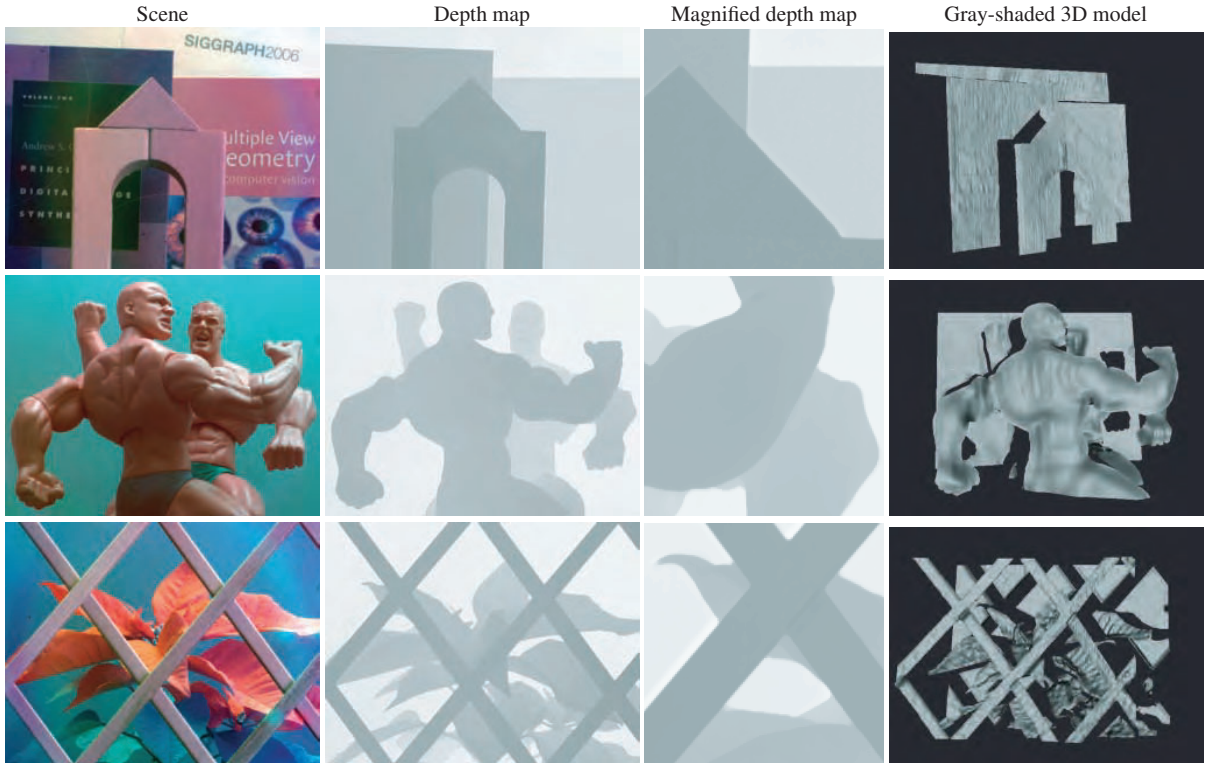


Figure 4: Depth recovery results for three scenes. From left to right: An image of the scene, the computed intensity-coded depth map, a close-up view of the depth map, a gray-shaded 3D model of the scene, as seen from a novel viewpoint. Notice the quality of the computed depth around scene discontinuities. **(Please see the companion video.)**

We tilt a foam board in front of the system and compute depth for each of its points. This is done by taking an image of the board with a few reference markers, as shown in Figure 3(a),[6] and computing the homography from the board to the projection plane. As detailed in [Zhang 2000], this homography allows us to estimate the position and orientation of the board[7], from which the depth of every point of the board can be easily computed. Figure 3(b) shows the intensity-coded depth map of the board. **Step 3**: We compute the θ values for all points on the board by shifting the stripe pattern and computing A_1 and A_2 for each pixel.

Figure 3(c) shows an image of the board under one of the shifted patterns, and (f) and (g) show the temporal radiance profiles of two points, **q** and **s**, on the board. As **s** is closer to the projector and more defocused than **q**, its temporal radiance profile is more blurred. Figures 3(h) and (i) show plots of the DFS coefficients A_k for the profiles in (f) and (g), respectively. Figure 3(d) shows the intensity-coded θ map and (e) shows a plot of the θ values (blue curve) for the column 350 in (d). The raw θ values include a small mount of noise and hence we fit a smooth curve which is shown as the green curve in (e). Using the depth estimation in Step 2, we can tabulate the mapping from θ to z. In our current implementation, we build a lookup table for each column, assuming that the defocus kernel is vertically invariant but has some horizontal variation. Figure 3(j) shows three θ–z mapping curves for the columns 100, 350 and 600. An even more comprehensive calibration would involve translating the board across the whole working volume and generating a lookup table for each pixel. Such a calibration would account for higher-order projector lens aberrations as well.

[6]The glare is produced by a number of dust particles on the half-mirror. Since the half-mirror is very close to the projector, these particles are lit with roughly 100 times the intensity compared to the scene, thereby producing bright spots in the image that have the shape of the camera aperture.

[7]Zhang [2000] takes several images of a board at different orientations to estimate a camera's intrinsic parameters. In our work, we approximately estimate the projector's intrinsics from the frustum specification in its manual. Therefore, we only take one image of the board to compute its pose.

3.5 Depth Recovery Results

Once we know the mapping between θ and z, the depth recovery is straightforward. We take 24 images while shifting the illumination pattern across the scene. From these images, we compute θ for each camera pixel using Eqs. (2,3) and then transform the θ image to a depth image using the pre-computed lookup table. Figure 4 shows several depth recovery results. In the first row, we show results for a scene with books and wooden blocks that has simple boundaries. Our method recovers the sharp depth discontinuities at the boundaries. In the second row, we show results for a scene with two toy wrestlers with similar skin color. Again, our method can separate the two figures with clean boundaries, even though the surfaces have a specular component. In the third row, we show results for a scene with leaves behind a fence. This result shows that the performance of our method is independent of occlusion complexity. In fact, it has been shown [Scharstein and Szeliski 2003] that it is very difficult to obtain a depth map that is complete in any camera view for a scene like the one with leaves using triangulation-based methods, even if multiple cameras and projectors are used. All the results are computed within a minute using Matlab.

We have also computed noise statistics for our depth estimation. Specifically, we chose 20 planar patches from the experimented scenes that are located at different depth within the working volume. We then computed the mean and standard deviation of the

Exhibit 1010 page 49 of 78
DENTAL IMAGING

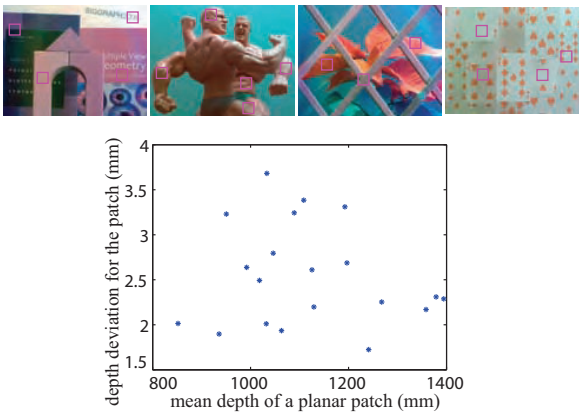


Figure 5: Noise statistics in the depth estimation. Twenty planar patches are chosen from the experimented scenes, shown on the top, and the mean and standard deviation of the depth for those patches are plotted at the bottom. Within a 600mm working volume, our depth estimation noise is around 4mm.



(a) (b) (c)

Figure 6: (a) Refocused images of the book scene in Figure 4. The defocus effects are synthesized using the recovered depth map. (b) A collection of playing cards placed at four depth layers and their estimated depth map. (c) A moving hand from another video is inserted into the card scene with all the desired occlusion effects. **(Please see the companion video.)**

depths inside these patches, as shown in Figure 5. From this figure, we can see that our depth estimation noise is about 4mm for a working volume of 600mm.

3.6 Applications: Refocusing & Video Composition

A distinctive feature of our depth recovery method is that it estimates depth at every pixel in the camera's view and the estimation is reliable near depth discontinuities. This makes our approach particularly well-suited for a variety of image-based applications in computer graphics. Here, we show two examples, namely, refocusing and object insertion to compose a new image or video.

Figure 6(a) shows two refocused images of the book scene in Figure 4: the top one is focused on the foreground wooden blocks and the bottom one is focused on the background letters. These two images were generated with the "lens blur" tool in Adobe PhotoShop.[8] This tool takes an image and its depth map as inputs to generate new images that are focused at any desired depth.

