Chapter 6

Conclusion

The goal of this dissertation has been to present a technique for modeling distributed parallel applications and a scheduling approach that can utilize enhanced modeling information to determine performance-efficient application schedules. We presented structural modeling, an adaptable, flexible approach to compositional modeling, and extended this technique to allow stochastic parameters. We then used structural models with stochastic predictions to tune a time balancing scheduling policy for data parallel applications. Our results showed that stochastic scheduling policies often resulted in faster execution times and more predictable application behavior.

6.A Critique and Extensions

As in any thesis, this work circumscribed a problem space that can be extended in many directions. A few of these areas include:

Extending the applicability of structural models. The structural modeling approach presented in Chapter 3 adequately predicted the behavior of two classes of applications applications when the performance of the application consisted of communication and computation, but did not account for other execution behaviors. Two execution behaviors we plan to add to the structural modeling approach are I/O behavior, perhaps building off work by Smirni [SRar, RSSS98], and memory behavior, possibly using work presented in PMH [ACFS93] or the LDA Model [SW98a] as starting points. In addition, we would like to expand the classes of applications that structural models have been shown to apply to.

Expanding the use of structural models. We demonstrated the use of structural models in predictions for scheduling, there are other areas in which such an approach might prove to be useful. Future work includes examining how application developers might use the modular design and extensibility of structural models as part of a large software design project in order to predict the execution time of a partially written application and steer development efforts. In addition, we would like to examine the use of structural models for porting efforts. A very common problem for application scientists is the fact that the platforms continually change as hardware resources are upgraded. Given an application running on soon-to-be-retired platform, applications developers would like to have a pragmatic way to decide where their application would best be ported to, and structural models may provide one way to determine this information.

Using other forms of extended information. We extended structural models, predictions and scheduling to take into account stochastic values. There are other sources of meta-data that might also be beneficial to include, such as error or bias predictions, or estimates of computational complexity. Another common problem seen in predictions is the lack of a complete set of information, so we would like to investigate how to compensate for incomplete information in predictions and scheduling.

Extending the applicability of stochastic predictions. In Chapter 4 we detailed a predictive approach to be used as part of an on-line scheduler. We would like to address the different needs of predictions used for resource selection, as this is also a very important part of achieving performance on parallel distributed systems. In addition, we would like to expand the systems that our predictions apply to, both in terms of being able to add other types of machines, for example including predictions for queue times so this approach can take into account queued systems, and by examining what factors would need to be added in order to use this approach for larger grid-style computing systems.

Extending stochastic scheduling. By far, the most room for future work can be seen in the stochastic scheduling approach. Some issues we plan to address are:

- Can stochastic data be used in other parts of the scheduling process besides data allocation?
- What other scheduling policies can take advantage of stochastic values, and how?
- Further extensions to time-balancing: definitions for other stochastic value representations; inclusion of networking information, memory data; other Tuning factor definitions
- Dynamic re-balancing approaches

All of these are logical extensions to the work presented.

6.B Conclusions and Contributions

Distributed parallel platforms are becoming more and more prevalent in research environments. As they grow in popularity, so does their user base. Because of this, in part, a growing problem in achieving performance on these platforms is resource contention. Our solution to this problem was the development of a more sophisticated scheduling approach that could factor in the variability of resource performance capabilities.

This thesis presented three main contributions: an approach to define distributed parallel application performance models called **structural modeling**, the ability to parameterize these models with **stochastic values** in order to meet the prediction needs of shared clusters of workstations, and a **stochastic scheduling policy** that can make use of stochastic predictions to achieve better application execution times and more predictable application behavior. We verified each of these contributions for an extensive test suite of applications, several platforms and a variety of load conditions.

In conclusion, we believe that structural modeling is a promising approach for predicting the behavior of parallel applications running on distributed platforms due to its extensibility and flexibility. This approach can be augmented to allow for a variety of sources and types of data, as we extended structural models to allow for stochastic valued parameters. The resulting predictions from these models can be used in a variety of settings, with out example of an on-line scheduler just one possibility.

