

Peter A. Dinda
pdinda@cs.cmu.edu
http://www.cs.cmu.edu/~pdinda

School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
412-268-3077 (voice)
412-268-5576 (fax)

6721 Thomas Blvd, Apt. 32
Pittsburgh, PA 15208
412-362-4239 (voice)
412-512-4395 (cell)

Research Interests

Host, LAN, and WAN resource availability measurement, statistical characterization and prediction; application workload measurement, statistical characterization, and prediction; statistical signal processing; adaptive applications; distributed interactive applications; distributed real-time systems; distributed object systems; compilers and run-time systems for parallel and distributed computing.

Education

Carnegie Mellon University

Ph.D. in Computer Science *Expected February, 2000*
Advisor: David R. O'Hallaron
Thesis: *Prediction-based Real-time Advisors*
M.S. in Computer Science *May, 1996*

University of Wisconsin — Madison

B.S. in Electrical and Computer Engineering *May, 1993*
Second major in Computer Science, minor in History of Science
Dean's list, all semesters

Employment

IBM Austin, Advanced Workstations Division

June, 1991 to December, 1991
Engineering co-op student
Performed timing analysis of Rios Single Chip and PowerPC 601 floating point units using ETA
Developed chip-level engineering changes in BDLS
Prototyped transforms for the IDSP hardware compiler in Rexx
Developed design management software for early consolidation

IBM Rochester, Storage Products Division

May, 1990 to September, 1990
Engineering co-op student
Programmed the Intellidex robotic vision system and a glue-dispense-and-apply robot
Assisted in debugging a tension control board for an HSA wire stringing robot
Developed core code for an SPC data acquisition system using a network of shop floor terminals

Madison Academic Computing Center*January, 1989 to August 1993*

Microcomputing consultant

Developed software distribution and update software with campus wide access

Sysoped bulletin board system

Provided public and contract consulting for OS/2, MS-DOS, MS-Windows, and Macintosh

Blue Moon Photography (self-employed)*1987 to 1990*

Weddings, environmental portraiture, and custom black and white printing

Teaching**Teaching assistant for operating systems, Carnegie Mellon University***Spring 1995, Spring 1996*

Course taught by M. Satyanarayanan and E. Elnozahy (2 semesters)

Designed homeworks

Assisted students with programming projects

Lectured on special topics

Co-advisor for undergraduate projects, Carnegie Mellon University

The iWarp Radio: software radio receiver for iWarp parallel computer (2 students)

Fall, 1995

An Extensible HTML Parser

Fall, 1996

A Distributed Object Naming System

Fall, 1996

CMCL Web Site

*Fall, 1996***Publications****Journal Article**

P. Dinda, "The Statistical Properties of Host Load," *Scientific Programming*, 7:3-4, pages 211–229, Winter, 1999.

Refereed Conference and Workshop Publications

P. Dinda, D. O'Hallaron, "An Evaluation of Linear Models for Host Load Prediction," *Proceedings of the 8th IEEE Symposium on High-Performance Distributed Computing (HPDC '99)* (Redondo Beach, California), pages 87–96, August, 1999.

A. Myers, P. Dinda, H. Zhang, "Performance Characteristics of Mirror Servers on the Internet," *Proceedings of Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Infocom 1999)* (New York, New York), pages 304–312, March, 1999.

P. Dinda, B. Lowekamp, L. Kallivokas, D. O'Hallaron, "The Case for Prediction-based Best-effort Real-time Systems," *Proceedings of the 7th International Workshop on Parallel and Distributed Real-time Systems (WPDRTS '99)* (San Juan, Puerto Rico), pages 309–318, March, 1999.

P. Dinda, "The Statistical Properties of Host Load," *Proceedings of the 4th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR '98)* (Pittsburgh, Pennsylvania), pages 319–334, May, 1998.

- P. Dinda, D. O'Hallaron, "Fast Message Assembly Using Compact Address Relations," *Proceedings of the ACM SIGMETRICS '96 Conference on Measurement and Modeling of Computer Systems* (Philadelphia, Pennsylvania), pages 47–56, May, 1996.
- P. Dinda, D. O'Hallaron, J. Subhlok, J. Webb, B. Yang, "Language and Run-time Support for Network Parallel Computing," *Proceedings of the 8th International Workshop on Languages and Compilers for Parallel Computing (LCPC '95)* (Columbus, Ohio), pages 534–550, August, 1995.
- P. Dinda, D. O'Hallaron, "The Performance Impact of Address Relation Caching," *Proceedings of the 3rd Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR '95)* (Troy, New York), pages 213–226, May, 1995.
- J. Subhlok, D. O'Hallaron, T. Gross, P. Dinda, J. Webb, "Communication and Memory Requirements as the Basis for Mapping Task and Data Parallel Programs," *Proceedings of Supercomputing '94* (Washington, DC), pages 330–339, November, 1994.

Unrefereed Conference Publication

- M. Aeschlimann, P. Dinda, L. Kallivokas, J. Lopez, B. Lowekamp, D. O'Hallaron, "Preliminary Report on the Design of a Framework for Distributed Visualization," *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'99)* (Las Vegas, Nevada), pages 1833–1839, June, 1999.

