

# Multiprocessing: An Annotated Bibliography

M. Satyanarayanan  
Carnegie-Mellon University

The years since the advent of multiprocessing have witnessed the proliferation of journal articles, conferences, and symposia on multiprocessing systems and closely related topics. The body of literature so generated is both vast and scattered—a situation that can be quite disheartening to one unfamiliar with this area of computer science.

In this guide to multiprocessing literature, I have divided the subject into seven major subareas and classified each paper into one of them. Of course, the decision to place a paper in one subarea rather than another was not always easy. In marginal cases, I assigned the paper to the subarea I felt most appropriate. My apologies, in advance, to anyone who feels his or her paper has been misclassified.

In all but a handful of cases, I have added a brief annotation; the exceptions are primarily those papers whose titles are sufficiently suggestive of their contents. Wherever possible, I have included page numbers. Articles and papers already cited in the preceding survey are not repeated here.

## Multiprocessor architectures and operating systems

These articles and papers describe the philosophy and design of multiprocessors and their operating systems. Almost all deal with experimental systems built as research projects in universities or in industry.

1. Anderson, G. L. and K. Bartlett, "Hardware Allocation of Data System Resources," *Computer Design*, Vol. 13, No. 7, July 1974, pp. 89-97.

Describes a multiprocessor architecture in which a set of microprocessors are initially treated as a pool of resources and dynamically assigned specific hardware functions as the need arises. The concept of functional partitioning exists here, but the assignment of functions to processors is on a temporary rather than permanent basis.

2. Arden, B. W. and A. D. Berenbaum, "A Multi-Microprocessor Computer System Architecture," *Proc. 5th Symp. Operating Systems Principles* (published as *ACM Operating Systems Review*, Vol. 9, No. 5), Austin, Tex., Nov. 1975, pp. 114-121.

Describes the design of a multiprocessor built with microprocessors and discusses the issues involved in building a distributed operating system for such a machine. The assignment of system processes to processors and the hierarchical connection scheme used to minimize interconnection requirements are of special interest.

3. Arnold, R. G. and E. W. Page, "A Hierarchical, Restructurable Multi-Microprocessor Architecture," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, pp. 40-45.\*

This multiprocessor system was built with bit-slice processors. It allows dynamic restructuring of the system to permit longer word lengths as well as reorganization into isolated groups of cooperating processors.

4. Barnwell, T. P., S. Gaglio, and R. M. Price, "A Multi-Microprocessor Architecture for Digital Signal Processing," *Proc. 1978 Int'l Conf. Parallel Processing*, Bellaire, Mich., Aug. 1978, pp. 115-121.\*

Describes the architecture of a multiprocessor dedicated to signal processing tasks and discusses techniques for generating efficient code for certain algorithms used in digital signal processing.

5. Baum, A. and D. Senzig, "Hardware Considerations in a Microcomputer Multiprocessor System," *Computer Technology to Reach the People—Digest of Papers—COMPCON Spring 75*, San Francisco, Calif., Feb. 1975, pp. 27-30.\*

A multiprocessor architecture suitable for LSI implementation is described.

6. Crick, A., "Scheduling and Controlling I/O Operations," *Data Processing*, Vol. 16, No. 3, May-June 1974, pp. 170-171.

A multiprocessor real-time operating system for the Univac 1110 permits every I/O device to communicate with

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

every processor in the system. Interrupt latency, reliability, error recoverability, and other related aspects of this system are discussed.

7. Davis, A. L., "The Architecture and System Method of DDM1: A Recursively Structured Data Driven Machine," *Conf. Proc. 5th Ann. Symp. Computer Architecture*, Palo Alto, Calif., Apr. 1978, pp. 210-215.\*  
This architecture presents a facade of unlimited concurrency and pipelining to programs, expressed in a formalism called *data driven nets*. Multiprocessing is used to approximate this unlimited parallelism, with the assignment of processes to available processors being done by the hardware.
8. de la Guardia, M. F. and J. A. Field, "A High Level Language Oriented Multiprocessor," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, pp. 256-262.\*
9. Ellis, C. A. and G. J. Nutt, "Preliminary Thoughts on Degrees of Security in Multiprocessor Systems," Tech. Report CU-CS-036-74, University of Colorado, Boulder, Colo., June 1974.  
Discusses protection and security issues in a multiprocessing environment and examines the impact of multiprocessing on problems such as deadlocks, scheduling, and equitable resource allocation.
10. England, D. M., "Architectural Features of the System 250," *Proc. Symp. Operational On-Line Computing for Defence*, Malvern, England, Nov. 1972.  
Describes the architecture and software of a multiprocessor system designed for reliable, real-time operation.
11. Filene, R. J. and A. I. Green, "A Simple Executive for a Fault-Tolerant Real-Time Multiprocessor," *Hardware, Software, Firmware and Tradeoffs—Digest of Papers—COMPCON Fall 71*, Boston, Mass., Sept. 1971.  
An operating system intended for a fault-tolerant multiprocessor in a real-time environment is described. Architectural features to support such an operating system are discussed.
12. Fuller, S. H. et al., "Multi-Microprocessors: An Overview and Working Example," *Proc. IEEE*, Vol. 66, No. 2, Feb. 1978, pp. 216-228.  
Discusses the problems and advantages of using microprocessors for building multiprocessors. Describes Cm\* from this viewpoint and gives some performance measurements.
13. Gilliland, M. C., B. J. Smith, and W. Calvert, "HEP: A Semaphore-Synchronized Multiprocessor with Central Control," *Proc. 1976 Summer Computer Simulation Conf.*, Washington, D.C., July 1976, pp. 57-62.  
HEP—heterogeneous element processor—is a multiprocessor designed specifically for high performance in simulation applications.
14. Grimsdale, R. L. and D. M. Johnson, "A Modular Executive for Multiprocessor Systems," *Proc. Conf. Trends in On-Line Computer Control Systems*, Sheffield, England, Apr. 1972.  
An operating system for a symmetric multiprocessor is described. The reliability and error recovery features of this system are emphasized.
15. Grushcow, M. S., "The Kernel of the SUE Operating System," *Proc. Canadian Computer Conf.*, Montreal, June 1972.  
SUE is an experimental multiprocessor operating system based on the notion of virtual parallel processors—each process considers itself to be running on its own, private processor. The multiplexing of the actual hardware between processes, the initiation and termination of processes, and low-level operating system functions such as I/O and timing are performed by the kernel.
16. Gula, J. L., "Operating System Considerations for Multiprocessor Architectures," *Proc. Seventh Texas Conf. Computing Systems*, Houston, Tex., Nov. 1978, pp. 7-1 to 7-6.\*
17. Holley, L. H. et al., "VM/370 Asymmetric Multiprocessing," *IBM Systems J.*, Vol. 18, No. 1, 1979, pp. 47-70.  
Describes the use of attached processing in IBM's Virtual Machine Operating System. Performance goals are discussed and compared with actual performance.
18. Jones, A. K. et al., "Software Management of Cm\*—A Distributed Multiprocessor," *AFIPS Conf. Proc.*, Vol. 46, 1977 NCC, pp. 657-663.  
Contains an exposition of the software issues involved in building an operating system for Cm\*. StarOS, one of the two operating systems for Cm\*, evolved from the ideas and strategies developed here.
19. Kamiuchi, T. and H. Nakanishi, "H-80—A Loadsharing, N:1 Backup Multisystem," *Computer Technology: Status, Limits, Alternatives—Digest of Papers—COMPCON Spring 78*, San Francisco, Calif., Mar. 1978, pp. 261-264.\*  
Intended for on-line control applications, this system uses multiprocessing for improving throughput and enhancing reliability and availability.
20. Kober, R., "The Multiprocessor System SMS 201—Combining 128 Microprocessors to a Powerful Computer," *Micros, Minis, & Maxis, Technology Thrust vs. User Requirement—Digest of Papers—COMPCON Fall 77*, Washington, D.C., Sept. 1977, pp. 225-230.\*  
A multi-microprocessor system designed by Siemens is described.
21. Kralej, M. F., "The Pluribus Multiprocessor," *Digest of Papers—1975 Int'l Symp. Fault-Tolerant Computing*, Paris, France, June 1975, p. 251.\*  
This fault-tolerant multiprocessor is designed primarily for message routing and buffering in the Arpanet. It uses multiprocessing for reliability as well as to allow modular performance growth.
22. Kuznia, C. H., R. Kober, and H. Kopp, "SMS 101—A Structured Multimicroprocessor System with Deadlock-Free Operation Scheme," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, p. 122.\*  
A multiprocessor design using microprocessors is described. A distinctive feature of this system is the use of simultaneously updatable private memories instead of a shared global memory.
23. Murakami, K., S. Nishikawa, and M. Sato, "Poly-Processor System Analysis and Design," *Conf. Proc. 4th Ann. Symp. Computer Architecture*, Silver Spring, Md., Mar. 1977, pp. 49-56.\*  
Describes a multiprocessor system built with microprocessors dedicated to performing specialized functions. Discusses the functional partitioning of the system and presents simulation results predicting its performance.
24. Newton, R. S., "The Design of a Reentrant Executive for a Multi-CPU Shared-Store Operating System," *Proc. Symp. Operational On-Line Computing for Defence*, Malvern, England, Nov. 1972.  
Examines the minimal hardware and software support needed for multiprocessing and describes those aspects of the software which differ from the uniprocessor case. Also suggests programming techniques for improving reliability.
25. Newton, R. S., "An Exercise in Multiprocessor Operating System Design," *AGARD Conf. Proc.*, No. 149, Real-Time Computer-Based Systems (NATO Advisory Group for Aerospace R&D), Athens, Greece, May 1974.