[8]In theory, without having the full light field available, refocusing can not be simulated exactly. The PhotoShop implementation of lens blur is proprietary. We think a practical way to implement it would be to heuristically impaint the occluded layers a little bit to obtain the full light field.

Figure 6(b) shows a collection of playing cards placed at four depth layers. The repetitive textures on the cards and the complex occlusions make it very hard to segment the scene into layers automatically, or even manually. Using projection defocus, our method produces a clean depth map that segments the scene into layers. This segmentation makes it easy to insert objects into the scene with correct occlusion effects. Figure 6(c) shows a real hand from another video inserted into the scene with all the desired occlusion effects. **Please see the video for more results.**

We believe that when a complete depth map is available for an image, a variety of image editing applications become possible or easier to implement. Examples other than the ones we have shown are creating layered representations of complex scenes for view morphing, segmentation for matting, object replacement, and shadow removal. In the past, such applications have not been easy to implement as "clean" depth maps have not been easy to obtain. For most image editing tasks, the depth maps need not be highly precise (as with time-of-flight sensors), but they need to be complete and they need to be reliable at discontinuities, as humans are particularly sensitive to edge artifacts. We hope our depth recovery method, although not as accurate as laser scanners or structured light sensors, will inspire new types of image editing applications.

4 Focused Projection at Multiple Depths

Due to their optical design, standard projectors can only be focused on a single fronto-parallel plane. In some applications, it is desirable to project images onto non-planar structures; for example, multiple planes or a dome to create a virtual environment [Raskar et al. 1998]. In such cases, most parts of the image are blurred due to defocus. One way to solve this problem is to design sophisticated optics. Even if this is possible, the optics cannot be modified to accommodate changes in the structure that is being projected onto. Another approach is to use multiple projectors [Bimber and Emmerling 2006], where the number of depths for which the projected image can be in focus equals the number of projectors. In this section, complementary to the multi-projector approach, we propose a computational method that processes an input image in a scene-dependent way to minimize the defocus blur at all points in the projected output image. Our method requires just a camera and can be applied to any off-the-shelf projector.

4.1 Defocus Compensation Algorithm

Consider the scenario where an image I needs to be projected onto a surface with given depth variation. The radiance of a point on the surface due to illumination by the projector is governed by the projection defocus equation Eq. (1). As in previous work, e.g. [Raskar et al. 2003; Fujii et al. 2005], we use a camera as a proxy for the human eye and seek to make the scene radiance captured by the camera be the same as the input image I. To do so, we can project a *compensation image*, P^*, by solving the projection defocus equation as

$$P^* = (\alpha f)^{-1} * (I - \beta), \tag{4}$$

where $(\alpha f)^{-1}$ is the inverse of the kernel αf. Theoretically, if P^* is projected, the scene brightness will be the same as the input image I. However, f is a low-pass filter and its inverse will have strong ringing effects. Therefore, Eq. (4) will not be feasible to implement unless the projector has an infinite dynamic range. Instead, we cast the problem of computing the compensation image as a constrained minimization problem, as follows:

$$P^* = \arg\min_P \{ d(\alpha f * P + \beta, I) \mid \forall \mathbf{x},\ 0 \le P(\mathbf{x}) \le 255 \}, \tag{5}$$

where $\mathbf{x}$ is the projector pixel coordinate and $d(\cdot,\cdot)$ is an image distance metric. Eq. (5) finds the compensation image P^*, with all its brightness values within the projector's dynamic range, that after defocus blurring most closely matches the input image I.

Exhibit 1010 page 50 of 78
DENTAL IMAGING

(a) (b)

Figure 7: Measuring the spatially-varying defocus kernel. (a) A portion of the dot pattern projected onto the scene. (b) The corresponding camera image with samples of the point spread function.

(a) (b)

Figure 8: Two test scenes for focused projection at multiple depths. (a) Projection onto three planes at different depths from the projector. (b) Projection onto a hemispherical dome.

In our implementation, we have used the sum-of-squared pixel differences for the image distance metric $d(\cdot,\cdot)$. The compensation image P^* is found by applying an iterative, constrained, steepest-descent algorithm [Nocedal and Wright 1999]. We represent the defocus convolution, $\alpha f * P$, as a matrix multiplication, FP, where each row of F is the defocus kernel of the corresponding pixel modulated by its albedo. The algorithm starts with $P_0 = I$ and iterates the following two steps:

$$\tilde{P}_{i+1} = P_i + \eta_i G_i \tag{6}$$

$$P_{i+1} = \texttt{CLAMP}(\tilde{P}_{i+1};\, 0, 255) \tag{7}$$

where $G_i = F^{\mathrm{T}}(I - \beta - FP_i)$, $\eta_i = \frac{\|G_i\|^2}{\|FG_i\|^2}$, and CLAMP is a pixel-wise clamping operation. Notice that G_i is the gradient of the image distance $\|FP + \beta - I\|^2$ with respect to P. Evaluating G_i is straightforward – it involves two image filterings with the kernel matrices F and F^{T}, respectively. Note that these filterings are spatially-varying and scene-dependent, unlike the standard sharpening operations that are often built into projectors by their manufacturers. Without Eq. (7), iterating Eq. (6) alone is a standard steepest-descent algorithm, which converges to the solution of Eq. (4). Combining Eq. (7) and Eq. (6) minimizes the difference between the defocused compensation image and the original input image within the dynamic range of the projector.

4.2 Kernel Estimation

We now describe how we obtain β and αf for any given scene. We obtain the ambient term β by turning off the projector and taking an image. To obtain αf for each pixel, we could shift a dot pattern across the scene and the temporal radiance profile for each pixel is then its kernel. This method is effective for a scene with intricate geometry, such as the scenes shown in Section 3. However, the scenes used in our experiments here are at least piece-wise smooth. Therefore, we simply project a sparse binary dot pattern, like the one in Figure 7(a). In our experiments, the distance between neighboring dots is 12 pixels. The point spread patterns captured by the camera, shown in Figure 7(b), are approximately the kernels for the projector pixels that are '1'. From these kernels, we interpolate the kernels for other pixels that are '0' in a bilinear manner. In the end, we have a per-pixel kernel map that accounts for spatially-varying defocus effects. We repeat the same procedure to capture kernels for each of the three color channels of the projector. Calibrating such a kernel map also helps to compensate lens aberrations as well. For the above computations of β and αf, we always warp the image from the camera to the projector's coordinate frame, which makes it convenient to compare the defocused compensation image, $FP + \beta$, and the input image I. The correspondences between the camera's and the projector's pixels are determined by shifting sinusoidal patterns in both horizontal and vertical directions, as in Section 3.

4.3 Focused Projection on Multiple Planes

In our first experiment, our goal was to project an image of three fashion models onto three planes that are at different depths, as shown in Figure 8(a). The projector is focused on the middle plane. We show the focus compensation results for the three face regions in the left three columns of Figure 9. The first row of Figure 9 shows the original image regions for the three faces. The third row shows the defocused original image regions, captured by a camera, without any compensation. This row represents the best the projector (without compensation) can produce. Notice that even though the middle plane is in focus, the face still looks a little blurred for this plane due to projector artifacts.[9] As expected, the faces on the left and right planes are even more blurred due to defocus. The second row shows the compensation image regions resulting from Eqs. (6,7), which represent the new input to the projector. These image regions look like high-pass filtered versions of the original input ones. This is expected as our method boosts high-frequency components to compensate for the defocus induced by scene geometry. The fourth row shows the defocused compensation image regions captured by the same camera, which represent our results. Notice that these image regions are less blurry and more closely resemble the input image regions. Among these three focused compensation image regions, the one on the middle plane (second column) is almost identical to the input. The ones on the left and right planes (first and third columns) are slightly blurred with respect to their inputs because defocus is more severe in these two cases, and certain high frequencies in the input image that are cut by the defocus cannot be fully compensated for due to the limited dynamic range of the projector.

4.4 Focused Projection on a Dome

In our second experiment, we project an image of a tiled flower texture onto the inside of a hemispherical dome, as shown in Figure 8(b). The projector is focused on the front plane of the dome. In the right two columns of Figure 9, we show the focus compensation results for two patches on the dome – one is from the top of the dome and the other is from the center. The top patch is less blurred than the center one because the former is closer to the focal plane of the projector. As with the three-plane experiment, the compensation results closely resemble the input images.

