

Curriculum Vitae

Michael D. Beynon

Department of Computer Science
A.V. Williams Bldg - Room 4166
University of Maryland
College Park, MD 20742 USA

Phone: 301-405-2756
Fax: 301-405-2744
Email: beynon@cs.umd.edu
Url: <http://www.cs.umd.edu/~beynon/>

Current Research Interests

Wide-area Data Intensive Applications: My dissertation addresses some of the problems in making data intensive applications more scalable and efficient in a distributed wide-area environment. Targetted applications include those that explore and process large distributed scientific datasets. Such applications typically perform significant processing on the data and oftentimes produce results smaller than the input dataset. The basic approach is to decompose the application's data-flow style processing chain into a series of discrete tasks. These tasks (or filters) can be then placed for execution on eligible hosts in the network according to some global optimization, such as reduction of traffic over a particular wide-area link or minimum application response time. One difficult issue is deciding how to organize and decompose an algorithm for such a runtime environment, which is still the responsibility of the programmer. My work deals with starting with a decomposed application, and using cost-based models to decide where and how to execute the application tasks to satisfy various execution goals. I have implemented DataCutter, which is a working prototype of the filter and streams application model for C++ and the runtime system that allows for different initial placements of filters. Several application have been written to make use of it, and initial results are very encouraging for the overall approach. We have been participating in the Grid Forum Advanced Programming Models group to advocate this approach, and the need to integrate storage retrieval and processing in a Grid environment.

Education

Ph.D. in Computer Science, University of Maryland, expected Summer 2001. (proposal Feb 3, 2000)

Thesis: *Supporting Resource-Intensive Applications with Remote Data Sources in a Grid Environment*

Advisor: Joel Saltz

Committee: Joel Saltz, Leana Golubchik, Pete Keleher, Ben Bederson

M.S. in Computer Science, 3.93 GPA, University of Maryland, May 1998.

B.S. in Computer Science, (*summa cum laude*) University of New York at Albany, May 1994.

Research experience

6/95 - present **Research Assistant**, University of Maryland

A member of the *High Performance Systems Lab* (HPSL) directed by Joel Saltz, working on various projects.

Wide-area Application Proxies:

I worked on improving the scalability and performance of client-server data-intensive applications over wide-area connections. Motivating applications include the Virtual Microscope, and other applications our group have been working with from the same general class. First, I experimented with a transparent caching proxy server for the Virtual Microscope. Since the benefits depend greatly on the degree of temporal commonality between concurrent clients, I have been considering non-transparent mechanisms to improve the situation. Task chain models of application behavior can be used in concert with online performance prediction models, to allow the processing to be moved from the server, onto the intermediary machine (proxy) and the client machine. For certain processing chains, we believe the prefixes will be similar enough to exploit by caching intermediate results at the intermediary.

The Virtual Microscope:

I am part of the effort to build a software system that emulates a physical microscope. We collaborated with Clinical Pathologists at Johns Hopkins Hospital, and developed a client-server system that delivered high resolution microscopy images of patient pathology slides to a specialized java client. One difficult aspect was structuring the server to handle the large (7-10GB) size of a single slide image data. The parallel server ran on our 16 node diskfull IBM SP-2. This has become an important data-intensive driving application for other projects in our group, and has recently been funded through the HUBS pathology institution collaboration.

Parallel I/O:

Earlier I was involved in the parallelization and characterization of data-intensive applications. The goal of this project is to find out what it takes to achieve good I/O performance for real-life data-intensive tasks on parallel machines suitable for I/O. The long term goal being to provide runtime and compiler support for such needs. We implemented a user-level library, Jovian-2, to perform Parallel I/O on our various parallel machines. The initial study was performed in a client-server fashion, and more work was needed to investigate what is needed to get good performance in a peer-to-peer mode where every node performs both I/O and computation. Following up on this work, I investigated the peer-to-peer question using the Jovian-2 library on our cluster of Digital Alpha SMPs.

Teaching experience

1/99 - 5/99 **Teaching Assistant**, University of Maryland, College Park, MD.

CMSC 417 Computer Networks (Spring 1999), taught by Prof. Leana Golubchik. I handled homework grading and provided solutions, helped design and grade exams, and gave a few class lectures. The main responsibility was in dealing with the large class project where the students designed and implemented the network and transport layer for a connection-oriented ATM-like system.

9/94 - 5/95 **Teaching Assistant** (half time), University of Maryland, College Park, MD.