Describes the design of general-purpose operating systems for multiprocessors, emphasizing those aspects of the design directly related to multiprocessing.

26. Nutt, G. J., "A Parallel Processor Operating System Comparison," *IEEE Trans. Software Eng.*, Vol. SE-3, No. 6, Nov. 1977, pp. 467-475.

Compares three different operating system designs for a multiprocessor and evaluates them for different system loads.

27. O'Grady, E. P., "A Multiprocessor for Continuous System Simulation," *Proc. 1979 Int'l Conf. Parallel Processing*, Bellaire, Mich., Aug. 1979, p. 306.\*

This multiprocessor system is designed to improve speed in simulation applications. The main point of interest here is the technique used for the efficient exchange of data between processors.

28. Ousterhout, J. et al., "Medusa: An Experiment in Distributed Operating System Structure," *Proc. 7th Ann. Symp. Operating Systems Principles*, Nov. 1979, pp. 115-116.

Discusses the issues involved in designing an operating system for the Cm\* multiprocessor. Describes the key features of Medusa and justifies its design in the context of Cm\*.

29. Ousterhout, J., "Partitioning and Communication in a Distributed Operating System," doctoral thesis, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., to appear, 1980.

Detailed description of Medusa, an operating system for the multiprocessor, Cm\*.

30. Pollard, L. H., "Multiprocessing with the TI 9900," *Conf. Record—11th Asilomar Conf. Circuits, Systems and Computers*, Pacific Grove, Calif., Nov. 1977, pp.461-465.\*

The design of a multiprocessor intended for data acquisition and processing is described.

31. Quaynor, N. and A. Bernstein, "Operating Systems for Hierarchical Multiprocessors," *Proc. Seventh Texas Conf. Computing Systems*, Houston, Tex., Nov. 1978, pp. 1-9 to 1-15.\*


Discusses the decomposition of operating systems for multi-microprocessors. Also describes a specific architecture supporting hierarchical software and discusses the mapping of a typical operating system onto this architecture.

32. Radoy, C. H. and G. J. Lipovski, "Switched Multiple Instruction, Multiple Data Stream Processing," *Conf. Proc. 2nd Ann. Symp. Computer Architecture*, Houston, Tex., Jan. 1975, pp. 183-187.\*

A multicomputer architecture combines the characteristics of multiprocessors and array processors. As in a multiprocessor, each processing unit is a computer capable of executing programs on its own. As in an array processor, however, each such processing unit periodically receives orders to compute specific program segments. This architecture may be viewed as a multiprocessor architecture with very rigid, hardware-enforced scheduling. The paper also discusses the criteria for determining the appropriate use of this architecture.

33. Rumbaugh, J., "A Data Flow Multiprocessor," *IEEE Trans. Computers*, Vol. C-26, No. 2, Feb. 1977, pp. 138-146.

A multiprocessor based on a data flow architecture is described and its characteristics explained. The primary advantages of such an interconnection scheme are the availability of a high degree of concurrency and the ability to execute programs described in data flow notation—a notation permitting the expression of parallelism and precedence constraints.



34. Siegel, H. J., P. T. Mueller, and H. E. Smalley, "Control of a Partitionable Multimicroprocessor System," *Proc. 1978 Int'l Conf. Parallel Processing*, Bellaire, Mich., Aug. 1978, pp. 9-17.\*

Describes a highly flexible architecture which may be dynamically configured as a multiprocessor or as an array processor.

35. Siewiorek, D., M. Canepa, and S. Clark, "C.vmp: The Architecture and Implementation of a Fault Tolerant Multiprocessor," *Proc. Seventh Ann. Int'l Conf. Fault-Tolerant Computing*, Los Angeles, Calif., June 1977, pp. 37-43.\*

This multiprocessor uses multiple processing units to improve reliability without software modification. The hardware uses bus-level voting on signals to detect and recover from failures.

36. Smith, B. J., "A Pipelined, Shared Resource MIMD Computer," *Proc. 1978 Int'l Conf. Parallel Processing*, Bellaire, Mich., Aug. 1978, pp. 6-8.\*

37. Smith, T. B. and A. L. Hopkins, "Architectural Description of a Fault-Tolerant Multiprocessor Engineering Prototype," *Digest of Papers—Eighth Ann. Int'l Conf. Fault-Tolerant Computing*, Toulouse, France, June 1978, p. 194.\*

Describes a multiprocessor system, for use in avionic systems, with high reliability characteristics.

38. Swan, R. J., "The Switching and Addressing Structure of an Extensible Multiprocessor: Cm\*," doctoral thesis, Dept. of Computer Science, Carnegie-Mellon University, Aug. 1978.

Gives a fairly detailed description of the Cm\* multiprocessor and its hardware implementation. Discusses in detail the problem of deadlock in Cm\*-like interconnection schemes and explains the strategies used to avoid it in Cm\*.

39. Taylor, J. R., "A Multiprocessor Configuration for High-Reliability Processing," *Proc. 1st European Seminar Computing with Real-Time Systems*, Harwell, England, 1972.

A multiprocessor configuration achieves high reliability using commercially available computers.

40. Walker, L. L., "Multiprocessor Operating System Design," in *Operating Systems, International Computer State of the Art Report*, Infotech Ltd., Maidenhead, England, 1972.

Identifies issues in, and derives guidelines for, multiprocessor operating system design.

41. Ward, S. A., "The MuNet: A Multiprocessor Message-Passing System Architecture," *Proc. Seventh Texas Conf. Computing Systems*, Houston, Tex., Nov. 1978, pp. 7-21 to 7-24.\*

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

Describes a multiprocessor whose performance may be varied widely without software modification, thus permitting the hardware designer to defer certain design decisions.

42. Wecker, S., "A Design for a Multiple Processor Operating Environment," *Computing Networks from Minis through Maxis—Are They for Real?—Digest of Papers—COMP-CON 73*, San Francisco, Calif., Mar. 1973, pp. 143-146.\*

Interprocess communication in multiprocessor systems and computer networks, the software analog of interprocessor communication, is discussed. Also presents a unified communication structure and techniques for using this structure to interface operating systems in multiple processor environments.

43. Widdoes, L. C., "The Minerva Multi-Microprocessor," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, pp. 34-39.\*

This multiprocessor consists of microprocessors connected by a shared bus. Discusses interrupt handling and techniques for reducing processor contention for the shared bus.

44. Wulf, W. A. and C. G. Bell, "C.mmp—A Multi-Mini-Processor," *AFIPS Conf. Proc.*, Vol. 41, Part II, 1972 FJCC, pp. 765-777.

This paper describes the original design of C.mmp and discusses issues such as its viability and the feasibility of building it using a minimum of custom-built hardware. The final C.mmp design differs significantly from the design described here, but the paper is worth reading for the picture of the designers' original motives.

45. Wulf, W. et al., "HYDRA: The Kernel of a Multiprocessor Operating System," *Comm. ACM*, Vol. 17, No. 6, June 1974, pp. 337-345.

Describes the original design of the operating system for C.mmp. The hardware structure assumed is the one described in Wulf and Bell's first C.mmp paper (44).