5 Depixelation

In this section, we show how the defocus compensation method discussed in Section 4 can be used to reduce projector pixelation.

[9] We found that even if we focus a projector on a fronto-parallel plane, when projecting a binary dot pattern, we see "light leakage" from one pixel into its neighbors. We observed this effect on different projector models.

Exhibit 1010 page 51 of 78
DENTAL IMAGING



Figure 9: Results for defocus compensation. The three faces on the left were projected onto the three planar surfaces in Figure 8(a). The two flower textures on the right were projected onto the hemispherical dome in Figure 8(b). The defocus compensation method results in less blurry image regions (fourth row) than the uncompensated image regions (third row). The original image regions are shown in the first row and the compensation image regions computed by the method and used as input to the projector are shown in the second row. The compensation method is scene independent as it can handle a spatially-varying defocus kernel. **(Please see the companion video.)**

Pixelation is a clearly noticeable artifact produced by all digital projectors. It is caused by two factors. The first is the spatial digitization due to the finite resolution of the projector. The second is the gap (dead-zone) between adjacent pixels on the projector's physical image plane that arises because the pixel fill-factor is never 100 %. The digitization is known to create jaggy boundaries when the resolution of the projector is not high enough for the given application. This effect is observed particularly when high quality captured images are projected, as the resolution of most LCD and DLP projectors has remained at 1024x768 or less for the last seven years, while digital cameras have tripled in their resolution during the same time period. The dead-zone between pixels does not generate light and produces thin black lines on the projection screen, known as the *screen-door effect*. This effect makes pixelation more pronounced as it clearly marks out pixel boundaries on the screen.

In our implementation of depixelation, we assume that the images are projected onto a single fronto-parallel screen, i.e., the traditional projection scenario. Our basic approach is to focus the projector slightly in front of (or behind) the projection screen so that the image on the screen is slightly blurred. In this case, a slight amount of light is leaked into the black gaps and neighboring pixels, which reduces the screen-door effect as well as smoothes out the jaggy pixel boundaries. However, doing so also creates blurry images. Since this purposely induced blur is very slight, we can use the compensation method in Section 4 to process the input image so that the defocused projection looks very similar to the original input image. Therefore, the compensation method can be used in conjunction with the minor induced defocus to achieve depixelation.

Specifically, we use the compensation method to compute an optimal image P^* at the projector resolution whose defocused projection most closely resembles the input image, which may have higher resolution. Note that if the input image resolution is higher than the projector's native resolution, the number of rows in F is larger than the number of columns. Each row of F is a defocus kernel modulated by both surface albedo and the screen-door effect.

We have tested the depixelation method on several high resolution input images (each one with 3 times the projector resolution). The capture of β and F is done using the method in Section 4, except that we capture a higher resolution kernel map by zooming-in our camera so that 3×3 camera pixels see about one projector pixel.

Exhibit 1010 page 52 of 78
DENTAL IMAGING

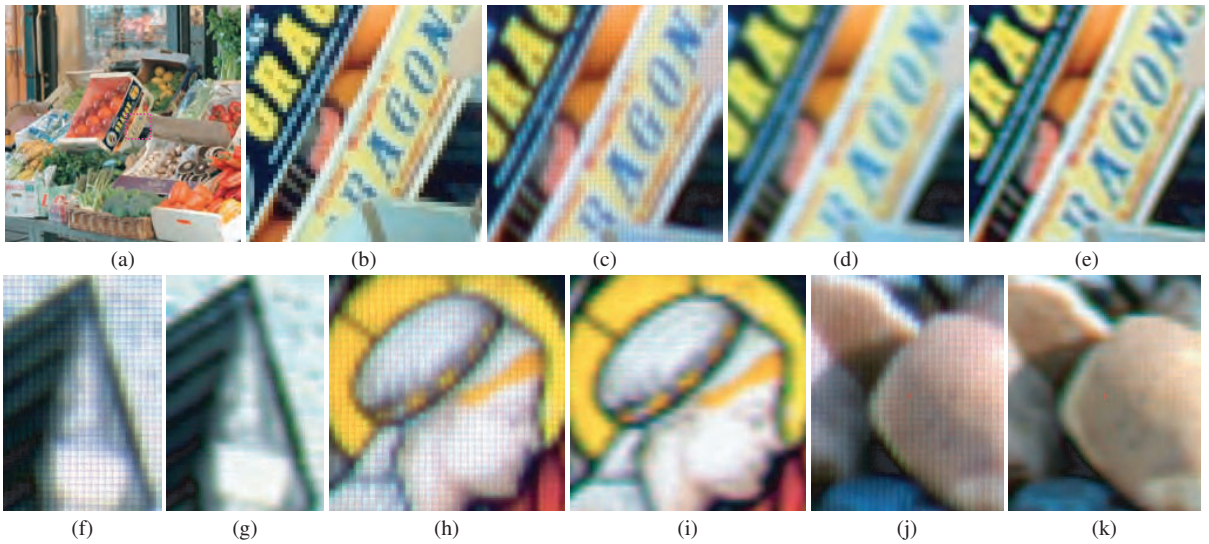


Figure 10: Examples of Projector Depixelation. (a) An original image of 3× projector resolution. (b) A patch in (a) of projector resolution. (c,d) Projection of (b) under focused and defocused settings, respectively. (e) Our depixelation result. (f-k) Three more examples of depixelation with (f,h,j) as focused projection and (g,i,k) as depixelated projection. In all examples, the compensation greatly reduces pixelation effects while preserving image quality. **(Please see the companion video.)**

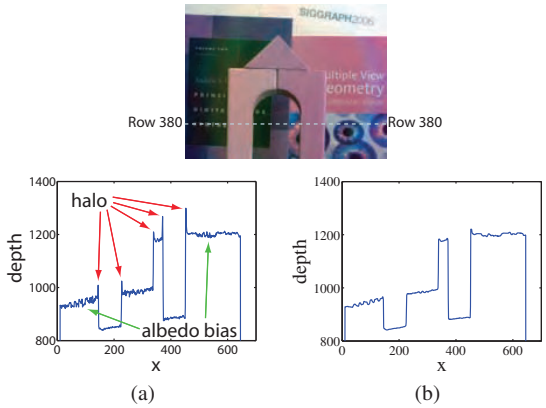


Figure 11: Albedo-correlated bias and halo effects in depth estimation. (a) The depth profile for row 380 of the scene. (b) The depth profile after removing noise and halo by a 7×7 median filter.

Figure 10 shows four examples of depixelation. In all the cases, spatial digitization and screen-door effects are greatly reduced and yet image sharpness is preserved. Our current implementation of the compensation algorithm uses Matlab. The algorithm takes about 10 iterations to converge and takes 3-5 minutes to produce the final compensation image. Since the compensation method only involves filtering with F and F^{T}, we believe it can be implemented on image processing chips or graphics hardware to achieve real-time (frame-rate) performance.

6 Discussion

In this paper, we have modelled projection defocus as a linear system and developed a temporal defocus method that recovers depth at all pixels in an image, regardless of the scene complexity. We also proposed a scene-dependent filtering algorithm that increases the depth of field of a projector without modifying its optics. Finally, we showed that defocus can be exploited to reduce pixelation artifacts produced by digital projectors. We now discuss the limitations of our work and open problems as topics for future research.

First, our temporal defocus method should, in principle, be invariant to surface albedo. In practice, we found that the surface albedo sometimes does affect the depth estimation, as shown in Figure 11. We believe that this is only an issue with our current implementation. Our beam-splitter not only reflects light from the scene to the camera, but also transmits light from a black backdrop placed behind it. Therefore, the defocus kernel f in Eq. (1) is actually a sum of the defocus kernels of both the scene and the backdrop. If the scene is much brighter than the black backdrop, Eq. (3) is dominated by the scene's defocus kernel. However, when the scene point is dark, due to low albedo or because its surface orientation is near-grazing for the projector, the light from the scene point to the camera is weak and comparable to the light from the backdrop. In this case, Eq. (3) will depend also on the defocus kernel of the backdrop, which causes errors in depth estimation. The backdrop kernel remains constant and can be calibrated to reduce depth estimation errors. Alternatively, the beam-splitter can be placed in a custom-designed chamber that absorbs all forms of stray light.