Led 2 weekly recitation sections for CMSC 104 Fortran Programming (Fall 1994), taught by Prof. Richard Austing. I taught new course material, graded programming assignments, helped design and grade exams.
Graded projects for CMSC 412 Operating Systems (Spring 1995), taught by Prof. A. Udaya Shankar. The class project was a semester long implementation of a threaded operating system under Nachos (Berkeley class project).

9/91 - 12/91 **Teaching Assistant**, University of New York, Albany, NY.

Led 3 weekly recitation sections for CSI 201Y Intro to Computer Science I (Fall 1991), taught by Prof. Dean Arden. I taught new course material, designed and graded homework and programming assignments, helped design and grade the exams. Led review sessions before exams for entire class (approx 250 students).

Industrial experience

1/98 - 8/98 **Website Consultant** (part time), Local companies

I designed and implemented websites for local companies. One involved an online database of price information, and online ordering of products.

1/92 - 7/94 **Research Group Intern** (half time), General Electric Corporate Research and Development, Schenectady, NY.

Member of the Engineous Research Group, that worked on constraint based optimization of engineering codes. My work with the group involved source code control, platform porting, bug fixes, etc. I also designed and implemented the online help system, printed documentation for users and developers, and designed and implemented the security system used for controlling licensing of the Engineous product. I organized and trained several G.E. engineers attending our in-house training sessions.

Honors

- Teaching Excellence Award 1999: awarded to a single nominated teaching assistant.
- Graduate School and Department travel grants to present paper at ICS'99 in Greece.
- DC ACM Master's Grant 1996: awarded to a single applicant in the Washington, D.C. area on basis of merit.
- NSF Graduate Research Fellowship 1995 honorable mention.
- University of Maryland Graduate School Fellowship 1994-1996: awarded to incoming full time graduate students on basis of merit.
- Member of Phi Beta Kappa - National Honor Society 1994.
- Member of Golden Key National Honor Society 1993.

Personal information

Non-curricular interests include: Scuba diving, sky-diving, skiing, motorcycles, securities investing, and my wife Michele.

Publications

Papers at Refereed Conferences

- Michael D. Beynon, Tahsin Kurc, Alan Sussman, Joel Saltz, "Optimizing Execution of Component-based Applications using Group Instances", *In Proceedings of the IEEE International Symposium on Cluster Computing and the Grid*, Brisbane, Australia, May 16-18, 2001.
- Michael D. Beynon, Tahsin Kurc, Alan Sussman, Joel Saltz, "Design of a Framework for Data-Intensive Wide-Area Applications", *In Proceedings of the 9th Heterogeneous Computing Workshop*, Cancun, Mexico, May 1-5, 2000.
- Michael Beynon, Renato Ferreira, Tahsin Kurc, Alan Sussman, Joel Saltz, "DataCutter: Middleware for Filtering Very Large Scientific Datasets on Archival Storage Systems", *In Proceedings of the 2000 Mass Storage Conference*, College Park, MD, March 27-30, 2000.
- Michael D. Beynon, Alan Sussman, Joel Saltz, "Performance Impact of Proxies in Data Intensive Client-Server Applications", *In Proceedings of the 1999 International Conference on Supercomputing*, Rhodes, Greece, June 20-25, 1999.
- Asmara Afework, Michael D. Beynon, Fabian Bustamante, Angelo Demarzo, M.D., Renato Ferreira, Robert Miller, M.D., Mark Silberman, M.D., Joel Saltz, M.D., Ph.D., Alan Sussman, Ph.D., Hubert Tsang, "Digital Dynamic Telepathology - the Virtual Microscope", *In Proceedings of the 1998 AMIA Annual Fall Symposium*, Orlando, Florida, November 7-11, 1998.
- Anurag Acharya, Mustafa Uysal, Robert Bennett, Assaf Mendelson, Michael D. Beynon, Jeffery K. Hollingsworth, Joel Saltz, Alan Sussman, "Tuning the Performance of I/O Intensive Parallel Applications", *4th ACM Workshop on I/O in Parallel and Distributed Systems*, Philadelphia, Pennsylvania, May 27, 1996.

Journal Papers

- Renato Ferreira, Tahsin Kurc, Michael D. Beynon, Chialin Chang, Joel Saltz, Alan Sussman, "Object-relational Queries into Multidimensional Databases with the Active Data Repository", *International Journal of Supercomputer Applications and High Performance Computing (IJSAC)*, 1999.