46. Wulf, W. A., S. P. Harbison, and R. Levin, *Hydra: An Experimental Operating System*, to be published by McGraw-Hill Inc., 1980.

A detailed description of the philosophy, design, and implementation of Hydra, similar in a sense to Organick's monograph on Multics. Highly recommended to anyone desiring an understanding of multiprocessor operating systems in general and Hydra in particular.

## Performance

The performance of multiprocessor systems has been a topic of considerable interest to both researchers and users. Some of the papers in this section attempt to define meaningful measures of performance for multiprocessor systems. Others use models based on analytical techniques such as queueing theory to predict performance. Yet others estimate performance by employing simulation and actual measurements. That multiprocessors share memory implies, almost by definition, that there will be conflicts between processors trying to access the same memory modules. Many studies have investigated the magnitude of this problem and have provided techniques to minimize the performance degradation caused by it. Most of the works cited here address the problem of conflicts, either directly or indirectly.

47. Baskett, F. and A. J. Smith, "Interference in Multiprocessor Computer Systems with Interleaved Memory," *Comm. ACM*, Vol. 19, No. 6, June 1976, pp. 327-334.

Investigates the effects of memory interference in multiprocessor systems and derives an analytical model to describe this. Simulation results validate this model.

48. Bhandarkar, D. P., "Some Performance Issues in Multiprocessor System Design," *IEEE Trans. Computers*, Vol. C-26, No. 5, May 1977, pp. 506-511.

Uses analytic and simulation models to study the effect of alternative multiprocessor designs on performance. Also presents guidelines for designing multiprocessor systems.

49. Chandy, K. M. and M. Reiser, eds., *Computer Performance*, North-Holland Publishers, Amsterdam, Netherlands, 1977.

Contains articles on the performance of multiprocessors and computer networks.

50. Chow, Yuan-Chieh and W. H. Kohler, "Performance of Several Queueing Models for Multiprocessor Multiprogramming Systems," *Computers... by the Millions, for the Millions—Digest of Papers—COMP-CON Fall 76*, Washington, D. C., Sept. 1976, pp. 66-71.\*

The performances of a single large processor and a number of cooperating processors are compared using three different models. Indicates the conditions, when using multiple processors, yielding a small performance degradation.

51. Covo, A. A., "Analysis of Multiprocessor Control Organizations with Partial Program Memory Replication," *IEEE Trans. Computers*, Vol. C-23, No. 2, Feb. 1974, pp. 113-120.

Replicating copies of a program run by different processors in a multiprocessing situation helps reduce memory interference and queueing delays, at the expense of increased memory requirements. Describes a dynamic programming solution which finds an optimal degree of replication.

52. Dal Cin, M., "Performance Evaluation of Self-Diagnosing, Multiprocessing Systems," *Digest of Papers—Eighth Ann. Int'l Conf. Fault-Tolerant Computing*, Toulouse, France, June 1978, pp. 59-64.\*

Uses a graph-theoretic model and queueing analysis to investigate the performance of self-diagnosing multiprocessor systems.

53. Ferrari, D., E. Gelenbe, and R. Mahl, "An Analytic Study of Memory Allocation in Multiprocessor Systems," *Proc. Conf. Computer Architecture and Networks*, Rocquencourt, France, Aug. 1974.

Discusses analytical techniques for optimally partitioning main memory among processes running in a multiprocessor environment and outlines the issues involved in applying these ideas to a practical situation.

54. Franta, W. R. and P. A. Houle, "Comments on Models of Multiprocessor Multi-Memory Bank Computer Systems," *Proc. 1974 Winter Simulation Conf.*, Vol. I, Washington, D.C., Jan. 1974, pp. 86-97.

Contains a discussion of the hardware and software factors contributing to memory interference in multiprocessor systems. Describes a simulation model of memory interference and compares it to other models.

55. Fuller, S. H., "Price/Performance Comparison of C.mmp and the PDP-10," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, pp. 195-202.\*

Compares the price/performance trade-offs involved in choosing between a large mainframe computer (the PDP-10) and a multiprocessor built with minicomputers (the C.mmp). Yields valuable insights into the economic and technical issues involved in such a choice.

56. Hoogendoorn, C. H., "A General Model for Memory Interference in Multiprocessors," *IEEE Trans. Computers*, Vol. C-26, No. 10, Oct. 1977, pp. 998-1005.

Presents a mathematical model of memory interference in

multiprocessor systems and validates it with simulation results and by comparison with models described in the literature.

57. Hoogendoorn, C. H., "Reduction of Memory Interference in Multiprocessor Systems," *Conf. Proc. 4th Ann. Symp. Computer Architecture*, Silver Spring, Md., Mar. 1977, pp. 179-183.\*

This paper describes various techniques for reducing memory interference in multiprocessor systems and evaluates them using a trace-driven simulation model. The techniques considered fall into two classes—memory placement schemes and private memory schemes.

58. Ishida, H., S. Nomoto, and H. Ozawa, "Graphic Monitoring of the Performance of a Large 4-CPU Multiprocessor System," *2nd USA-Japan Computer Conf. Proc.*, Tokyo, Japan, Aug. 1975, pp. 271-275.

Describes a software system for collecting and displaying performance data for a large multiprocessor system. Presents performance figures for this system and discusses the use of these results to improve the design of its operating system.

59. Jensen, J. E. and J. L. Baer, "A Model of Interference in a Shared Resource Multiprocessor," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, pp. 52-57.\*

Considers contention for resources in multiprocessor systems. A noteworthy aspect of this paper is that it considers general shared resources and not just contention for memory.

60. Jordan, H. F., M. Scalabrin, and W. Calvert, "A Comparison of Three Types of Multiprocessor Algorithms," *Proc. 1979 Int'l Conf. Parallel Processing*, Bellaire, Mich., Aug. 1979, pp. 231-238.\*

Evaluates the performance, on a multiprocessor of specified architecture, of three different algorithms for a specific problem. Also discusses hardware extensions to improve performance.

61. Kinney, L. L. and R. G. Arnold, "Analysis of a Multiprocessor System with a Shared Bus," *Conf. Proc. 5th Ann. Symp. Computer Architecture*, Palo Alto, Calif., Apr. 1978, pp. 89-95.\*

Investigates the performance, on a set of tasks having little interaction, of a multiprocessor system with a shared bus interconnection scheme.

62. Kurinckx, A. and G. Pujolle, "Analytic Methods for Multiprocessor Modelling," *Conf. Pre-Prints 4th Int'l Symp. Modelling and Performance Evaluation of Computer Systems*, Part II, Vienna, Austria, Feb. 1979.


Reviews mathematical results useful for modeling multiprocessor systems. Presents an approximate probabilistic model of such systems and describes its use in evaluating multiprocessor performance.

63. Kurtzberg, J. M., "On the Memory Conflict Problem in Multiprocessor Systems," *IEEE Trans. Computers*, Vol. C-23, No. 3, Mar. 1974, pp. 286-293.

Considers the problem of memory interference in multiprocessor systems where each processor runs a different job. Presents analytical models for such situations and develops algorithms for assign jobs to memory modules so as to minimize memory interference.

64. Liu, J. W. S., C. L. Liu, E. Gelenbe, and R. Mahl, "Performance Analysis of Heterogeneous Multiprocessor Computer Systems," *Proc. Conf. Computer Architecture and Networks*, Rocquencourt, France, Aug. 1974.

Discusses criteria for comparing the performances of different multiprocessor systems and provides information on resource utilization in a multiprocessor environment.



65. Marathe, M. V., "Performance Evaluation at the Hardware Architecture Level and the Operating System Kernel Design Level," doctoral thesis, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., Dec. 1977.

One of the early experimental investigations of the hardware and operating system performance of C.mmp. Gives measured data on the observed instruction mix in the operating system and on overheads due to synchronization. Develops a model which describes the performance speedup due to multiprocessing.

66. Mason, P. H., "Design Tools for Evaluating Multiprocessor Programs," doctoral dissertation, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., July 1976.

Describes the design and implementation of a performance evaluation system called STEPPS; this system evaluates multiprocessor programs along different dimensions.

67. Nessett, D. M., "The Effectiveness of Cache Memories in a Multiprocessor Environment," *Australian Computer J.*, Vol. 7, No. 1, Mar. 1975, pp. 33-38.

A simulation model studies the effectiveness of cache memories in reducing memory access conflicts.