We demonstrated that stochastic values contribute valuable additional information about the performance of applications running on shared clusters of workstations. Both using interval and normal distribution representations of stochastic data can be used efficiently, with differing results that will meet different prediction needs.

In addition, we defined a stochastic scheduling approach that defined an initial approach but also opens up a broad new area of possible schedulers that take into account additional available information. Our initial instantiation shows promise in reducing execution time for systems with a high variably in available performance, as seen in many shared cluster environments.

Bibliography

[ACFS93] B. Alpern, L. Carter, E. Feig, and T. Selker. The uniform memory hierarchy model of computation. *Algorithmica*, 1993.

[Adv93] Vikram S. Adve. *Analyzing the Behavior and Performance of Parallel Programs.* PhD thesis, University of Wisconsin-Madison, December 1993. Also available as University of Wisconsin Computer Sciences Technical Report #1201.

[AHWC98] Scott Anderson, David Hart, David Westbrook, and Paul Cohen. A toolbox for analyzing programs. *International Journal on AI Tools*, 1998. to appear.

[AL77] A. Arakawa and V. R. Lamb. Computational design of the basic dynamical processes of the ucla general circulation model. *Methods in Comput. Phys.*, 17:173–265, 1977.

[And97] Kelsey Anderson. Personal communication, Summer 1997.

[ASWB95] Cosimo Anglano, Jennifer Schopf, Richard Wolski, and Francine Berman. Zoom: A hierarchical representation for heterogeneous applications. Technical Report #CS95-451, University of California, San Diego, Computer Science Department, 1995.

[AV93] Vikram Adve and Mary Vernon. The influence of random delays on parallel execution times. In *Proceedings of Sigmetrics '93*, 1993.

[Bai72] A. Baines. Introduction. In Owen L. Davies and Peter L. Goldsmith, editors, *Statistical Methods in Research and Production*, chapter One. Published for Imperial Chemical Industries Limited by Oliver and Boyd, 1972.

[Bai95] D. H. Bailey, editor. *Seventh SIAM conference on parallel processing for scientific computing*, 1995.

[Bar78] B. Austin Barry. *Errors in Practical Measurement in Science, Engineering and Technology.* John Wilsey & Sons, 1978.

[BBB+91] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T.A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrishnam, and S. K. Weeratunga. The nas parallel benchmarks. *International Journal of Supercomputer Applications*, 5(3):63–73, Fall 1991.

[BBC+94] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*. SIAM, Philadelphia, PA, 1994.

[BC88] J. Bruno and P.R. Cappello. Implementing the beam and warming method on the hypercube. In *Proceedings of 3rd Conference on Hypercube Concurrent Computers and Applications*, 1988.

[BCF+98] S. Brunett, K. Czajkowski, S. Fitzgerald, I. Foster, A. Johnson, C. Kesselman, J. Leigh, and S. Tuecke. Application experiences with the globus toolkit. In *Proceedings of the 7th IEEE Symp. on High Performance Distributed Computing (HPDC-7)*, 1998.

[BDGM93] A. Beguelin, J. Dongarra, A. Geist, and R. Manchek. Hence: A heterogeneous network computing environment. *Scientific Programming*, 3(1):49–60, 1993.

[Ber98] Francine Berman. *The Grid: Blueprint for a New Computing Infrastructure, edited by Ian Foster and Carl Kesselman*, chapter 12: High-Performance Schedulers. Morgan Kaufmann Publishers, Inc., 1998.

[BFM91] Marcos A. G. Brasileiro, James A. Field, and Antao B. Moura. Modeling and analysis of time critical applications on local area networks. In *Proc. of Chilean Computer Science Society*, pages 459–71, 1991.

[BFVW93] E. Barszcz, R. Fatoohi, V. Venkatakrishnan, and S. Weeratunga. Solution of regular, sparse triangular linear systems on vector and distributed-memory multiprocessors. Technical Report Technical Report NAS RNR-93-007, NASA Ames Research Center, Moffett Field, CA, April 1993.