The projector defocus is for most cases independent of scene geometry, and this independence makes is superior to camera defocus for depth estimation. However, there is a special case that this independence does not hold. Specifically, when a closer object occludes a distant object, there will be a narrow band of points near the occluding contour on the distant object that "see" only part of the projector's aperture; so the projected pattern will be more infocus than it would be if the entire aperture were visible. This partial aperture effect results in a halo artifact: depth estimation biased toward the focal plane, as shown in Figure 11. However, even in this special case, projector defocus is still better than camera defocus. Camera defocus will cause estimation errors for both the foreground and background pixels near the occluding boundary. In contrast, projector defocus will only cause errors at the background pixels, while leaving the foreground pixels unaffected. Furthermore, in practice, because the solid angle subtended by the projector aperture is only about 1 degree with respect to a scene point, the number of affected background pixels is small. Therefore, as shown in Figure 11, the halo effect can be greatly reduced by a simple median filtering.

Exhibit 1010 page 53 of 78
DENTAL IMAGING

In practice, using projector defocus for depth estimation does not require moving detectors or changing optical settings as traditional camera defocus methods do [Pentland 1987; Nayar and Nakagawa 1994]. Therefore, our method can be more conveniently implemented. However, a coaxial implementation of depth from camera defocus, as in [Nayar et al. 1996], allows for capturing dynamic scenes, which our method currently is incapable of, because we use 24 images for depth estimation. We believe this number can be greatly reduced, which would enable us to develop a real-time sensor that can handle dynamic scenes. In Eq. (1), there are three unknowns: α, β, and z. If we can calibrate the defocus kernel f a priori, then, in theory, 3 images are enough to solve for these unknowns. Interestingly, this minimum number is the same for other shape acquisition methods, like the phase-shift method [Huang et al. 2003], the time-of-flight method [Gonzales-Banos and Davis 2004], and photometric stereo [Horn and Brooks 1989].

Our temporal defocus method can deal with specular surfaces, whenever the BRDF's are approximately constant within the solid angle subtended by the projector aperture from the scene point. Nevertheless, the extreme case of near-mirror reflection, as well as saturated highlight pixels, will cause errors. Although our method is developed under the direct illumination assumption, it's not very sensitive to the indirect illumination. This is because indirect illumination usually creates only very low frequency intensity fluctuation, which will mainly affect Fourier coefficient A_0. Our depth is however estimated via A_1 and A_2. Of course, for extremely specular reflection, the global illumination can change dramatically with the projected patterns and hence our method will break down.

One challenging open problem for defocus-based methods is handling translucent objects. A translucent material introduces additional blur due to subsurface scattering, which would cause a systematic bias in our methods. The amount of bias depends on the amount of translucency blur compared to defocus blur. If the defocus blur dominates the translucency blur, the error will be small. Otherwise, the error will be large. It would be interesting to extend our analysis to handle blur induced by translucency. In practice, the depth of mixture pixels at boundaries is often a blend of its neighboring depth values. It would be interesting to accurately model the mixture pixels, e.g., combining our approach with the defocus matting method [McGuire et al. 2005], to estimate mattes for scenes with continuous depth variations.

Our method for defocus compensation essentially extends the depth of field of the projector. However, when the blur is very severe, our method cannot fully compensate for the blur due to the limited dynamic range of the projector. In such a case, combining our method with multiple-projector solutions [Bimber and Emmerling 2006] could be beneficial. Finally, we plan to extend our work on depixelation to make it a practical and useful projector feature. How objectionable pixelation artifacts are, is really a subjective issue. For our technique to be incorporated into commercial projectors, a few preset options can be provided so that the user can decide how much focus quality they are willing to tradeoff for depixelation.

References

BIMBER, O., AND EMMERLING, A. 2006. Multi-focal projection. *IEEE Trans. on Visualization and Computer Graphics to appear*.

BIMBER, O., WETZSTEIN, G., EMMERLING, A., AND NITSCHKE, C. 2005. Enabling view-dependent stereoscopic projection in real environments. In *Proc. Int. Symp. on Mixed and Augmented Reality*, 14–23.

CURLESS, B., AND LEVOY, M. 1995. Better optical triangulation through spacetime analysis. In *Proc. Int. Conf. on Computer Vision*, 987–994.

DAVIS, J., NEHAB, D., RAMAMOOTHI, R., AND RUSINKIEWICZ, S. 2005. Spacetime stereo : A unifying framework for depth from triangulation. *IEEE Trans. on Pattern Analysis and Machine Intelligence 27*, 2, 296–302.

FAUGERAS, O. 1993. *Three-Dimensional Computer Vision*. MIT Press.

FAVARO, P., AND SOATTO, S. 2005. A geometric approach to shape from defocus. *IEEE Trans. on Pattern Analysis and Machine Intelligence* (in press).

FUJII, K., GROSSBERG, M., AND NAYAR, S. 2005. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 814–821.

GIROD, B., AND SCHEROCK, S. 1989. Depth from defocus of structured light. In *Proc. SPIE Conf. on Optics, Illumination, and Image Sensing for Machine Vision*.

GONZALES-BANOS, H., AND DAVIS, J. 2004. A method for computing depth under ambient illumination using multi-shuttered light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 234–241.

GROSSBERG, M., PERI, H., NAYAR, S., AND BELHUMEUR, P. 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. I, 452–459.

HORN, B., AND BROOKS, M. 1989. *Shape from Shading*. MIT Press.

HUANG, P. S., ZHANG, C. P., AND CHIANG, F. P. 2003. High speed 3-d shape measurement based on digital fringe projection. *Optical Engineering 42*, 1, 163–168.

JAIN, A. K. 1989. *Fundamentals of Digital Image Processing*. Prentice Hall.

JIN, H., AND FAVARO, P. 2002. A variational approach to shape from defocus. In *Proc. Eur. Conf. on Computer Vision*, 18–30.

KANADE, T., GRUSS, A., AND CARLEY, L. 1991. A very fast vlsi rangefinder. In *Proc. Int. Conf. on Robotics and Automation*, vol. 39, 1322–1329.

KONINCKX, T. P., PEERS, P., DUTR, P., AND GOOL, L. V. 2005. Scene-adapted structured light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 611–619.

LEVOY, M., CHEN, B., VAISH, V., HOROWITZ, M., MCDOWALL, I., AND BOLAS, M. 2004. Synthetic aperture confocal imaging. In *SIGGRAPH Conference Proceedings*, 825–834.

LJUNG, L. 1998. *System Identification: A Theory for the User*, 2 ed. Prentice Hall.

MAJUMDER, A., AND WELCH, G. 2001. Computer graphics optique: Optical superposition of projected computer graphics. In *Proc. Eurographics Workshop on Virtual Enviroment/Immersive Projection Technology*.

MCGUIRE, M., MATUSIK, W., PFISTER, H., HUGHES, J. F., AND DURAND, F. 2005. Defocus video matting. In *SIGGRAPH Conference Proceedings*, 567–576.

NAYAR, S. K., AND NAKAGAWA, Y. 1994. Shape from focus. *IEEE Trans. on Pattern Analysis and Machine Intelligence 16*, 8, 824–831.

NAYAR, S. K., WATANABE, M., AND NOGUCHI, M. 1996. Real-time focus range sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence 18*, 12, 1186–1198.

NOCEDAL, J., AND WRIGHT, S. J. 1999. *Numerical Optimization*. Springer.

OPPENHEIM, A. V., AND WILLSKY, A. S. 1997. *Signals and Systems*, 2 ed. Prentice Hall.

PENTLAND, A. 1987. A new sense for depth of field. *IEEE Trans. on Pattern Analysis and Machine Intelligence 9*, 4, 523–531.

RAJ, A., AND ZABIH, R. 2005. A graph cut algorithm for generalized image deconvolution. In *Proc. Int. Conf. on Computer Vision*.

RAJAGOPALAN, A. N., AND CHAUDHURI, S. 1997. A variational approach to recovering depth from defocused images. *IEEE Trans. on Pattern Analysis and Machine Intelligence 19*, 10, 1158–1164.

RASKAR, R., WELCH, G., CUTTS, M., LAKE, A., STESIN, L., AND FUCHS, H. 1998. The office of the future : A unified approach to image-based modeling and spatially immersive displays. In *SIGGRAPH Conference Proceedings*, 179–188.

RASKAR, R., WELCH, G., LOW, K., AND BANDYOPADHYAY, D. 2001. Shader lamps. In *Proc. Eurographics Workshop on Rendering*.

RASKAR, R., VAN BAAR, J., BEARDSLEY, P., WILLWACHER, T., RAO, S., AND FORLINES, C. 2003. ilamps: geometrically aware and self-configuring projectors. In *SIGGRAPH Conference Proceedings*, 809–818.

RASKAR, R., HAN TAN, K., FERIS, R., YU, J., AND TURK, M. 2004. Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. In *SIGGRAPH Conference Proceedings*, 679–688.