68. Pearce, R. C. and J. C. Majithia, "Upper Bounds on the Performance of Some Processor-Memory Interconnections," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, p. 303.\*

Derives the best possible performance for each of a variety of multiprocessor interconnection schemes, such as a crossbar switch and a timeshared bus.

69. Pearce, R. C. and J. C. Majithia, "Performance Results for an m.i.m.d. Computer Organisation Using Pipelined Binary Switches and Cache Memories," *Proc. Inst. Electronic Engineers* (England), Vol. 125, No. 11, Nov. 1978, pp. 1203-1207.

Presents simulation results on the use of cache memories in multiprocessors and on the effect of cache parameters on performance. A low-cost alternative to a crossbar switch, called a pipelined binary switch, is investigated for its suitability as an interconnection mechanism.

70. Pearce, R. C. and J. C. Majithia, "Analysis of a Shared Resource MIMD Computer Organization," *IEEE Trans. Computers*, Vol. C-27, No. 1, Jan. 1978, pp. 64-67.

Analyzes the performance of a multiprocessor system in which one resource is shared.

71. Raskin, L., "Performance Evaluation of Multiple Processor Systems," doctoral thesis, Dept. of Computer Science,

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

Carnegie-Mellon University, Pittsburgh, Pa., Aug. 1978.

The major published source of information on Cm\*'s performance. Describes experiments done on an early, one-cluster version of Cm\* as well as analytic models for larger configurations.

72. Raynor, R. J. and J. M. Gwynn, "Minimization of Supervisor Conflict for Multiprocessor Computer Systems," *Proc. Symp. Simulation of Computer Systems*, Boulder, Colo., Aug. 1976.
73. Sastry, K. V. and R. Y. Kain, "On the Performance of Certain Multiprocessor Computer Organizations," *IEEE Trans. Computers*, Vol. C-24, No. 11, Nov. 1975, pp. 1066-1074.

Develops analytic models for evaluating the performance of a multiprocessor computer system. Studies how varying the number of processors, interleaving memory, and varying the number of memory modules affects the instruction execution rate.
74. Sethi, A. S. and N. Deo, "Interference in Multiprocessor Systems with Localized Memory Access Probabilities," *IEEE Trans. Computers*, Vol. C-28, No. 2, Feb. 1979, pp. 157-163.

Using more realistic assumptions than earlier work, considers the problem of memory interference. Develops a model of the memory reference pattern of typical programs and uses it to analyze the performance of a multiprocessor system. Also presents simulation data validating these analytical results.
75. Smith, A. J., "Multiprocessor Memory Organization and Memory Interference," *Comm. ACM*, Vol. 20, No. 10, Oct. 1977, pp. 754-761.

Examines the effect of various memory organizations on interference. Demonstrates that localizing each processor's memory references to one or more memory modules results in lowered memory interference. These results are derived analytically and verified by simulation.
76. Towsley, D. F., "The Effects of CPU: I/O Overlap on Computer System Configurations," *Conf. Proc. 5th Ann. Symp. Computer Architecture*, Palo Alto, Calif., Apr. 1978, pp. 238-241.\*

Presents a model describing the overlap of processing and I/O in multiprocessor systems. The model allows an evaluation of the effectiveness of a number of multiprocessor designs.
77. Andrews, G. J. and J. R. McGraw, "Language Features for Parallel Processing and Resource Control," *Proc. Conf. Design and Implementation of Programming Languages*, Ithaca, N.Y., Oct. 1976.

Describes requirements that a language for parallel processing should meet. A set of language features for describing processes, and the interactions between them, is presented and evaluated in light of these requirements. Also presents language features to control the use of shared resources in a parallel processing environment.
78. Chen, Shyh-Ching, "Speedup of Interactive Programs in Multiprocessing Systems," Tech. Report UIUCDCS-R-75-694, Dept. of Computer Science, University of Illinois, Urbana, 1975.

Discusses certain aspects of recognizing parallelism in sequential programs. Presents techniques for exposing vector operations in loop computations, for use in multiprocessors as well as pipelined systems. Derives theoretical results for the time and processor bounds of these techniques and for the maximum speedup achievable for different loop computations.
79. Cohen, T., "Structured Flowcharts for Multiprocessing," *Computer Languages*, Vol. 3, No. 4, 1978, pp. 209-226.

Describes flow chart constructs for parallel processing primitives and discusses the transformation of programs using these primitives into equivalent structured programs.
80. Dowsing, R. D., "Processor Management in a Multiprocessor System," *Electronic Letters*, Vol. 12, No. 24, Nov. 1976.

Compares memory management in uniprocessor systems with processor management in multiprocessor systems and shows that many results apply to both cases.
81. Flon, L. and N. Suzuki, "Nondeterminism and the Correctness of Parallel Programs," *Proc. IFIP Working Conf. Formal Descriptions of Programming Concepts*, St. Andrews, N.B., Canada, Aug. 1977.

Describes a technique for transforming parallel programs into equivalent nondeterministic programs and then proving the former by proving the latter. Presents the proofs of some commonly encountered parallel programs.
82. Friedman, D. P. and D. S. Wise, "The Impact of Applicative Programming on Multiprocessing," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, pp. 263-272.\*
83. Gonzalez, M. J. and C. V. Ramamoorthy, "Recognition and Representation of Parallel Processable Streams in Computer Programs," in *Parallel Processor Systems, Technologies and Applications*, Macmillan Ltd., London, England, 1970.

Two techniques for exploiting parallelism in programs are developed here. One tries to decompose a sequential program into tasks that may be executed in parallel; the other tries to recognize parallelism at a finer grain, within one task.
84. Gonzalez, M. J. and C. V. Ramamoorthy, "Parallel Task Execution in a Decentralized System," *IEEE Trans. Computers*, Vol. C-21, No. 12, Dec. 1972, pp. 1310-1322.

Develops a technique for representing program segments that may be executed in parallel and describes the use of this technique in multiprocessor systems. The overheads involved in executing tasks in parallel are investigated in both centralized and decentralized control environments.
85. Graham, R. L., "Bounds on Multiprocessing Anomalies and Related Packing Algorithms," *AFIPS Conf. Proc.*, Vol. 40, 1972 SJCC, pp. 205-217.

Surveys theoretical results in multiprocessing and discusses algorithms for improving multiprocessor performance.
86. Gries, D., "An Exercise in Proving Parallel Programs Correct," *Comm. ACM*, Vol. 20, No. 12, Dec. 1977, pp. 921-930.

## Theoretical results

One goal of multiprocessing is to speed up program execution by exploiting parallelism. Some of the papers included in this section present results on the recognition of parallelism in sequential programs. Others deal with issues such as synchronization of cooperating processes and equitable sharing of resources between competing processes. The verification of parallel programs forms yet another dimension of research in this area. The phrase "multiprocessing," in the context of this section, is often used as a synonym for "parallel processing," rather than in the more restricted sense used in the foregoing sections. Since multiprocessors are just a particular type of parallel processing system, the results reported in these papers apply to multiprocessor systems too.

Provides a formal proof of a parallel algorithm and discusses the problems in generating and understanding such proofs.

87. Kotov, V. E., A. Ershov, and V. A. Nepomniaschy, "Toward Automatic Construction of Parallel Programs," *Proc. Int'l Symp. Theoretical Programming*, Novosibirsk, U.S.S.R., 1974.

Discusses issues in recognizing parallelism in sequential programs. Presents models of sequential and parallel programs for use in high-level language constructs.

88. Lamport, L., "The Synchronization of Independent Processes," *Acta Informatica*, Vol. 7, No. 1, 1976, pp. 15-34.

Considers the problem of building a robust system using failure-prone processes. Defines a synchronization mechanism to accomplish this.

89. Lamport, L., "Proving the Correctness of Multiprocess Programs," *IEEE Trans. Software Eng.*, Vol. SE-3, No. 2, Mar. 1977, pp. 125-143.

Extends the inductive assertion technique of proving sequential programs to programs employing multiple processes.

90. Preparata, F. P., "Parallelism in Sorting," *Proc. 1977 Int'l Conf. Parallel Processing*, Detroit, Mich., Aug. 1977, pp. 202-206.\*

Describes a family of sorting algorithms for use in multiprocessor systems.

91. Ramamoorthy, C. V. and W. H. Leung, "A Scheme for the Parallel Execution of Sequential Programs," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, pp. 312-316.\*

92. Rosenfeld, J. L., "A Case Study in Programming for Parallel Processors," *Comm. ACM*, Vol. 12, No. 12, Dec. 1969, pp. 645-665.