[BG98] J. Bolliger and T. Gross. A framework-based approach to the development of network-aware applications. *IEEE Transactions on Software Engineering (Special Issue on Mobility and Network-Aware Computing)*, 24(5), May 1998.

[Bha96] Karan Bhatia. Personal communication, 1996.

[BHS+95] David Bailey, Tim Harris, William Saphir, Rob van der Wijngaart, Alex Woo, and Maurice Yarrow. The nas parallel benchmarks 2.0. Technical Report Report NAS-95-020, Numerical Aerospace Simulation Facility at NASA Ames Research Center, December 1995. Also available at `http://science.nas.nasa.gov/Software/NPB/Specs/npb2_0/npb2_0.html`.

[BK91] Scott B. Baden and Scott R. Kohn. A comparison of load balancing strategies for particle methods running on mimd multiprocessors. Technical Report CS91-199, University of California, San Diego, Computer Science and Engineering Dept., 1991.

[Bla80] Leland Blank. *Statistical Procedures for engineering, management, and science*. McGraw-Hill Book Company, 1980.

[Bri87] William L. Briggs. *A Multigrid Tutorial*. Society for Industrial and Applied Mathematics, Lancaster Press, 1987.

[BW96] Francine Berman and Richard Wolski. Scheduling from the perspective of the application. In *Proceedings of the Fifth IEEE International Symposium on High Performance Distributed Computing*, 1996.

[BW97] Francine Berman and Richard Wolski. The apples project: A status report. In *Proceedings of the 8th NEC Research Symposium*, 1997.

[BW98] Francine Berman and Richard Wolski. Performance prediction engineering. DARPA report, available from `http://www.cs.ucsd.edu/groups/hpcl/apples/PPE/report.html`, 1998.

[BWF+96] Francine Berman, Richard Wolski, Silvia Figueira, Jennifer Schopf, and Gary Shao. Application-level scheduling on distributed heterogeneous networks. In *Proceedings of SuperComputing '96*, 1996.

[BWS97] Francine Berman, Richard Wolski, and Jennifer Schopf. Performance prediction engineering for metacomputing. `http://www.cs.ucsd.edu/groups/hpcl/apples/PPE/index.html`, 1997.

[Car98] Larry Carter. Personal communication, 1998.

[CGS95] Brad Calder, Dirk Grunwald, and Amitabh Srivastava. The predictability of branches in libraries. In *Proceedings of the 28th International Symposium on Microarchitecture*, 1995. Also available as WRL Research Report 95.6.

[Cha94] K. M. Chandy. Concurrent program archetypes. In *Proceedings of the 1994 Scalable Parallel Libraries Conference*, 1994.

[CI98] Alex Cereghini and Wes Ingalls. Personal communication, 1998.

[Col89] Murray Cole. *Algorithmic Skeletons: Structured Management of Parallel Computation.* Pitman/MIT Press, 1989.

[Cox84] M. D. Cox. A primitive, 3-dimensional model of the ocean. Technical Report Technical Report 1, GFDL Ocean Group, 1984.

[CQ93] Mark J. Clement and Michael J. Quinn. Analytical performance prediction on multicomputers. In *Proceedings of SuperComputing '93*, 1993.

[Cra68] J. M. Craddock. *Statistics in the Computer Age.* The English Universities Press Limited, 1968.

[Cro94] Mark Edward Crovella. *Performance prediction and tuning of parallel programs.* PhD thesis, University of Rochester, 1994.

[CS95] Robert L. Clay and Peter A. Steenkiste. Distributing a chemical process optimization application over a gigabit network. In *Proceedings of SuperComputing '95*, 1995.

[Das97] Chita Das. Personal communication with rich wolski, 1997.

[DDH+98] Ewa Deelman, Aditya Dube, Adolfy Hoisie, Yong Luo, Richard L. Oliver, David Sundaram-Strukel, Harvey Wassermann, Vikram S. Adve, Rajive Bagrodia, James C. Browne, Elias Houstis, Olaf Lubeck, John Rice, Patricia J. Teller, and Mary K. Vernon. Poems: End-to-end performance design of large parallel adaptive computational systems. In *Proceedings of the First International Workshop on Software and Performance*, October 1998.