Non-overlapping Technical Reports

- P. Dinda, D. O'Hallaron, "An Extensible Toolkit for Resource Prediction In Distributed Systems," Technical Report CMU-CS-99-138, School of Computer Science, Carnegie Mellon University, July, 1999.
- P. Dinda, B. Garcia, K. Leung, "The Measured Network Traffic of Compiler-Parallelized Programs," Technical Report CMU-CS-98-144, School of computer Science, Carnegie Mellon University, July, 1998.
- P. Dinda, G. Nacula, M. Price, "MacFS: A Portable Macintosh File System Library," Technical Report CMU-CS-98-145, School of computer Science, Carnegie Mellon University, July, 1998.
- P. Dinda, T. Gross, D. O'Hallaron, E. Segall, J. Stichnoth, J. Subhlok, J. Webb, and B. Yang, "The CMU task parallel program suite," Technical Report CMU-CS-94-131, School of Computer Science, Carnegie Mellon University, March, 1994

Presentations

Contributed Presentations

- "An Evaluation of Linear Models for Host Load Prediction," *8th IEEE Symposium on High-Performance Distributed Computing (HPDC '99)*, Redondo Beach, California, August 3, 1999.
- "The Case for Prediction-based Best-effort Real-time Systems," *7th International Workshop on Parallel and Distributed Real-time Systems (WPDRTS '99)*, San Juan, Puerto Rico, April 13, 1999.
- "The Statistical Properties of Host Load," *4th Workshop on Languages, Compilers, and Run-time Systems for Scalable Computers (LCR '98)*, Pittsburgh, Pennsylvania, May 30, 1998.

“Fast Message Assembly Using Compact Address Relations,” *ACM SIGMETRICS '96 Conference on Measurement and Modeling of Computer Systems*, Philadelphia, Pennsylvania, May 24, 1996.

“Language and Run-time Support for Network Parallel Computing,” *8th International Workshop on Languages and Compilers for Parallel Computing (LCPC '95)*, Columbus, Ohio, August 12, 1995.

“The Performance Impact of Address Relation Caching,” *3rd Workshop on Languages, Compilers, and Run-time Systems for Scaleable Computers (LCR '95)*, Troy, New York, May 23, 1995.

Invited Presentations

“Load Prediction for Best-effort Real-time,” *Poster for ARPA Quorum PI Meeting*, July 11, 1998.

“Load Analysis and Prediction for Responsive Interactive Applications,” *ARPA Site visit (Gary Koob)*, April 14, 1998.

“Responsive Interactive Applications by Dynamic Mapping of Activation Trees,” *BBN Technologies*, February 20, 1998.

“Distributed User-centric Applications,” *Intel Fellowship Forum*, Santa Clara, California, May, 1996.

“PVM-based Implementations of Fx and Archimedes,” *PVM Users' Group Meeting*, Pittsburgh, Pennsylvania, May 7, 1995.

Projects and Systems

Prediction-based Real-time Advisors

January, 1998 to present

Real-time advisors simplify application scheduling of resilient real-time tasks in a shared, unreserved computing environment. This is a problem that occurs in the context of distributed interactive applications such as scientific visualization. Using explicit resource-oriented prediction, a real-time advisor can inform the application of the host on which a task will most likely meet its deadline. For my thesis, I developed and evaluated a real-time advisor based on host load prediction. <http://www.cs.cmu.edu/~pdinda/thesis.html>

Resource Prediction System Toolkit

January, 1998 to present

The RPS Toolkit simplifies the development of fast, low overhead on-line prediction systems for resource signals. A resource signal is a time series of values (a discrete-time signal) associated with the availability of an underlying resource. Using RPS, a researcher can study the performance of different predictive models on representative traces of the resource signal of interest, and then quickly construct an on-line prediction system which uses the most appropriate model. RPS-based prediction systems have been incorporated into the CMU Remos resource measurement system and into BBN's QuOin framework for distributed object quality of service. RPS is currently distributed with the Remos system.

<http://www.cs.cmu.edu/~pdinda/RPS.html>

Dv

July, 1998 to present

The Dv project is developing a framework for heavy-weight Internet services, particularly scientific visualizations for large datasets such as those resulting from earthquake simulations. I have contributed to the design of this framework, especially with an eye toward enabling real-time scheduling of Dv tasks (called active frames) using prediction-based real-time advisors.

<http://www.cs.cmu.edu/~dv>

Host load trace playback and public archive*August, 1997 to present*

Research into adaptive applications running on shared, unreserved distributed computing environments requires the ability to construct realistic “background” traffic on hosts and networks. I developed tools for playing back the workloads captured in load traces collected on real machines and created a public archive of such traces.

