

The Psyche Parallel Operating System

Computer Science Department
University of Rochester

Personnel

Principal investigators	Tom LeBlanc	leblanc@cs.rochester.edu
	Michael Scott	scott@cs.rochester.edu
Students	Brian Marsh	marsh@cs.rochester.edu
	Evangelos Markatos	markatos@cs.rochester.edu
	Cezary Dubnicki	dubnicki@cs.rochester.edu
	Mark Crovella	crovella@cs.rochester.edu
Professional staff	Tim Becker	becker@cs.rochester.edu

Additional contributions by Chris Brown, Sean Colbath, Alan Cox, Prakash Das, Rob Fowler, Neal Gafter, Yenjo Han, Loring Holden, Kurt Jones, Jonas Karlsson, John Kerber, John Leadley, John Mellor-Crummey, Cesar Quiroz, Neil Smithline, Dave Tilley, and Jack Veenstra.

Project Description

Shared-memory multiprocessors provide an extremely flexible hardware base for high-performance parallel computing. They admit the efficient implementation of a wide range of parallel programming models, with widely varying notions of what a process is, how processes communicate, and how they create, use and protect shared data. Programmers need different models for different kinds of programs. Unfortunately, the traditional approach to operating system design presents a virtual machine in which many of the policies and mechanisms on which these models depend are outside the control of the user.

Psyche is a multiprocessor operating system designed to facilitate *multi-model parallel programming*—the simultaneous use of many different styles of parallelism, both across programs and within a single program. Psyche presents a user interface in which processes, communication, sharing, and protection are managed by user-level code. It uses shared data structures to lower the cost of communication between the kernel and the user. It provides a set of software interrupts that can be used to implement *first-class user-level threads*, capable of performing all the functions of a traditional kernel-supported process. It

establishes addressing and protection conventions that encourage the sharing of code and data between applications, and defines interface conventions for user-level schedulers that allow threads of different kinds to synchronize in user space.

An implementation of Psyche for the BBN Butterfly Plus multiprocessor has been in use since early 1989. Kernel implementation studies have focussed on the organization of data structures and division of labor among nodes of a NUMA machine. Application work has focussed on collaboration with departmental researchers in computer vision and robotics.

Contact

Tom LeBlanc or Michael Scott
Computer Science Department
University of Rochester
Rochester, NY 14627-0226
{leblanc,scott}@cs.rochester.edu
FAX (716) 461-2018

Report Availability

Most of the following are available via anonymous ftp (user anonymous, password anything) from cayuga.cs.rochester.edu; directory pub/systems-papers. Hard copy is available for a modest fee from Peg Meeker, Computer Science Department, U of R, etc.

References

- [1] M. L. Scott and T. J. LeBlanc. Psyche: A General-Purpose Operating System for Shared-Memory Multiprocessors. BPR 19, TR 223, Computer Science Department, University of Rochester, July 1987.
- [2] T. J. LeBlanc, M. L. Scott, and C. M. Brown. Large-Scale Parallel Programming: Experience with the BBN Butterfly Parallel Processor. In *Proceedings of the First ACM Conference on Parallel Programming: Experience with Applications, Languages and Systems*, pages 161-172, New Haven, CT, 19-21 July 1988. Also available as BPR 22, Computer Science Department, University of Rochester, September 1988.
- [3] M. L. Scott, T. J. LeBlanc, and B. D. Marsh. Design Rationale for Psyche, a General-Purpose Multiprocessor Operating System. In *Proceedings of the 1988 International Conference on Parallel Processing*, pages II-Software:255-262, 15-19 August 1988. Also published in the *University of*

Rochester 1988-89 Computer Science and Computer Engineering Research Review.

- [4] M. L. Scott, T. J. LeBlanc, and B. D. Marsh. Evolution of an Operating System for Large-Scale Shared-Memory Multiprocessors. TR 309, Computer Science Department, University of Rochester, March 1989.
- [5] T. J. LeBlanc, B. D. Marsh, and M. L. Scott. Memory Management for Large-Scale NUMA Multiprocessors. TR 311, Computer Science Department, University of Rochester, March 1989.
- [6] M. L. Scott, T. J. LeBlanc, B. D. Marsh, T. G. Becker, C. Dubnicki, E. P. Markatos, and N. G. Smithline. Implementation Issues for the Psyche Multiprocessor Operating System. *Computing Systems*, 3(1):101-137, Winter 1990. Earlier version presented at the *First Workshop on Experiences with Building Distributed and Multiprocessor Systems*, Ft. Lauderdale, FL, 5-6 October 1989, 227-236.
- [7] M. L. Scott, T. J. LeBlanc, and B. D. Marsh. A Multi-User, Multi-Language Open Operating System. In *Proceedings of the Second Workshop on Workstation Operating Systems*, pages 125-129, Pacific Grove, CA, 27-29 September 1989.
- [8] M. L. Scott, T. J. LeBlanc, and B. D. Marsh. Multi-Model Parallel Programming in Psyche. In *Proceedings of the Second ACM Symposium on Principles and Practice of Parallel Programming*, pages 70-78, Seattle, WA, 14-16 March 1990. In *ACM SIGPLAN Notices* 25:3.
- [9] E. Chaves, T. J. LeBlanc, B. D. Marsh, and M. L. Scott. Kernel-Kernel Communication in a Shared-Memory Multiprocessor. In *Proceedings of the Second USENIX Symposium on Experiences with Distributed and Multiprocessor Systems*, pages 105-116, Atlanta, GA, 21-22 March 1991. Also published as TR 368, Computer Science Department, University of Rochester, January 1991.
- [10] C. Brown, B. Marsh, T. Becker, M. Crovella, P. Das, J. Karlsson, T. LeBlanc, C. Quiroz, and M. Scott. A Multi-Model Software Architecture for Computer Vision. TR 374, Computer Science Department, University of Rochester, March 1991.
- [11] B. D. Marsh, M. L. Scott, T. J. LeBlanc, and E. P. Markatos. First-Class User-Level Threads. In *Proceedings of the Thirteenth ACM Symposium on Operating Systems Principles*, Pacific Grove, CA, 14-16 October 1991 (to appear).
- [12] B. D. Marsh. Multi-Model Parallel Programming. Ph.D. Thesis. Computer Science Department, University of Rochester, August 1991 (to appear).