SCHARSTEIN, D., AND SZELISKI, R. 2003. High-accuracy stereo depth maps using structured light. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 195–202.

SCHECHNER, Y. Y., KIRYATI, N., AND BASRI, R. 2000. Separation of transparent layers using focus. *Int. J. on Computer Vision 39*, 1, 25–39.

SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. In *SIGGRAPH Conference Proceedings*, 745–755.

TAPPEN, M. F., RUSSELL, B. C., AND FREEMAN, W. T. 2004. Efficient graphical models for processing images. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 2, 673–680.

ZHANG, L., SNAVELY, N., CURLESS, B., AND SEITZ, S. M. 2004. Spacetime faces: High-resolution capture for modeling and animation. In *ACM Annual Conference on Computer Graphics*, 548–558.

ZHANG, Z. 2000. A flexible new technique for camera calibration. *IEEE Trans. on Pattern Analysis and Machine Intelligence 22*, 11, 1330–1334.

Exhibit 1010 page 54 of 78
DENTAL IMAGING

APPENDIX A-3

Pages

WorldCat

Home

Transfer old account | Create account | Sign in

It... libraries

Search for books, articles, and more

Topics

Lists

About

For Librarians

Exhibit 1010 page 56 of 78
DENTAL IMAGING



Proceedings

Authors: ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, SIGGRAPH, Association for Computing Machinery

Journal, Magazine, English, c2005-
Edition: View all formats and editions
Publisher: Association for Computing Machinery, New York, NY, c2005-

Genre:
Conference papers and proceedings

Physical Description:
v. : ill. ; 28 cm.

ISSN:
1931-4027

OCLC Number / Unique Identifier:
62096329

Subjects:
Computer graphics
Computer graphics Congresses
Conference papers and proceedings
Three-dimensional display systems
Three-dimensional display systems Congresses

Show more

More Information:
portal.acm.org

Show less information

Find a Copy at a Library Filter by: Any format Any edition Distance within 200+ mi

***14 editions** in 459 libraries*

Exhibit 1010 page 57 of 78
DENTAL IMAGING

Featured libraries All libraries

Showing 388 libraries near Washington, District of Columbia Only public libraries

Academic



George Washington University

1.2 miles from your current location.

Melvin Gelman Libr, 2130 H St NW
Washington, DC 20052-2521
United States

Get Directions

Favorite Library

Borrow

Academic



Howard University

1.4 miles from your current location.

Founders Library, 500 Howard Place NW
Washington, DC 20059
United States

Get Directions

Favorite Library

Borrow

Regional or National



Library of Congress

2.2 miles from your current location.

101 Independence Avenue, SE
Washington, DC 20540
United States

Get Directions

Favorite Library

Borrow

Academic

American University

Exhibit 1010 page 58 of 78
DENTAL IMAGING



American University

American University Library; Bender Library

3.5 miles from your current location.

4400 Massachusetts Avenue Northwest
Washington, DC 20016-8046
United States

Get Directions

Favorite Library

Borrow

Academic



Marymount University

Emerson G. Reinsch Library

5 miles from your current location.

2807 North Glebe Road, 2807 N Glebe Rd
Arlington, VA 22207-4299
United States

Get Directions

Favorite Library

Borrow

Other



Institute for Defense Analyses Library

IDA Library

6.8 miles from your current location.

4850 Mark Center Dr., 4850 Mark Ctr Dr
Alexandria, VA 22311
United States

Get Directions

Favorite Library

Borrow

Academic



University of Maryland Libraries

UMD Libraries

Favorite

Exhibit 1010 page 59 of 78
DENTAL IMAGING

Library

7.3 miles from your current location.

McKeldin Library - ILL, University of Maryland, 7649 Library Lane
College Park, MD 20742-7011
United States

Get Directions

Borrow

Academic

George Mason University

Fenwick Library

Favorite Library

16 miles from your current location.

4400 University Drive
Fairfax, VA 22030-4444
United States

Get Directions

Borrow

Academic

Divine Mercy University

Mary S. Thelen Library

Favorite Library

22 miles from your current location.

45154 Underwood Lane
Sterling, VA 20166
United States

Get Directions

Borrow

Academic

Christendom College Libraries

St. John the Evangelist Library

Favorite Library

60 miles from your current location.

263 St. Johns Way
Front Royal, VA 22630
United States

Get Directions

Borrow

Exhibit 1010 page 60 of 78
DENTAL IMAGING

1 2 3 4 5 6 ... Results per Page: 10

Borrow from **George Washington University** near Missouri, United States

Borrow

1.2 miles away

WorldCat is the world's largest library catalog, helping you find library materials online.

Home
Libraries
Topics
Lists
About
For Librarians
Help & FAQ
Contact Us

Language

English

Copyright © 2001-2022 OCLC. All rights reserved.
Terms and Conditions Help

9631fcf3-7541-437b-9bd5-26ccad4649f2

© 2022 OCLC Domestic and international trademarks and/or service marks of OCLC, Inc. and its affiliates. Cookie notice
Manage my cookies

Privacy Policy Accessibility Statement ISO 27001 Certificate

Exhibit 1010 page 61 of 78
DENTAL IMAGING

APPENDIX A-4

LIBRARY OF CONGRESS | CATALOG

1. Proceedings

LCCN	2005205010
Type of material	Periodical or Newspaper
Meeting name	ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games.
Main title	Proceedings / ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games.
Published/Created	New York, NY : Association for Computing Machinery, c2005-
Publication history	2005-
Description	v. : ill. ; 28 cm.
Current frequency	Biennial
ISSN	1931-4027
Linking ISSN	1931-4027
LC classification	T385 .A2783
Continues	ACM SIGGRAPH Symposium on Interactive 3D Graphics. Proceedings 1551-1669 (DLC) 2004213756 (OCoLC)55859011
Serial key title	Proceedings (ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games)
Related names	SIGGRAPH. Association for Computing Machinery.
LC Subjects	Computer graphics--Congresses. Three-dimensional display systems--Congresses.
Other Subjects	Computer graphics. Three-dimensional display systems.
Form/Genre	Conference papers and proceedings.
Notes	Sponsored by: ACM SIGGRAPH
Reproduction no./Source	ACM Order Dept., P.O. Box 11405, Church Street Station, New York, NY 10286-1405
Additional formats	Available also online as: SI3D Proceedings; TOC and abstracts freely available; full text available only to ACM members and subscribers of ACM's Portal Digital Library.
Dewey class no.	006
Other system no.	(OCoLC)ocm62096329
Additional Links	http://portal.acm.org/browse%5Fdl.cfm?linked=1&part=series&idx=SERIES749&coll=ACM&dl=ACM&CFID=37202478&CFTOKEN=24749567

CALL NUMBER	T385 .A2783 Set 1
Request in	Jefferson or Adams Building Reading Rooms
Latest receipts	2009 2008 2007 2006 2005

LC CATALOG
Library of Congress
101 Independence Ave., SE

Exhibit 1010 page 63 of 78
DENTAL IMAGING

Washington, DC 20540

Questions? Ask a Librarian:
https://ask.loc.gov/

Exhibit 1010 page 64 of 78
DENTAL IMAGING

APPENDIX B

JAMES L. MULLINS, PhD

Prior Art Documentation Librarian Services, LLC

205 Saint Cuthbert, Williamsburg, VA 23188

jlmullins@priorartdoclib.com

ph. 765 479 4956

Prior Art Documentation Librarian Services, LLC. (PADLS). Founded January 2018. As of August, 2022, 103 declarations have been completed for 35 law firms; four depositions scheduled, three cancelled a few days prior, and only one was held in February, 2019.

Library Experience:

2018-present Dean Emeritus, Libraries & Norton Professor Emeritus, Purdue University

2011 - 2017 Dean of Libraries & Esther Ellis Norton Professor, Purdue University

2004 - 2011 Dean of Libraries & Professor, Purdue University, West Lafayette, IN

2000-2004 Assistant/Associate Director for Administration, MIT Libraries, Massachusetts Institute of Technology, Cambridge, MA.

1996-2000 University Librarian & Director, Falvey Memorial Library. Villanova University, Villanova, PA.

1978-1996 Director of Library Services, Indiana University South Bend.

1974-1978 Associate Librarian, Indiana University Bloomington, School of Law. 1973-1974 Instructor/Catalog Librarian. Georgia Southern College (now University).

Teaching Experience:

1977-1996 Associate Professor (part-time), School of Library and Information Science, Indiana University. Subjects taught: Cataloging, Management, and Academic Librarianship.