The use of multiprocessing to decrease the execution time of a single job is investigated here.

93. Siewiorek, D. P., R. Hartenstein, and R. Zaks, "Process Coordination in Multi-Microprocessor Systems," *Proc. Workshop Microarchitecture of Computer Systems*, Nice, France, June 1975.

Discusses the coordination of multiple tasks in a multiprocessing environment. Describes different synchronization techniques and their applicability in a multi-microprocessor environment.

94. Smith, J. W., "Cooperation and Competition: An Approach to Parallel Computation," *Proc. SOUTHEASTCON 1979*, Roanoke, Va., Apr. 1979.

Reviews the use of parallelism in computing systems and discusses resource allocation in parallel systems.

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

## COMPUTER ARCHITECTS & PERFORMANCE ANALYSTS

The Future is Now at  
Fischer & Porter

# 1980

Doing Distributed  
System Design  
Using Modern Technology

GROWTH WITH A PURPOSE  
through worldwide excellence in instrumentation

At Fischer & Porter we are using modern computer technology to develop distributed digital systems for the process control industry such as our highly successful DCI 4000<sup>®</sup> based on multiple microprocessors. We have several key positions immediately available in our Advanced Systems Development Group for creative software professionals.

Distributed System Architecture  
Mathematical System Analysis & Modeling  
Performance Analysis/Monitoring

You should have a MS or PhD in Computer Science, Electrical Engineering or equivalent. Candidates should possess a good working knowledge of operating systems, and performance/tools methodology. Experience with finite-state definitions of communication protocols is desirable.

The working environment will include:

Interactive UNIX\* based computer network  
CRT terminals near your desk  
Structured Analysis & Design  
Programming in High Level Languages (such as C)

We offer an excellent starting salary, a complete benefits package which includes a continuing education program, a modern systems center surrounded by attractive suburban communities ideal for family living with exceptional cultural activities and educational facilities and easy access to the Pocono Mountains and New Jersey seashore resorts.

If your educational experience and interests match our needs, we would like to talk with you. To apply, please send your resume to Neal McDonnell, Dept. 292, Fischer & Porter Company, Warminster, PA 18974 or call collect (215) 674-6205 between 8 AM and 4 PM Monday through Friday.

\*UNIX IS A TRADEMARK OF BELL LABORATORIES.

FISCHER  
& PORTER 

95. Strovink, E. F., "Compilation Strategies for Multiprocessor Message-Passing Systems," *Proc. Seventh Texas Conf. Computing Systems*, Houston, Tex., Nov. 1978, pp. 7-15 to 7-20.\*

Describes the syntax and implementation of a language for the MuNet multiprocessor. An important design goal of this language is the easy expression of parallelism.

96. Subrahmanyam, P. A. and R. D. Kieburz, "Interprocess Communication and Blockage Propagation in Multiprocessor Configurations," *Proc. Twelfth Hawaii Int'l Conf. System Sciences*, Part I, Honolulu, Hawaii, Jan. 1979.  
Considers the interaction of processes and classifies interprocess communication techniques according to the degree of interprocess synchronization.

## Multiprocessor applications

Papers in this section examine the use of multiprocessors in specific applications. The potential reliability and robustness of multiprocessors make them ideal candidates for aerospace applications, where repairs are difficult if not impossible. Some works cited here describe such multiprocessor applications. Others investigate applications such as digital signal processing, simulation, and message switching. A few describe and evaluate specific algorithms implemented on multiprocessors.

97. Barnwell, T. P., S. Gaglio, and C. J. M. Hodges, "Efficient Implementation of One and Two Dimensional Digital Signal Processing Algorithms on a Multi-Processor Architecture," *Rec. IEEE Int'l Conf. Acoustics, Speech and Signal Processing*, Washington, D. C., Apr. 1979, pp. 698-701.

Describes the efficient implementation of certain digital signal processing algorithms on a multiprocessor architecture specifically designed for such applications.

98. Baudet, G. M., "Asynchronous Iterative Methods for Multiprocessors," *J. ACM*, Vol. 25, No. 2, Apr. 1978, pp. 226-244.

Presents techniques for exploiting the parallelism of multiprocessing systems in solving a set of equations. An interesting feature of the algorithms described is that they do not require synchronization among the cooperating processes working on the problem.

99. Davis, E. W., "A Microprocessor-Based Simulation Machine," *Proc. SOUTHEASTCON '78 Region 3 Conf.*, Atlanta, Ga., Apr. 1978.

This multiprocessor system was intended specifically for simulation applications. Presents techniques for using its parallelism to speed up simulations.

100. Erman, L. D. et al., "System Organizations for Speech Understanding: Implications of Network and Multiprocessor Computer Architectures for AI," *IEEE Trans. Computers*, Vol. C-25, No. 4, Apr. 1976, pp. 414-421.

Discusses the implementation of a speech-understanding system on multiprocessors and on network architectures.

101. Jones, A. K. et al., "Programming Issues Raised by a Multiprocessor," *Proc. IEEE*, Vol. 66, No. 2, Feb. 1978, pp. 229-237.

There are problems inherent in writing software for Cm\*. Describes the StarOS operating system and the implementation of a specific application under it.

102. Mazaré, G., "A Few Examples of How to Use a Symmetrical Multi-Micro-Processor," *Conf. Proc. 4th Ann. Symp. Computer Architecture*, Silver Spring, Md., Mar. 1977, pp. 57-62.\*

Describes a microprocessor-based multiprocessor system and illustrates its use with two programs that exploit the parallelism present in the system.

103. McGill, R. and J. Steinhoff, "A Multiprocessor Approach to Numerical Analysis: An Application to Gaming Problems," *Conf. Proc. 3rd Ann. Symp. Computer Architecture*, Clearwater, Fla., Jan. 1976, pp. 46-51.\*

This multiprocessor organization is oriented towards solving a certain class of numerical analysis problems. The system uses a minicomputer host and a set of microprocessor modules. Describes algorithms that exploit the structure of this system to speed up certain problem solutions.

104. Oleinick, P. N., "The Implementation and Evaluation of Parallel Algorithms on C.mmp," doctoral thesis, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa.

Describes experimental performance results and discusses the factors affecting the performance.

105. Robinson, J. T., "Analysis of Asynchronous Multiprocessor Algorithms with Applications to Sorting," *Proc. 1977 Int'l Conf. Parallel Processing*, Detroit, Mich., Aug. 1977, pp. 128-135.\*

106. Soubiron, J., "Multiple-Microprocessor Systems in Attitude and Orbit Control Systems," *Proc. First Int'l Conf. Attitude and Orbit Control Systems*, Noordwijk, Netherlands, Oct. 1977.

Examines the use of multi-microprocessor systems for signal and data processing in space applications, where reliability is crucial.

107. Steele, G. L., "Multiprocessing Compactifying Garbage Collection," *Comm. ACM*, Vol. 18, No. 9, Sept. 1975, pp. 495-508.

Presents algorithms for concurrently doing useful computation and garbage collection in a list-processing environment, such as Lisp.

108. Thomas, T. B. and W. L. Arbuckle, "Multiprocessor Software: Two Approaches," *Proc. 6th Ann. Conf. Use of Digital Computers in Process Control*, Baton Rouge, La., Feb. 1971.

Two uses of multiprocessing in real-time systems are discussed. One approach uses multiple processors as backups to provide enhanced reliability, while the other uses them for improving performance. Discusses the characteristics of the operating system in each case.

109. Tippie, J. W. and J. E. Kulaga, "Design Considerations for a Multiprocessor Based Data Acquisition System," *IEEE Trans. Nuclear Science*, Vol. NS-26, No. 4, Aug. 1979, pp. 4548-4551.

Discusses the use of a multiprocessor system to achieve a high throughput in data acquisition applications.

110. Zellweger, A., "Computer Architectures for Advanced Air Traffic Control Applications," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, pp. 132-139.\*

These multiprocessors are used in highly reliable air traffic control systems; two specific examples are described.