[DFH+93] J. Darlington, A. J. Field, P. G. Harrison, P. H. J. Kelly, D. W. N. Sharp, Q. Wu, and R. C. While. Parallel programming using skeleton functions. In *Proceedings of Parallel Architectures and Languages Europe (PARLE) '93*, 1993.

[Din98] P. Dinda. The statistical properties of host load. In *Proceedings of the 4th Workshop on Languages, Compilers, and Runtime Systems for Scalable Computers*, 1998. Also available from `http://www.cs.cmu.edu/~pdinda/html/papers.html`.

[DMP97] Ralf Diekmann, Burkahrad Monien, and Robert Pries. Load balancing strategies for distributed memory machines. In F. Karsh and H. Saltz, editors, *Multi-Scale Phenomena and their Simulation.* World Scientific, 1997.

[DN95] J. Dongarra and Peter Newton. Overview of vpe: A visual environment for message-passing parallel programming. In *Proceedings of the 4th Heterogeneous Computing Workshop*, 1995.

[DO98] P. Dinda and D. O'Hallaron. The statistical properties of host load. Technical Report Technical Report CMU-CS-98-143, School of Computer Science, Carnegie Mellon University, 1998. Also available from `http://www.cs.cmu.edu/~pdinda/html/papers.html`.

[DP96a] Jay Devore and Roxy Peck. *Statistics: The Exploration and Analysis of Data*, page 272. Duxbury Press, 1996.

[DP96b] Jay Devore and Roxy Peck. *Statistics: The Exploration and Analysis of Data*, page 88. Duxbury Press, 1996.

[DP96c] Jay Devore and Roxy Peck. *Statistics: The Exploration and Analysis of Data*, page 567. Duxbury Press, 1996.

[DS95] J. Demmel and S. L. Smith. Performance of a parallel global atmospheric chemical tracer model. In *Proceedings of SuperComputing '95*, 1995.

[DTMH96] D. M. Deaven, N. Tit, J. R. Morris, and K. M. Ho. Structural optimization of lennard-jones clusters by a genetic algorithm. *Chemical Physical Letters*, 256:195, 1996.

[FB96] Silvia Figueira and Francine Berman. Modeling the effects of contention on the performance of heterogeneous applications. In *Proceedings of the Fifth IEEE Symposium on High Performance Distributed Computing*, 1996.

[FJ98] A. Fersha and J. Johnson. Performance oriented development of spmd program based on task structure specifications. to appear, also available from `http://sokrates.ani.univie.ac.at/~ferscha/publications.html`, 1998.

[FK97] Ian Foster and Carl Kesselman. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications*, Summer 1997. Also available at `ftp://ftp.mcs.anl.gov/pub/nexus/reports/ globus.ps.Z`.

[Fos95] Ian Foster. *Designing and Building Parallel Programs.* Addison-Wesley, 1995.

[FOW84] J. Ferrante, K. J. Ottenstein, and J. D. Warren. The program dependence graph and its use in optimization. In *Proceedings of the International Symposium on Programming 6th Colloquium*, 1984.

[Geh95] Jörn Gehring. Dynamic program description as a basis for runtime optimization. Technical Report PC^2/TR-002-97, Paderborn Center for Parallel Computing, 1995.

[Ger98] Apostolos Gerasoulis. Personal communication, 1998.

[GKM82] S. Graham, P. Kessler, and M. McKusick. gprof: A call graph execution profiler. In *Proceedings of the SIGPLAN '82 Symposium on Compiler Construction*, 1982. Also published in SIGPLAN Notices, volume 17, number 6, pages 120–126.

[GLFK98] Andrew Grimshaw, Michael Lewis, Adam Ferrari, and John Karpovich. Architectural support for extensibility and autonomy in wide-area distributed object systems. Technical Report UVa CS Technical Report CS-98-12, University of Virginia, Computer Science Dept., June 1998.

[GR96] Jörn Gehring and Alexander Reinefeld. Mars - a framework for minimizing the job execution time in a metacomputing environment. *Future Generation Computer Systems*, Spring 1996.