Education:

1972 - University of Iowa. Honors Bachelor of Arts, History, Religion & Political Science.

1973 - University of Iowa. Master of Arts in Library Science.

1984 - Indiana University. Doctor of Philosophy. Concentration: Academic Library Administration. Emphasis: Law Librarianship.

Awards and Recognition:

2017 Wilmeth Active Learning Center/Library of Engineering and Science, Grand Reading Room, was announced by President Mitch Daniels, Purdue University, that it would be re-named the James L. Mullins Reading Room to honor his leadership and reputation in the academic library profession. September 2017. Portrait unveiled December 2017.

2017 Distinguished Alumnus Award by the School of Informatics and Computing, Indiana University, Bloomington. Given June 25, 2017.

2016 Hugh C. Atkinson Memorial Award, jointly sponsored by the four divisions of the American Library Association (ALA), June 27, 2016.

2015 ACRL Excellence in University Libraries Award, April 23, 2015.

Named Esther Ellis Norton Professor of Library Science by Purdue Trustees, December 11, 2011.

International Review Panel to evaluate the University of Pretoria Library, February 20 – 24, 2011. Pretoria, South Africa.

Publications: (selected)

Designing the Academic Library, Chapter Contributor. Association of College and Research Libraries, American Library Association, 2022.

A Purdue Icon: Creation, Life, and Legacy, edited by James L. Mullins, Founder's Series, Purdue University Press, 138pp., August 2017.

"The policy and institutional framework." In *Research Data Management, Practical Strategies for Information Professionals,* edited by Ray, J M. Purdue University Press, pp.25-44, 2014.

"DataCite: linking research to data sets and content." In Benson, P and Silver, S. *What Editors Want: An Author's Guide to Scientific Journal Publishing*. University of Chicago Press, pp. 21-23, December 2012.

"Library Publishing Services: Strategies for Success," with R. Crow, O. Ivins, A. Mower, C. Murray-Rust, J. Ogburn, D Nesdill, M. Newton, J. Speer, C. Watkinson. *Scholarly Publishing and Academic Resources Coalition (SPARC),* version 2.0, March 2012.

"The Changing Definition and Role of Collections and Services in the University Research Library." *Indiana Libraries,* Vol 31, Number 1 (2012), pp.18-24.

"Are MLS Graduates Being Prepared for the Changing and Emerging Roles that Librarians must now assume within Research Libraries?" *Journal of Library Administration.* Volume 52, Issue 1, 2012, p. 124-132

Baykoucheva, Svetla. What Do Libraries Have to Do with e-Science?: An Interview with James L. Mullins, Dean of Purdue University Libraries. Chem. Inf. Bull. [Online] 2011, 63 (1), 45-49. http://www.acscinf.org/publications/bulletin/63-1/mullins.php (accessed Mar 16, 2011).

"The Challenges of e-Science Data-set Management and Scholarly Communication for Domain Sciences and Technology: a Role for Academic Libraries and Librarians," chapter in, *The Digital Deluge: Can Libraries Cope with e-Science?"* Deanna B. Marcum and Gerald George, editors, Libraries Unlimited/Teacher Ideas Press, 2009. (a monograph publication of the combined proceedings of the KIT/CLIR proceedings).

"Bringing Librarianship to e-Science," *College and Research Libraries.* vol. 70, no. 3, May 2009, editorial.

"The Librarian's Role in e-Science" *Joho Kanri (Journal on Information Processing and Management),* Japan Science and Technology Agency (formerly Japan Information Center of Science and Technology), Tokyo, Japan. Translated into Japanese by Taeko Kato. March 2008.

The Challenge of e-Science Data-set Management to Domain Sciences and Engineering: a Role for Academic Libraries and Librarians," KIT (Kanazawa Institute of Technology)/CLIR (Council of Library and Information Resources) International Roundtable for Library and Information Science, July 5-6, 2007. Developments in e-science status quo and the challenge, The Japan Foundation, 2007.

"An Administrative Perspective," Chapter 14, *Proven Strategies for Building an Information Literacy Program*, Susan Curzon and Lynn Lampert, editors, Neal-Schuman Publishers, Inc., New York, 2007. pp. 229-237.

Library Management and Marketing in a Multicultural World, proceedings of the IFLA Management and Marketing (M&M) Section, Shanghai, China, August 16-17, 2006, edited. K.G. Saur, Munchen, Germany, June 2007. 390 pp.

Top Ten Assumptions for the Future of Academic Libraries and Librarians: a report from the ACRL Research Committee, with Frank R. Allen and Jon R. Hufford. *College & Research Libraries,* April 2007, vol.68, no.4. pp.240-241, 246.

To Stand the Test of Time: Long-term Stewardship of Digital Data Sets in Science and Engineering. A report to the National Science Foundation from the ARL Workshop on New Collaborative Relationships: the Role of Academic Libraries in the Digital Data Universe. September 26-27, 2006, Arlington, VA. p.141. http://www.arl.org/bm~doc/digdatarpt.pdf

"Enabling Interaction and Quality in a Distributed Data DRIS," *Enabling Interaction and Quality: Beyond the Hanseatic League*. 8th International Conference on Current Research Information Systems, with D. Scott Brandt and Michael Witt. Promoted by euro CRIS. Leuven University Press, 2006. pp.55-62. Editors: Anne Garns Steine Asserson and Eduard J. Simons.

"Standards for College Libraries, the final version approved January 2000," prepared by the ACRL College Libraries Standards Committee (member), *C&RL News*, March 2000, p.175-182.

"Standards for College Libraries: a draft," prepared by the ACRL College Libraries Section, Standards Committee (member), *C&RL News*, May 1999, p. 375-381.

"Statistical Measures of Usage of Web-based Resources," *The Serials Librarian,* vol. 36, no. 1-2 (1999) p. 207-10.

"An Opportunity: Cooperation between the Library and Computer Services," in *Building Partnerships: Computing and Library Professionals.* Edited by Anne G. Lipow and Sheila D. Creth. Berkeley and San Carlos, CA, Library Solutions Press, 1995. p. 69-70.

"Faculty Status of Librarians: A Comparative Study of Two Universities in the United Kingdom and How They Compare to the Association of College and Research Libraries Standards, " in *Academic Librarianship, Past, Present, and Future: a Festschrift in Honor of David Kaser.* Englewood, Colorado; Libraries Unlimited, 1989. p. 67-78. Review in: *College & Research Libraries,* vol. 51, no. 6. November 1990, p. 573-574.

Presentations: (Representative)

"How Long the Odyssey? Transitioning the Library and Librarians to Meet the Needs and Expectations of the 21st Century University," David Kaser Lecture, School of Informatics & Computing, Indiana University, Bloomington, IN, November 16, 2015.

Presentation at University of Cape Town, Cape Town, South Africa, August 20, 2015.

"The Challenge of Discovering Science and Technology Information," Moderator, International Federation of Library Associations (IFLA) Science and Technological Libraries Section Program, Cape Town, South Africa, August 18, 2015.

"An Odyssey in Data Management: Purdue University," International Federation of Library Associations (IFLA) Research Data Management: Finding Our Role – A program of the Research Data Alliance, Cape Town, South Africa, August 17, 2015.

Presentation at University of Pretoria, Pretoria, South Africa, August 11, 2015.

Co-Convener with Sarah Thomas, Harvard University, at the Harvard Purdue Symposium on Data Management, Harvard University, Cambridge, MA, June 15-18, 2015.

"Strategic Communication," panel discussion on the Director's role and perspective on library communications at Committee on Institutional Cooperation (CIC) Center for Library Initiatives (CLI) Annual Conference, University of Illinois Urbana-Champaign, May 20, 2015.

"Issues in Data Management," panel discussion moderated by Catherine Woteki, United States Undersecretary for Research, Education & Economics at 20th Agriculture Network Information Collaborative (AgNIC) Annual Meeting in the National Agricultural Library, Beltsville, MD, May 6, 2015.

"Active learning/IMPACT & the Active Learning Center at Purdue University," Florida Institute of Technology, Melbourne, FL, February 11, 2015.

"Science+art=creativity: libraries and the new collaborative thinking," panel moderator, International Federation of Library Associations (IFLA) 80th General Conference and Assembly, Lyon, France, August 19, 2014.

"Purdue University The Active Learning Center—A new concept for a library," Association of University Architects 59th Annual National Conference, University of Notre Dame, South Bend, IN, June 23, 2014.