## Scheduling

Multiprocessor scheduling has been a fertile source of interesting problems for researchers in computer science, mathematics, and operations research. In its most general form, the problem involves the scheduling of a set of tasks with arbitrary precedence constraints on a set of processors

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.



with arbitrary characteristics, in order to optimize some objective function (usually the total execution time or the total number of processors used). This general problem yields a variety of specific problems, e.g., whether all processors should be considered identical, whether tasks should be independent of each other (have no precedence constraints), and so on. Most of the papers cited here present results for specific cases of the general scheduling problem. Others investigate the computational complexity of scheduling algorithms. Since almost all the interesting scheduling problems are known to be NP-complete, considerable interest has been shown in heuristic algorithms. Note that most of the results assume the existence only of multiple processors, and not necessarily of shared memory. Obviously, these more general results are applicable to multiprocessors.

111. Bilal, A. Y. and O. I. El-Dessouki, "An Optimization Technique for Parallel Processing in a Generalized Multiprocessor Computer Network," *Proc. Symp. Computers, Electronics and Control*, Alta., Canada, May 1974.  
Presents a heuristic branch and bound technique for obtaining the optimal schedule in a multiprocessor system having processors of varying speeds. Describes and evaluates alternative heuristics to be used in this algorithm.
112. Bokhari, S. H., "Dual Processor Scheduling with Dynamic Reassignment," *IEEE Trans. Software Eng.*, Vol. SE-5, No. 4, July 1979, pp. 341-349.  
Discusses the use of network flow algorithms in the optimal assignment of program subtasks to processors in a dual-processor environment. An extension of earlier work by Stone, this approach differs from Stone's in that modules are dynamically rather than statically assigned.
113. Buten, R. E. and V. Y. Shen, "A Scheduling Model for Computer Systems with Two Classes of Processors," *Proc. 1973 Sagamore Computer Conf. Parallel Processing*, Sagamore, N.Y., Aug. 1973, pp. 130-138.\*
114. Chandy, K. M., "Models for the Recognition and Scheduling of Parallel Tasks on Multiprocessor Systems," *Bull. Operations Research Soc. America*, Vol. 23, Suppl. 1, Spring 1975, p. B-117.  
Examines the problems of recognizing parallelism in programs and of partitioning such programs to minimize execution time. Also discusses the effect of multiprocessor architecture on such partitionings.
115. Chen, N. F. and C. L. Liu, "On a Class of Scheduling Algorithms for Multiprocessor Computing Systems," *Proc. Conf. Parallel Processing*, Raquette Lake, N.Y., Aug. 1974.  
Describes a class of scheduling algorithms called *level algorithms* and derives the conditions under which this class yields optimal schedules. Also presents quantitative estimates of the quality of these algorithms in the cases where they are suboptimal.
116. Chen, N. F., "An Analysis of Scheduling Algorithms in Multiprocessor Computer Systems," Tech. Report UIUCDCS-R-75-724, Dept. of Computer Science, University of Illinois, Urbana, 1975.
117. Coffman, E. G. and R. Sethi, "Algorithms Minimizing Mean Flow Time: Schedule-Length Properties," *Acta Informatica*, Vol. 6, No. 1, 1976, pp. 1-14.
118. Dhall, S. K. and C. L. Liu, "On a Real-Time Scheduling Problem," *Operations Research*, Vol. 26, No. 1, Jan.-Feb. 1978, pp. 127-140.  
Provides two heuristic algorithms for deciding the number of processors needed to service periodically arriving time-critical requests.
119. Ecker, K., "On Task Scheduling in an Multiprocessor Environment," *Micros, Minis, & Maxis, Technology Thrust vs. User Requirement—Digest of Papers—COMPCON Fall 77*, Washington, D. C., Sept. 1977, pp. 297-298.\*  
Discusses heuristic algorithms for scheduling multiple processors.
120. Fernandez, E. B. and T. Lang, "Computation of Lower Bounds for Multiprocessor Schedules," *IBM J. Research and Development*, Vol. 19, No. 5, Sept. 1975, pp. 435-444.  
Presents computationally efficient techniques for calculating lower bounds on the number of processors needed to complete a set of tasks within a given time, and on the minimum execution time using a fixed number of processors. The analysis assumes non-preemptive scheduling, identical processors, and a set of tasks with precedence constraints.
121. Garey, M. R. and D. S. Johnson, "Complexity Results for Multiprocessor Scheduling under Resource Constraints," *SIAM J. Computing*, Vol. 4, No. 4, Dec. 1975, pp. 397-411.  
Considers an abstract model of a multiprocessor system and obtains lower bounds on the complexity of scheduling algorithms for this model. Presents an important but rather depressing result—almost all interesting scheduling problems are NP-complete.
122. Garey, M. R. and R. L. Graham, "Bounds for Multiprocessor Scheduling with Resource Constraints," *SIAM J. Computing*, Vol. 4, No. 2, June 1975, pp. 187-200.  
Presents upper and lower bounds for non-preemptive scheduling algorithms applicable to a multiprocessor system composed of identical processors. The tasks are assumed to have precedence constraints and nonidentical resource requirements.
123. Garey, M. R. and D. S. Johnson, "Scheduling Tasks with Nonuniform Deadlines on Two Processors," *J. ACM*, Vol. 23, No. 3, July 1976, pp. 461-467.
124. Garey, M. R. and D. S. Johnson, "Two-Processor Scheduling with Start-Times and Deadlines," *SIAM J. Computing*, Vol. 6, No. 3, Sept. 1977, pp. 416-426.
125. Gonzalez, T. and S. Sahni, "Open Shop Scheduling to Minimize Finish Time," *J. ACM*, Vol. 23, No. 4, Oct. 1976, pp. 665-679.  
Examines the problem of obtaining a minimum finish time schedule for a set of preemptible tasks. Presents a linear time algorithm for the two-processor case and a polynomial algorithm for the general case. The problem for non-preemptible tasks is proved NP-complete.
126. Gonzalez, M. J., "Deterministic Processor Scheduling," *Computing Surveys*, Vol. 9, No. 3, Sept. 1977, pp. 173-204.  
Surveys processor scheduling and presents results on scheduling in a multiprocessor environment. Describes a wide range of scheduling constraints and examines results for each constraint. This is an excellent starting point for anyone interested in detailed information on scheduling.
127. Gonzalez, T., O. H. Ibarra, and S. Sahni, "Bounds for LPT Schedules on Uniform Processors," *SIAM J. Computing*, Vol. 6, No. 1, Mar. 1977, pp. 155-166.  
Using an LPT (largest processing time) algorithm as an approximation to the optimal schedule, investigates the scheduling of independent, non-preemptible tasks. The processors are identical except for their speeds.
128. Gonzalez, T. and S. Sahni, "Preemptive Scheduling of Uniform Processor Systems," *J. ACM*, Vol. 25, No. 1, Jan. 1978, pp. 92-101.