[Gus90] John Gustafson. Fixed time, tiered memory, and superlinear speedup. In *Proceedings of the Fifth Distributed Memory Computing Conference (DMCC5)*, 1990.

[GW] Lawrence L. Green and Robert Weston. Hpccp cooperation on high speed civil transport. `http://hpccp-www.larc.nasa.gov/~llgreen/hpccp-hsct.html`.

[GWtLt97] Andrew S. Grimshaw, William A. Wulf, and the Legion team. The legion vision of a worldwide virtual computer. *Communications of the ACM*, 40(1), January 1997.

[Han75] Eldon R. Hansen. A generalized interval arithmetic. *Interval Mathematicals, proceedings of the International Symposium*, 1975.

[Han92] Eldon R. Hansen. Global optimization using interval analysis. *Monographs and textbooks in pure and applied mathematics*, 1992.

[HK97] David F. Hegarty and M. Tahar Kechadii. Topology preserving dynamic load balancing for parallel molecular simulations. In *Proceedings of Supercomputing '97*, 1997.

[HM96] Jeffrey K. Hollingsworth and Barton P. Miller. An adaptive cost model for parallel program instrumentation. In *Proceedings of EuroPar'96*, 1996. Also available at `http://www.cs.umd.edu/~hollings/papers/`.

[HS91] Reinhard V. Hanxleden and L. Ridgeway Scott. Load balancing on message passing architectures. *Journal of Parallel and Distributed Computing*, 13, 1991.

[HT] John Helly and John Truong. San diego bay project perspective. `www.sdsc.edu/SDBAY/sdbay.html`.

[HT97] M. Heath and V. Torczon, editors. *Eighth SIAM conference on parallel processing for scientific computing*, 1997.

[IKP96] S. Irani, A. R. Karlin, and S. Phillips. Strongly competitive algorithms for paging with locality of reference. *SIAM Journal on Computing*, 25(3):477–97, June 1996.

[Jam87] Leah Jamieson. *The Characteristics of Parallel Algorithms*, chapter 3, *Characterizing Parallel Algorithms*. MIT Press, 1987. eds. L. Jamieson, D. Gannon and R. Douglass.

[Kar96] John F. Karpovich. Support for object placement in wide area heterogeneous distributed systems. Technical Report CS-96-03, University of Virginia, Department of Computer Science, January 1996.

[Kas96] N. K. Kasabov. *Foundations of neural networks, fuzzy systems and knowledge engineering*. The MIT press, 1996.

[KB97] C. Kesselman and J. Bannister. Qualis: Guaranteed quality of service for metacomputing, 1997. Private communication.

[KME89] A. Kapelnikov, R. R. Muniz, and M. D. Ercegovac. A modeling methodology for the analysis of concurrent systems and computations. *Journal of Parallel and Distributed Computing*, 6:568–597, 1989.

[KMMO94] A. R. Karlin, M. S. Manasse, L. A. McGeoch, and S. Owicki. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6):542–71, June 1994.

[Koh97] James Arthur Kohl. Personal communication, 1997.

[KRB83] Samuel Kotz, Campbell B. Read, and David L. Banks, editors. *Encyclopedia of Statistical Sciences*. John Wiley & Sons, 1983.

[LG97] Richard N. Lagerstrom and Stephan K. Gipp. Psched: Political scheduling on the cray t3e. In *Proceedings of the Job Scheduling Strategies for Parallel Processing Workshop, IPPS '97*, 1997.

[LH95] Johannes Lüthi and Günter Haring. Mean value analysis for queuing network models with intervals as input parameters. Technical Report TR-950101, Institute of applied science and information systems, university of Vienna, July 1995.

[LH97] Johannes Lüthi and Günter Haring. Fuzzy queuing network models of computer systems. In *Proceedings of the 13th UK Performance Engineering Workshop*, 1997.

[LLKS85] Lawler, Lenstra, Kan, and Shmoys. *The Traveling Salesman Problem.* John Wiley & Sons, 1985.