"Big Data & Implications for Academic Libraries," keynote speaker, Greater Western Library Alliance (GWLA) Cyber-infrastructure Conference, Kansas City, MO, May 28, 2014.

"Research Infrastructure," panel moderator, Association of Research Libraries (ARL) 164th Membership Meeting, Ohio State University, Columbus, OH, May 7, 2014.

"An Eight Year Odyssey in Data Management: Purdue University," International Association of Scientific and Technological University Libraries (IATUL) 2013 Workshop Research Data Management: Finding Our Role, University of Oxford, UK, December 2013.

"Purdue University Libraries & Press: from collaboration to integration," Ithaka Sustainable Scholarship, The Evolving Digital Landscape: New Roles and Responsibilities in Higher Education, libraries as publishers, New York, New York, October 2013.

"Tsinghua and Purdue: Research Libraries for the 21st Century," Tsinghua University, Tsinghua, China, August 2013.

"Purdue Publishing Experience in the Libraries Publishing Coalition," Association of American University Presses Annual Meeting, Press-Library Coalition Panel, Boston, Massachusetts, June 21, 2013.

"Indiana University Librarians Day: Purdue University Libraries Ready for the 21st Century," Indiana University Purdue University Indianapolis (IUPUI), June 7, 2013.

"Purdue University Libraries and Open Access; CNI Project Update," Coalition for Networked Information, San Antonio, TX, April 5, 2013.

Memorial Resolution, honoring Joseph Brannon, to the Board of the Association of College & Research Libraries, Seattle, WA, January 2013.

"An overview of sustaining e-Science collaboration in an Academic Research Library—the Purdue experience," Duraspace e-Science Institute webcast, October 17, 2012.

"The Role of Libraries in Data Curation, Access, and Preservation: an International Perspective, " Panel Moderator, 78th General Conference and Assembly, International Federation of Library Associations, Helsinki, Finland, August 15, 2012.

"21st Century Libraries," moderator of First Plenary Session, International Association of Technological University Libraries 33rd Annual Conference, Singapore, June 4, 2012.

"Planning for New Buildings on Campus," panel presenter, University of Calgary Building Symposium on Designing Libraries for the 21st Century, Calgary, Alberta, Canada, May 17, 2012.

"Data Management and e-Science, the Purdue Response." Wiley-Blackwell Executive Seminar-2012, Washington, DC, March 23, 2012.

"An overview of Sustaining e-Science Collaboration in Academic Research Libraries and the Purdue Experience." Leadership & Career Development Program Institute, Association of Research Libraries (ARL). Houston, TX, March 21, 2012.

"An overview of Data Activities at Purdue University in response to Data Management Requirements." Coalition for Academic Scientific Computation (CASC). Arlington, VA, September 8, 2011.

"Getting on Track with Tenure," Association of College and Research Libraries (ACRL) Research Program Committee. Washington, DC, June 26, 2011.

"Integration of the Press and Libraries Collaboration to Promote Scholarly Communication," Association of Library Collections & Technical Services (ALCTS) Scholarly Communication Interest Group – American Library Association, New Orleans, Louisiana, June 25, 2011.

"Cooperation for improving access to scholarly communication," with N. Lossau (Germany), C. Mazurek (Poland), J. Stokker (Australia), panel moderator and presenter, Second Plenary Session, International Association of Scientific and Technological University Libraries (IATUL) 32nd Conference 2011, Warsaw, Poland. May 29-June 2, 2011.

"Riding the Wave of Data," STM Annual Spring Conference 2011. Trailblazing & transforming scholarly publishing **2011**. Washington, D.C., April 28, 2011.

"Confronting old assumptions to assume new roles: physical and operational integration of the Press and Libraries at Purdue University," keynote speaker, 2011 BioOne Publishers & Partners Meeting. Washington, D. C., April 22, 2011.

"Are MLS Graduates Being Prepared for the Changing and Emerging Roles that Librarians must now assume within Research Libraries?" University of Oklahoma Libraries Seminar, March 4, 2011, Oklahoma City, Oklahoma.

"The Future Role of University Librarians," the University of Cape Town, South Africa, February 25, 2011.

"New Roles for Librarians: the Application of Library Science to Scientific/Technical Research – Purdue University – a case study. International Council for Science and Technology (ICSTI); Ottawa, Canada. June 9, 2009.

"Reinventing Science Librarianship: Models for the Future," Association of Research Libraries / Coalition for Networked Information. October 16-17th, 2008, Arlington, VA. Moderator and convener of Data Curation: Issues and Challenges.

"Practical Implementation and Opportunities Created at Purdue University," African Digital Curation Conference, Pretoria, South Africa, (live video transmission), February 12, 2008.

Keynote speaker. "*Scholarly Communication & Academe: The Winter of Our Discontent,"* XXVII Charleston Conference on Issues in Book and Serial Acquisition, Charleston, South Carolina. November 8, 2007.

Keynote speaker. "*Enabling Access to Scientific & Technical Data-sets in e-Science: a role for Library and Archival Sciences,"* Greater Western Library Alliance (GWLA), Tucson, Arizona. September 17, 2007. A meeting of library directors and vice presidents for research of member institutions.

"*The Challenge of e-Science Data-set Management to Domain Sciences and Engineering: a Role for Academic Libraries and Librarians,"* KIT (Kanazawa Institute of Technology)/CLIR (Council of

Library and Information Resources) International Roundtable for Library and Information Science, July 5-6, 2007. Invited to participate by the Deputy Librarian of Congress.

International Association of Technological University Libraries (IATUL), Stockholm, Sweden. June 8, 2007**.** Invited paper**,** *Enabling International Access to Scientific Data-sets: creation of the Distributed Data Curation Center (D2C2).*

"A New Collaboration for Librarians: The Principles of Library and Archival Sciences Applied to the Curation of Datasets," Symposium of the Libraries and the College of Engineering, University of Louisville, April 6, 2007.

"Purdue University Libraries: Through Pre-eminent Innovation and Creativity, Meeting the Challenges of the Information Age," Board of Trustees, Purdue University, February 15, 2007.

ARL Workshop on New Collaborative Relationships: The Role of Academic Libraries in the Digital Data Universe, September 26-27, 2006, Arlington, VA. Invited participant.

NARA and SDSC: A partnership. A panel before the National Science Foundation, June 27, 2006. Arlington, VA. Invited participant.

"Kaleidoscope of Scientific Literacy: fusing new connections," with Diane Rein, American Library Association, Association of College and Research Libraries, Science & Technology Section, Annual Conference, New Orleans, June 26th, 2006.

"Leadership for Learning: Building a Culture of Teaching in Academic Libraries – an administrative perspective," American Library Association, Association of College and Research Libraries, Instruction Section, Annual Conference, New Orleans, June 25th, 2006.

"Building an interdisciplinary research program in an academic library:
how the Libraries'associate dean for research makes a difference at Purdue University," International Association of Technological University Libraries (IATUL), Porto, Portugal, May 23rd, 2006.

"Enabling Interaction and Quality in a Distributed Data DRIS," *Enabling Interaction and Quality: Beyond the Hanseatic League*. 8th International Conference on Current Research Information Systems, with D. Scott Brandt and Michael Witt. Promoted by euro CRIS, Bergen, Norway, May 12th, 2006, Brandt, and Witt presented in person

"Interdisciplinary Research," with D. Scott Brandt, Coalition for Networked Information (CNI) Spring Meeting: Project Briefing, Washington, D.C., April 3rd, 2006.

"An Interview with Purdue's James Mullins," a podcast submitted by Matt Pasiewicz, on *Educause Connect,* http://connect.educause.edu/James_L_Mullins_Interview_CNI_2005

"Managing Long-Lived Digital Data-sets and their Curation: Interdisciplinary Policy Issues," Managing Digital Assets Forum, Association of Research Libraries (ARL), Washington, D.C., October 28th, 2005.
"The Odyssey of a Librarian." Indiana Library Federation (ILF), District 2 Meeting, South Bend, Indiana. October 4th, 2005.
"New College Library Standards," Standards Committee Presentation, ALA, Chicago, July 7, 2000.
SUNY Library Directors, Lake George, New York. "*The College Library Standards: a Tool for Assessment."* April 5, 2000.
Tri-State College Library Association, *Finding You Have Talents You Never Knew You Had,* Penn State Great Valley, March 25, 2000.
Using Web Statistics, American Library Association, New Orleans, June 24, 1999.
Keynote speaker at the JSTOR Workshop, January 29, 30, 1999. University of Pennsylvania, Philadelphia, PA.
"The New Standards for Electronic Resources Statistics," Society of Scholarly Publishers, Washington, D.C., September 17, 1998.
"Evaluating Online Resources: Now that you've got them what do you do?," joint presenter with Chuck Hamaker, LSU, at the NASIG Conference, Boulder, Colorado. June 1998.
"What Employers Are Looking for in New Librarians?" Pennsylvania Library Association, Philadelphia. September 26, 1997.
"The Theory of Matrix Management" panel presentation of the Comparative Library Organization Committee of the Library Organization and Management Section of the Library Administration and Management Association, a division of the American Library Association, Annual Meeting, Chicago, June 24, 1990.