<http://www.cs.cmu.edu/~pdinda/LoadTraces>

Remos*March, 1997 to present*

The Remos project is developing a system that permits a network-aware application to measure the hosts and network of any heterogeneous distributed computing environment using a common query interface. I have contributed to Remos in several ways. First, I have helped to incorporate RPS-based prediction services so that applications can ask about future resource availability. Second, I am helping to design the next implementation of Remos, including a new query model that incorporates prediction and time series information, and a “tool-oriented” interface for conversing with third-party tools such as those developed by other participants in the Grid Forum. Finally, I am involved in developing new, extensible representations for queries and responses, with the goal of influencing the standards process within the Grid Forum.

<http://www.cs.cmu.edu/~cmcl/remulac/remos.html>

Application-aware Anycast*March, 1998 to April, 1999*

Information such as web documents is often replicated on several servers, which are called mirrors. The application-aware anycast problem is for the client to choose the mirror which can supply the requested document most quickly. This project evaluated the prospects for solving this problem via a large scale, multi-site measurement study of the performance of existing mirrors on the world wide web. My contribution was a detailed statistical analysis of the measurement data that looked for correlations that could be exploited to increase the client’s chances of choosing the appropriate server.

<http://www.cs.cmu.edu/~acm/research/anycast.html>

Lightweight Distributed Object System*May, 1996 to July, 1997*

I built LDOS to simplify the creation of fast distributed C++ objects, and as a research platform for distributed objects. LDOS uses a CORBA IDL compiler to generate code that connects instances of a C++ class to a high-performance ORB-like run-time system. I developed and implemented extensions to IDL, the compiler, and the LDOS run-time system to allow object state to be efficiently shared among multiple instances of the object.

<http://www.cs.cmu.edu/~pdinda/ldo.html>

Address Relation Toolbox*August, 1994 to May, 1996*

Parallel run-time systems often copy data, both between buffers and also across a network. These copy operations are often complex — even an innocuous-appearing High Performance Fortran array assignment statement can result in a complex relation between source and destination addresses that can take a long time to compute. I developed ART to facilitate computing and using such address relations within a run-time system. ART provides tools for computing, compressing, and caching arbitrary address relations (eg, in an inspector), and then reusing these cached compact address relations to perform copies at full memory bandwidth (eg, in an executor). ART was used to implement inter-task communication in the Fx parallelizing Fortran compiler.

<http://www.cs.cmu.edu/~pdinda/codes.html>

Fx*August, 1993 to May, 1996*

The Fx project focused on parallelizing compilers and run-time systems. The Fx compiler supports

a variant of the High Performance Fortran language. I contributed to the Fx project in several ways. I developed a stereo vision application in Fx that could be adapted to use different degrees of task- and data-parallelism according to latency and throughput requirements. This became part of the CMU task-parallel program suite. I ported the Fx compiler and run-time to PVM and the Intel Paragon. I used my Address Relation Toolbox to implement high performance inter-task communication in the Fx run-time and studied its performance. I developed SCANMACS, a set of C macros for creating fast parallel prefix operators on distributed arrays and used it to implement parallel prefix support for Fx. I helped to design NetFx, a variant of Fx to support adaptation on commodity networks.

<http://www.cs.cmu.edu/~fx>

MacFS

August, 1993 to February, 1994

MacFS is a portable library that implements the B*-tree-based Macintosh HFS filesystem. It has been incorporated into several open source and commercial products.

<http://www.cs.cmu.edu/~pdinda/codes.html>

Service

Referee for SIGMETRICS, HPDC, IPDPS, IPPS, ICPP, Supercomputing, SPAA, ISCA, IEEE Network, SIGMM, and others.

Contributor to two ARPA Quorum PI Meetings.

Program committee member for Carnegie Mellon University's Second Symposium On Computer Systems (SOCS2).

Participant in Carnegie Mellon University's Immigration Course Research Symposium, 1994–1996.

Member of ACM, IEEE, and History of Science Society.

Co-founder and secretary of CMU Photographic Society, 1993-1995.

Honors

Intel Foundation Fellowship, 1996-1997

Carnegie Mellon School of Computer Science Fellowship, 1993–1999

George P. Ryan Scholarship, 1988–1992

Alexander DeLorenzo Scholarship, 1988–1989

Villas Scholarship, 1988

Personal

Born February 13, 1970 in Racine, Wisconsin. Speaks passable German. U.S. Citizen.

References

David O'Hallaron

Associate Professor
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891
droh@cs.cmu.edu
(412) 268-8199 (voice)
(412) 268-5576 (fax)

Peter Steenkiste

Senior Research Scientist
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891
prs@cs.cmu.edu
(412) 268-3261 (voice)
(412) 268-5576 (fax)

David Bakken

Assistant Professor
School of Electrical Engineering and Computer Science
PO Box 642752
Washington State University
Pullman, Washington 99164-2752
bakken@eecs.wsu.edu
(509) 335-2399 (voice)
(509) 335-3818 (fax)

Thomas Gross

Associate Professor (CMU)
Professor (ETH Zuerich)
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891
trg@cs.cmu.edu
(412) 268-7661 (voice)
(412) 268-5576 (fax)

Jaspal Subhlok

Associate Professor
Department of Computer Science
University of Houston
Houston, TX 77204
jaspal@cs.uh.edu
(713) 743-3340 (voice)
(713) 743-3335 (fax)