Tech Reports

- Michael D. Beynon, Tahsin Kurc, Alan Sussman, Joel Saltz, "Design of a Framework for Data-Intensive Wide-Area Application", Tech Report CS-TR-4104, Feb 2000.
- Michael D. Beynon, Alan Sussman, Joel Saltz, "Performance Impact of Proxies in Data Intensive Client-Server Parallel Applications", Tech Report CS-TR-3951 and UMIACS-98-70, Nov 1998.
- Michael D. Beynon, Renato Ferreira, Asmara Afework, Ganti Krishna Mohan, "Performance Evaluation of Client-Server Architectures for Large-Scale Image Processing Applications", Tech Report CS-TR-3970 and UMIACS-TR-98-17, Dec 1998.

Misc Papers

- Michael D. Beynon, "Experiments with Parallel I/O", University of Maryland, Department of Computer Science, Dec 1995, CMSC818K
- Robert Bennett, Michael D. Beynon, "Validation of an Indirect Network Module for PROTEUS", University of Maryland, Department of Computer Science, May 1995, CMSC818J

Computer experience

Languages:

Proficient in: C, C++, Csh, HTML, Web Servers, cgi-bin, ASP, Latex, SQL, Visual Basic, Awk

Know and can use: Java, Pascal, Fortran, Tcl/Tk, Perl, Lisp, Prolog, Modula-2, Assembler (80x86, Vax, 6809), Texinfo, Lex, Yacc

Platforms: Unix (SunOS 4.1.x, Solaris 2.+, Ultrix, Aix, OSF/1, Linux), MS-DOS, Win9x/NT/2000, Vax/Vms

Software Systems: POSIX threads, MPI, RCS, CVS, Embedded-SQL, Stanford SUIF, Wisconsin Minirel

Large class projects

Compilers: (CMSC838T, Spring 1996, Prof. Chau-Wen Tseng)

Developed optimizing compiler passes for use with the Stanford SUIF compiler system. Added local optimizations such as value numbering, and global optimizations utilizing an iterative data flow framework.

Database: (CMSC624, Prof. Michael Franklin, Fall 1995)

Developed a multi-user concurrent relational database based on the Wisconsin Minirel class project. The front end parsed standard SQL (both data definitions and queries), and returned the result while providing ACID semantics. Benchmarking was performed against the Wisconsin Benchmark, a suite of standard SQL queries to evaluate performance of queries. Overall, our project ranked in the top two for speed. Multiuser support included locking the data files and index for concurrent access.

Graphics: (CMSC828M, Fall 1995, Prof. Dave Mount)

Three projects were developed during the course. First was a fractal generator, with the user interface designed from scratch in C++. Second was a three dimensional flight simulator. Last was a ray tracing system for simple geometric objects. The user interface for my first project also appeared in the Fall 1995 issue of *Maryland Windows*, a departmental newsletter.

Database: (CMSC424, Spring 1995, Prof. Michael Franklin)

Developed a complete management system for a fictitious dining club. The front end was written using Motif, and the back end used embedded SQL and the Ingres relational database. The entire cycle from ER diagrams to complete design specification documents were produced.

Compilers: (CMSC430, Fall 1994, Prof. John Gannon)

Developed scanner, parser, and code generator for a subset of a Pascal-like language. Lex, Yacc, and C++ were used for implementation.

Professional activities

Reviewed paper submissions for

- IPDPS'2000, HCW'2000
- ICPP'99
- PDP'97
- ICPP'96

Memberships and positions

- Student member of ACM Computer Society
- Graduate Council Student Lecture Series coordinator 1995-1996
- Panel member for "New Graduate Student" orientation, 1995
- Host for visiting graduate students, 1996

References

Joel Saltz, Professor
Department of Computer Science
University of Maryland
College Park, MD 20742 USA
Phone: 301-405-2669
Email: saltz@cs.umd.edu

Leana Golubchik, Assistant Professor
Department of Computer Science
University of Maryland
College Park, MD 20742 USA
Phone: 301-405-2751
Email: leana@cs.umd.edu

Peter J. Keleher, Assistant Professor
Department of Computer Science
University of Maryland
College Park, MD 20742 USA
Phone: 301-405-0345
Email: keleher@cs.umd.edu

Alan Sussman, Assistant Research Scientist
Department of Computer Science
University of Maryland
College Park, MD 20742 USA
Phone: 301-405-3360
Email: als@cs.umd.edu

Tahsin Kurc, Instructor
Department of Pathology
Informatics Division
The Johns Hopkins Medical Institutions
2024 East Monument Street, Suite 2-613
Baltimore, MD 21205 USA
Phone: 410-502-7666
Email: kurc@cs.umd.edu

Anurag Acharya, Assistant Professor
Department of Computer Science
University of California
2121 Engineering I
Santa Barbara, CA 93106 USA
Phone: 805-893-3411
Email: acha@cs.ucsb.edu

Michael D. Beynon

Last modified: Sun Jun 3 13:20:04 EDT 2001