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

- Presents a linear time algorithm for computing the optimal finish time for a set of independent tasks, ordered by length, on a set of processors, ordered by speed.
129. Gwynn, J. M. and R. J. Raynor, "Scheduling in a Multiprocessor Environment," *Proc. 1973 Sagamore Computer Conf. Parallel Processing*, Sagamore, N.Y., Aug. 1973, p. 139.\*  
Examines interrupt handling techniques in multiprocessor systems. Describes schemes such as master-slave control and floating executive control and discusses their impact on the queueing delays in handling interrupts.
  130. Horvath, E. C., S. Lam, and R. Sethi, "A Level Algorithm for Preemptive Scheduling," *J. ACM*, Vol. 24, No. 1, Jan. 1977, pp. 32-43.  
Extends earlier work on the preemptive scheduling of a set of tasks whose precedence graph is a tree, on processors with different speeds. Presents an algorithm which is optimal for the two-processor case and shows that, in other cases, the algorithm is optimal only when the tasks are independent.
  131. Jensen, J. E., "A Fixed-Variable Scheduling Model for Multiprocessors," *Proc. 1977 Int'l Conf. Parallel Processing*, Detroit, Mich., Aug. 1977, pp. 108-117.\*  
This adaptive scheduling algorithm selects a scheduling strategy appropriate to the load on the system. The algorithm uses a simplistic but computationally cheap approach when processor cycles are scarce. It adopts more complex scheduling strategies as processor cycles become available.
  132. Kafura, D. G., "Relationship Between Worst-Case and Expected Scheduling Performance for a Model of a Multiprocessor System," *Proc. Computer Science Conf.*, Washington, D.C., Feb. 1975.
  133. Kafura, D. G. and V. Y. Shen, "Task Scheduling on a Multiprocessor System with Independent Memories," *SIAM J. Computing*, Vol. 6, No. 1, Mar. 1977, pp. 167-187.  
Considers scheduling in a multiple processor system where all processors are identical and each processor has its own, private memory. Examines preemptive and non-preemptive scheduling schemes for tasks with known resource requirements, and evaluates these schemes both analytically and by simulation.
  134. Kohler, W. H., "A Preliminary Evaluation of the Critical Path Method for Scheduling Tasks on Multiprocessor Systems," *IEEE Trans. Computers*, Vol. C-24, No. 12, Dec. 1975, pp. 1235-1238.  
Expresses the problem of scheduling tasks on a set of independent, identical processors in a graph-theoretic framework and, using a critical path method, derives an approximation to the optimal solution.
  135. Lang, T. and E. B. Fernandez, "Scheduling of Unit-Length Independent Tasks with Execution Constraints," *Information Processing Letters*, Vol. 4, No. 4, Jan. 1976, pp. 95-98.
  136. Liu, J. W. S. and C. L. Liu, "Performance Analysis of Multiprocessor Systems Containing Functionally Dedicated Processors," *Acta Informatica*, Vol. 10, No. 1, 1978, pp. 95-104.  
Presents models to describe the scheduling of multiprocessing systems having specialized processors.
  137. Martin-Vega, L. A. and H. D. Ratliff, "Scheduling Parallel Processors," *Bull. Operations Research Soc. America*, Vol. 23, Suppl. 2, Nov. 1975, p. B-418.  
Investigates the differences between preemptive and non-preemptive scheduling algorithms. Shows that, for a subclass of scheduling problems, the relationship between preemptive and non-preemptive optimal solutions is isomorphic to the relationship among optimal solutions to a certain type of linear programming problem.
  138. Millen, J. K., "Parallel Derivatives for Multiprocessor Task Scheduling," Report No. MTR-2922, Mitre Corp., McLean, Va., May 1975.  
This technique transforms an arbitrary flowchart into the equivalent maximally parallel flowchart and uses it to derive alternative scheduling strategies for a multiprocessor.
  139. Rowicki, A., "A Note on Optimal Scheduling for Two-Processor Systems," *Information Processing Letters*, Vol. 4, No. 2, Nov. 1975, pp. 27-30.  
Describes an algorithm for the optimal non-preemptive scheduling of identical tasks on identical processors.
  140. Schindler, S. and H. Lüdtkke, "An Approach to a Restricted Scheduling-Problem for Multiprocessor Systems," *Proc. 1973 Sagamore Computer Conf. Parallel Processing*, Sagamore, N.Y., Aug. 1973, pp. 121-129.\*  
Discusses the preemptive scheduling of tasks with precedence constraints on a fixed set of identical processors, with all other resources unlimited.
  141. Sethi, R., "Scheduling Graphs on Two Processors," *SIAM J. Computing*, Vol. 5, No. 1, Mar. 1976, pp. 73-82.  
Presents two algorithms for scheduling equal-length tasks with precedence constraints on two-processor systems.
  142. Stone, H. S., "Multiprocessor Scheduling with the Aid of Network Flow Algorithms," *IEEE Trans. Software Eng.*, Vol. SE-3, No. 1, Jan. 1977, pp. 85-93.  
These scheduling algorithms, based on the Ford-Fulkerson network flow algorithm, minimize interprocessor communication in a multiprocessor system. The two-processor case is treated in detail, and partial results are presented for the general case.
  143. Yang, Chao-Chih, "Fast Algorithms for Bounding the Performance of Multiprocessor Systems," *Proc. 1976 Int'l Conf. Parallel Processing*, Walden Woods, Mich., Aug. 1976, pp. 73-82.\*  
Presents efficient algorithms for finding the lower bound on performance when scheduling a set of tasks. The performance measures here are the number of processors used and the total execution time for the set of tasks.
  144. Yao, A. Chi-Chih, "Scheduling Unit-Time Tasks with Limited Resources," *Proc. Conf. Parallel Processing*, Raquette Lake, N.Y., Aug. 1974.  
Gives heuristic algorithms for optimally scheduling a set of tasks with different resource requirements, but with identical processing times.

## Reliability and error recovery

Papers in this section investigate the reliability of multiprocessor systems and describe techniques to utilize available redundancy. Some try to define figures of merit for such entities as "reliability" and "fault-tolerance," so that quantitative comparisons may be made. Techniques to handle hardware and software failures, and to permit graceful degradation in case of such failures, are also described. Reliability models and their use in predicting failure rates are the focus of a number of the works cited here.

145. Bertera, F., A. Boccalatte, and M. Di Manzco, "A Fault-Tolerant Approach to Processor Synchronization in a Multi-

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

Microprocessor Environment," *Proc. Conf. Microprocessors and Microprogramming*, Amsterdam, Netherlands.

Considers the problem of software synchronization in the presence of hardware failure in a multiprocessor system. Describes a fault-tolerant synchronization technique for critical applications.

146. Hayes, J. P. and R. Yanney, "Fault Recovery in Multiprocessor Networks," *Digest of Papers—Eighth Ann. Int'l Conf. Fault-Tolerant Computing*, Toulouse, France, June 1978, pp. 123-128.\*  
Uses graph theory to describe dynamic reconfiguration and error recovery in multiple processor networks.
147. Ingle, A. and D. P. Siewiorek, "Reliability Modeling of Multiprocessor Structures," *Computers... by the Millions, for the Millions—Digest of Papers—COMPCON Fall 76*, Washington, D.C., Sept. 1976, pp. 24-29.\*  
Discusses the reliability and robustness of multiprocessor systems. Reliability models of two multiprocessor systems, C.mmp and Cm\*, are presented and used to compare the reliability of these systems with that of their uniprocessor counterparts.
148. Ingle, A. and D. P. Siewiorek, "Reliability Models for Multiprocessor Systems with and without Periodic Maintenance," *Proc. Seventh Ann. Int'l Conf. Fault-Tolerant Computing*, Los Angeles, Calif., June 1977, pp. 3-9.\*  
Examines the effects of integrity checks and periodic maintenance on the reliability of multiprocessor systems whose performance degrades gracefully in the event of failure.
149. Katsuki, D. et al., "Pluribus—An Operational Fault-Tolerant Multiprocessor," *Proc. IEEE*, Vol. 66, No. 10, Oct. 1978, pp. 1146-1159.  
Describes the error recovery aspects of a multiprocessor with high availability requirements; the system discussed is used for message switching in the Arpanet communication subnetwork.
150. Losq, J., "Effects of Failures on Gracefully Degradable Systems," *Proc. Seventh Ann. Int'l Conf. Fault-Tolerant Computing*, Los Angeles, Calif., June 1977, pp. 29-34.\*  
Considers multiprocessor systems designed to gracefully degrade after hardware or software failure. Presents a model of graceful degradation and describes quantitative measures for comparing "gracefully degradable" systems.
151. MacWilliams, W. H., "Reliability of Large Real-Time Control Software Systems," *Proc. IEEE Symp. Computer Software Reliability*, New York, May 1973.  
Discusses issues in designing reliable software for real-time multiprocessor systems.
152. Miller, J. S., "Fault-Tolerance Features of an Aerospace Multiprocessor," *AGARD Conf. Proc. No. 149, Real-Time Computer-Based Systems (NATO Advisory Group for Aerospace R&D)*, Athens, Greece, May 1974.  
This system uses multiprocessing primarily to improve reliability. Results of multiple, concurrently executing processors are compared on a per-instruction basis, and automatic reloading of one processor by another is used in case of processor failure.
153. Murray, N. D., A. L. Hopkins, and J. H. Wensley, "Highly Reliable Multiprocessors," in *Integrity in Flight Control Systems*, Report AGARD-AG-224, NATO Advisory Group for Aerospace R&D, Neuilly sur Seine, France, Apr. 1977.  
Discusses the reliability requirements of multiprocessors for critical avionic systems and describes two designs meeting these requirements.
154. Owen, G. J., "Rollback—A Method of Process and System Recovery," *Proc. Conf. Software Eng. for Telecommunica-*
- tion Switching Systems*, Colchester, England, Apr. 1973.  
The GEC Mark II multiprocessor is described, and a software technique for error recovery, called *rollback*, is discussed in detail. Rollback essentially consists of restarting an errant process with fresh data.
155. Quillin, W. E., "Reliability of Multiprocessor Systems," *Proc. 1st European Seminar Computing with Real-Time Systems*, Harwell, England, 1972.  
Contains a general discussion of the use of multiprocessing in achieving high reliability. Also contains a brief description of some systems that use multiprocessing primarily for this purpose.
156. Robinson, J. G. and E. S. Roberts, "Software Fault-Tolerance in the Pluribus," *AFIPS Conf. Proc.*, Vol. 47, 1978 NCC, pp. 563-569.  
Describes software features that enhance reliability in the Pluribus multiprocessor.
157. Robinson, J. G., "The Pluribus Fault-Tolerant Multiprocessor," *Exploding Technology, Responsible Growth—Digest of Papers—COMPCON Spring 79*, San Francisco, Calif., Feb. 1979, pp. 45-48.\*  
Describes the fault tolerance features of, and operational experience with, Pluribus.
158. Saheban, F., L. Simoncini, and A. D. Friedman, "Concurrent Computation and Diagnosis in Multiprocessor Systems," *Digest of Papers—Ninth Ann. Int'l Symp. Fault-Tolerant Computing*, Madison, Wisc., June 1979, pp. 149-156.\*  
Discusses issues involved in designing systems which do concurrent computation and perform self-diagnosis. The discussion is motivated by, and oriented toward, multiprocessor systems.
159. Shrivastava, S. K. and J. P. Banâtre, "Reliable Resource Allocation Between Unreliable Processes," *IEEE Trans. Software Eng.*, Vol. SE-4, No. 3, May 1978, pp. 230-241.  
Examines error recovery in interacting processes and describes error recovery techniques for cooperating processes. Presents programming language constructs to express these techniques.
160. Siewiorek, D. P. et al., "A Case Study of C.mmp, Cm\*, and C.vmp: Part I—Experiences with Fault Tolerance in Multiprocessor Systems," *Proc. IEEE*, Vol. 66, No. 10, Oct. 1978, pp. 1178-1199.  
Describes the fault tolerance features of three multiprocessors and presents data on their reliability.
161. Siewiorek, D. P. et al., "A Case Study of C.mmp, Cm\*, and C.vmp: Part II—Predicting and Calibrating Reliability of Multiprocessor Systems," *Proc. IEEE*, Vol. 66, No. 10, Oct. 1978, pp. 1200-1220.  
Presents reliability models for multiprocessor systems and compares the predictions of these models with observed data on three multiprocessor systems.
162. Siewiorek, D. P., "Multiprocessors: Reliability, Modelling and Graceful Degradation," in *System Reliability and Integrity, State of the Art Report*, Infotech Ltd., Maidenhead, England.  
Discusses the reliability and fail-soft characteristics of multiprocessor systems and describes quantitative techniques for evaluating systems along these dimensions.
163. Waters, S. J., "Majority Verdicts in Multi-Processing—Any Two From Three," *Computer J.*, Vol. 20, No. 3, Aug. 1977, pp. 207-212.