Professional Involvement: (summary of recent emphasis)

The focus for my professional involvement and research has moved recently toward managing massive data-sets. This has resulted in working with faculty in the sciences and technology to determine how librarians can collaborate in managing, curating, and preserving data-sets for future access and documentation. This has included various speaking opportunities as well as participation in planning with the National Science Foundation (NSF) on ways in which librarians can be integrated more completely into the funded research process. Participation in the Kanazawa Institute of Technology/Council of Library Resources Roundtable was particularly rewarding and provided new opportunities to share with international colleagues the issues surrounding data-set management. I was

the champion for the creation of the Distributed Data Curation Center (D2C2) at Purdue University (http://d2c2.lib.purdue.edu/)

Throughout my career, beginning with my dissertation, I have been actively involved with assessing and evaluating libraries. In the fall of 1999, I contacted twenty-two academic library directors to determine whether the need was also felt by others. The response was overwhelmingly affirmative. This resulted in a meeting at ALA Midwinter, January 2000. A formal meeting followed at Villanova University in April 2000. As convener, I helped to form the University Libraries Group (ULG), modeled after the Oberlin Group for college libraries. The ULG is made up of university libraries that support diverse wide-ranging programs through doctoral level and have a level of support that places them in the top tier of academic institutions. A few of the member libraries, along with Villanova, are William and Mary, Wake Forest, Lehigh, Carnegie-Mellon, Tufts, Marquette, Miami of Ohio, and Southern Methodist.

In 1994 appointed to the Standards Committee, College Section, Association of College and Research Libraries. During the next six years, the Committee concentrated on changing the focus of the standards from quantitative analysis of input and output factors to emphasis on assessment of the outcome. Culmination of the work was a re-issue of the *Standards for College Libraries* in 2000. The knowledge gained through my work experience enabled me to formulate the changes needed in the standards. This work allowed for close collaboration with accrediting agencies, both professional and regional.

During this same time another focus emerged, the impact of digital resources. Through my work on the JSTOR Statistics Task Force, standards were developed on the collection of use of electronic databases. This Standard was later adopted in 1998 by the International Consortium of Library Consortia (ICOLC).

In 2002, the American Library Association appointed me to serve as the liaison to the Marketing and Management Section of the International Federation of Library Associations (IFLA).

Professional Service: (representative list)

Nominations Committee, Association of Research Libraries (ARL), 2016.

Steering Committee, Scholarly Publishing and Academic Resources Coalition (SPARC), 2016 – 2017.

"Excellence in Library Services," Chair, Review Team, University of Hong Kong, Hong Kong, August 24-27, 2015.

Chair, Management Advisory Board, 2015-2017; Member, Scientific Advisory Board, arXiv, Cornell University, 1/1/2013 – present.

Advisory Board for the Wayne State University School of Library and Information Science, July 2012 – present.

Advisory Board for Microsoft Academic Search, 2012 – 2015. Redmond, WA.

Transforming Research Libraries, a Strategic Direction Steering Committee of the Association of Research Libraries (ARL), 2012-2015.

Science and Technology section, representing ARL, International Federation of Library Associations (IFLA), Chair, 2013 – 2017; Member, 2011 to present.

Member of University of Pretoria, South Africa, Library Review Committee. August 2013.

Co-chair, Local Arrangements Planning Committee for 2013 Conference, Association of College and Research Libraries (ACRL), a division of the American Library Association (ALA).

Association of Research Libraries Leadership & Career Development Program Mentor, 2011-2017.

e-Science Task Force, Association of Research Libraries. July 2006 – present. Chair, October 2011 – October 2012.

Board of Directors, International Association of Technological University Libraries (IATUL). January 2008 – December 2014.

Midwest Collaborative for Library Services (MCLS); Board Member, October 2010 – December 2012.

Chair, Library Directors, Committee on Institutional Cooperation (CIC), July 2010 – June 2012.

Board of Directors, Association of Research Libraries (ARL); October 2008 – October 2011.

Scholarly Communication Steering Committee, Association of Research Libraries (ARL) 2008-2011.

Editorial Board, *College and Research Libraries*, Association of College and Research Libraries, American Library Association. January 2008 – December 2014.

Chair, Organizing Committee for IATUL Conference 2010, June 21-24, 2010, Purdue University, West Lafayette, Indiana/Chicago, Illinois.

Conference Planning Committee for National Conference of the Association of College and Research Libraries, 2009, Seattle, Washington.

Research Committee, Association of College and Research Libraries, ACRL, division of ALA. 2002-2007, chair, 2005-2007.

Association of Research Libraries, Search and Screen Committee, Executive Director. March – January 2008.

Center for Research Libraries, Board of Directors. April 2006 – April 2012.

Academic Libraries of Indiana, Board of Directors, 2004 – present. Vice-president, 2005-2007. President, 2007- 2009.

ALA Representative to the International Federation of Library Associations (IFLA), Marketing and Management (M&M) Section, initial term 2003-2007, re-appointed for second term, 2007-2011.
Invited to represent Research Libraries at the ACRL/3M Wonewok Retreat to assess Marketing of Academic Libraries, October 2002.
Hugh A. Atkinson Award Committee, LAMA Representative, ALA, 2001-2005.
Program Committee, Library Administrators and Management Association (LAMA), a division of ALA. 1996-2001.
ACRL, Standards and Accreditation Committee, a division of ALA. Liaison to RBMS Section of ACRL. 1997-2002.
Elected to the Executive Committee of LAMA, LOMS, a division of the American Library Association, 1998-2000. Nominated as Chair/Elect for 2003 – 2005.
Columbia University Press Advisory Committee. 1996 - 2000.
LITA/LAMA Conference Evaluation Committee, Pittsburgh, Pennsylvania, October 1996.
"New Learning Communities," Coalition for Networked Information, Indianapolis. November 19-21, 1995. Facilitator for invitational, national conference committed to developing collaborative learning and teaching techniques, involving librarians.
Planning Committee-Evaluation. LITA/LAMA 1996 Conference, Pittsburgh. This first conference, to be held jointly between two divisions of ALA, will focus on new technologies within libraries.
Indiana Cooperative Library Services Authority (InCoLSA), elected to Executive Committee, April 1991, served as President in 1993-94. InCoLSA is a statewide network of academic, public, school, and special libraries that supports library cooperation for cataloging, interlibrary loan, collection development and application of new technologies.
Governor’s Conference on Libraries and Information Services. Served on Planning Committee, Academic Libraries Representative, appointed by the Governor to represent academic libraries in Indiana, Chair, Finance Committee, April 1989-July 1991.
Indiana Library Endowment Foundation Board, 1984-92. Charter Member, 1984, President, 1988-1992. 2004-2005.

University Service: (Summary)

Served on search and screen committees for senior positions including chancellor, dean, and directors; most recently I have been asked to serve on the search committee for the provost of Purdue University. At MIT service included the Library Council & appointment to the Administrative Council by President

Vest, 2001-2003 & Member of the Faculty Committee on the Library System. At Purdue appointed by the President to the Search Committee for the Provost, October 2007 to May 2008; member of the Capital Projects Committee, and IT Operational Oversight Committee as senior academic dean, 2008-2014.

Global Council, Global Policy Institute, 2012 – 2016.

Academic Program Excellence and Rankings (APER) project team, 2014.

Representative of the Academic Deans on the Re-engineering Business Operations, Purdue University, 2016 –

Academic Deans Council chaired by Provost – 2004 – 2017.

University Promotion and Tenure Committee – 2006 – 2017.

"Outstanding Team Award, Electronic Reserve Project," served as Chair, recognition awarded by the President of Villanova University to one team who made an outstanding contribution to the operations of the University, selected by a committee of administrators, faculty, and staff. Awarded September 9, 1999.

Nominated for the IUSB Lundquist Award, 1995 & 1996. The Lundquist award is given to faculty who have "exhibited excellence in teaching, scholarly or artistic achievement, and diversified relevant service..."