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

A software technique uses voting among three processors performing identical tasks to detect computational errors and processor failure. Different approaches to this problem are proposed and their relative merits discussed.

## Miscellaneous topics in multiprocessing

Included here are works I felt were of interest, but which did not fall into any of the earlier categories. Some are general surveys of multiprocessing. Others describe experiences in using multiprocessors. A few give advice on the appropriate use of multiprocessing.

164. Baer, J. L., "Multiprocessing Systems," *IEEE Trans. Computers*, Vol. C-25, No. 12, Dec. 1976, pp. 1271-1277.  
Examines and classifies array processors and multiprocessors according to the tightness of coupling between individual processors. Discusses software enhancements needed to exploit the potential of multiprocessor systems.
165. Childs, R. E., "Multiple Microprocessor Systems: Goals, Limitations and Alternatives," *Exploding Technology, Responsible Growth—Digest of Papers—COMPCON Spring 79*, San Francisco, Calif., Mar. 1979, pp. 94-97.\*  
Surveys the merits and appropriateness of using multiprocessing in microprocessor systems.
166. Enslow, P. H., ed., *Multiprocessors and Parallel Processing*, John Wiley and Sons, New York, 1974.  
Describes hardware and software issues involved in multiprocessing. An appendix includes brief descriptions of a number of multiprocessing systems.
167. Enslow, P. H., "Multiprocessor Organization—A Survey," *Computing Surveys*, Vol. 9, No. 1, Mar. 1977, pp. 103-129.  
A well-written and oft-quoted survey of multiprocessing. Describes a variety of processor-memory interconnection schemes as well as different operating system structures.
168. Ferreira, R. C. and D. Vojnovic, "Multimicrocomputers: A Perspective on the Next Five Years," in *Future Systems, State of the Art Report*, Infotech Ltd., Maidenhead, England, 1977.  
Predicts the future of multiprocessors that are built with microcomputers, e.g., C.mmp.
169. Franklin, M. A., S. A. Kahan, and M. J. Stucki, "Design Issues in the Development of a Modular Multiprocessor Communications Network," *Conf. Proc. 6th Ann. Symp. Computer Architecture*, Philadelphia, Pa., Apr. 1979, pp. 182-187.\*  
Describes a modular, easily expandable crossbar network for use in multiprocessor systems.
170. Hardcastle, A. R. K., "Multi-Minis versus Large Mainframes," in *Minis versus Mainframes, State of the Art Report*, Infotech Ltd., Maidenhead, England, 1978.  
Discusses the pros and cons of using multi-microprocessors in lieu of large mainframes for commercial data processing.
171. Harris, J. A. and D. R. Smith, "Hierarchical Multiprocessor Organizations," *Conf. Proc. 4th Ann. Symp. Computer Architecture*, Silver Spring, Md., Mar. 1977, pp. 41-48.\*  
Examines the benefits of multi-microprocessor organizations and the problems encountered in such systems. Describes a specific multiprocessor design, based on microprocessors, and investigates its suitability for some classes of problems.
172. Jones, A. K. and P. Schwarz, "Experience Using Multiprocessor Systems: A Status Report," Tech. Report CMU-CS-79-146, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., Oct. 1976.  
The discussion here is based on actual experience in building and using multiprocessors; it thus falls into a relatively small group of experience-oriented papers.
173. Leiberman, M., "To Multiprocess or Not to Multiprocess," *Proc. Fourteenth Meeting Computer Performance Evaluation Users Group*, Boston, Mass., Oct. 1978.  
Considers the decision to run multiple processors as a multiprocessing system or as independent systems.
174. Mazare, G., "Multiprocessor Systems," *Proc. 1974 CERN School of Computing* (European Organization for Nuclear Research), Godysund, Norway, Aug. 1974.  
A detailed description of the development and principles of multiprocessing systems.
175. Patel, J. H., "Processor-Memory Interconnections for Multiprocessors," *Conf. Proc. 6th Ann. Symp. Computer Architecture*, Philadelphia, Pa., Apr. 1979, pp. 168-177.\*  
Describes an interconnection structure—more cost-effective than a crossbar switch—that permits every processor to be connected to every memory module.
176. Siewiorek, D. P. and M. R. Barbacci, "Modularity and Multiprocessor Structures—Some Open Problems in the Construction and Utilization of Mini- and Microprocessor Networks," in *Distributed Systems, International State of the Art Report*, Infotech Ltd., Maidenhead, England, 1976.  
Presents a taxonomy of the multiprocessor design space and discusses three multiprocessor designs in this light. Discusses both hardware and software issues related to the three designs.
177. Srodawa, R. J., "Positive Experiences with a Multiprocessing System," *Computing Surveys*, Vol. 10, No. 1, Mar. 1978, pp. 73-82.  
Examines the experience gained from using multiprocessing in the Michigan Terminal System. Based on this experience, evaluates the options available to the designer of a multiprocessing system. This paper is one of the few pieces of widely circulated literature describing actual multiprocessor experience rather than system design or modeling.
178. White, C. H., ed., *Multiprocessor Systems, International Computer State of the Art Report*, Infotech Ltd., Maidenhead, England, 1976.  
A collection of articles on the state of the art of multiprocessing in 1976. Most of the information is still valid and relevant.
179. Wulf, W. A. and S. P. Harbison, "Reflections in a Pool of Processors," Tech. Report No. CMU-CS-78-103, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., Feb. 1978.  
A candid evaluation of the successes and shortcomings of the C.mmp project. The report is refreshingly frank and pulls no punches.

\*These digests and proceedings are available from the IEEE Computer Society Publications Office, 5855 Naples Plaza, Suite 301, Long Beach, CA 90803.

M. Satyanarayanan is the author of the survey accompanying this bibliography; his photograph and biography appear on page 96.