[LM86] Richard J. Larsen and Morris L. Marx. *An Introduction to Mathematical Statistics and Its Applications.* Prentice-Hall, 1986.

[LMH96] Johannes Lüthi, Shikharesh Majumdar, and Günter Haring. Mean value analysis for computer systems with variabilities in workload. In *Proceedings of the IEEE International Computer Performance and Dependability Symposium, IPDS '96*, 1996.

[LMS+98] B. Lowekamp, N. Miller, D. Sutherland, T. Gross, P. Steenkiste, and J. Subhlok. A resource query interface for network-aware applications". In *Proceedings of the 7th IEEE Smposium on High-Performance Distributed Computing*, 1998.

[LS93] S. Leuttenegger and X. Sun. Distributed computing feasibility in a non-dedicated homogeneous distributed system. Technical Report Technical Report Number 93-65, NASA - ICASE, 1993.

[Lue98] Johannes Luethi. Histogram-based characterization of workload parameters and its consequences on model analysis. In *Proceedings of the Workshop on Workload Characterization in High-Performance Computing Environments, jointly with MASCOTS'98*, 1998.

[MBC84] M. Ajmone Marson, G. Balbo, and G. Conte. A class of generalized stochastic petri nets for the performance analysis of multiprocessor systems. *ACM TOCS*, pages 93–122, May 1984.

[MFS+95] C. R. Mechoso, J. D. Farrara, J. A. Spahr, K. Sklower, and M. Stonebraker. Development of an earth system model: A progress report. Technical report, Univ. of California, Los Angeles, April 1995.

[ML90] V. W. Mak and S. F. Lundstrom. Predicting the performance of parallel computations. *IEEE Transactions on Parallel and Distributed Systems*, pages 257–270, July 1990.

[MLH95] Shikharesh Majumdar, Johannes Lüthi, and Günter Haring. Histogram-based performance analysis for computer systems with variabilities or uncertainties in workload. Technical Report Research Report SCE-95-22, Department of systems and computer engineering, Carleton University, Ottawa, Canada, November 1995.

[MMFM93] C. R. Mechoso, C. Ma, J. D. Farrara, and R. W. Moore. Parallelization and distribution of a coupled ocean-atmosphere general circulation model. *Monthly Weather Review*, 121(7), July 1993.

[Moh84] J. Mohan. *Performance of Parallel Programs: Model and Analyses.* PhD thesis, Carnegie Mellon University, July 1984.

[Mol90] M.K. Molloy. Performance analysis using stochastic petri nets. *IEEE Transactions on Parallel and Distributed Systems*, C-31:913–7, September 1990.

[mp] Solaris SunOS 5.5 man page. dis - object code disassembler.

[MR95] Shikharash Majumdar and Ravathy Ramadoss. Interval-based performance analysis of computing systems. In *Proceedings of the Third International Workshop on Model Analysis and Simulation of Computer and Telecommunication Systems*, 1995.

[MR98] Celso L. Mendes and Daniel A. Reed. Integrated compilation and scalability analysis for parallel systems. In *Proceedings of PACT '98*, 1998.

[MRW92] A. Malony, D. Reed, and H. Wijshoff. Performance measurement intrusion and perturbation analysis. *IEEE Transactions on Parallel and Distributed Systems*, 3(4), July 1992.

[MSW81] William Mendenhall, Richard L. Scheaffer, and Dennis D. Wackerly. *Mathematical Statistics with Applications*, pages 4–6. Duxbury Press, 1981.

[NB92] P. Newton and J. C. Browne. The code 2.0 graphical parallel programming language. In *Proceedings of the ACM International Conference on Supercomputing*, July 1992.

[Neu90] Arnold Neumaier. *Interval methods for systems of equations*, chapter Basic Properties of Interval Arithmetic. Cambridge University Press, 1990.

[PM94] Jan Pedersen and John Moult. Determination of the structure of small protein fragments using torsion space monte carlo and genetic algorithm methods. In *Proceedings of Meeting on Critical Assessment of Techniques for Protein Structure Prediction, Asilomar Conference Center*, December 1994.

[PRW94] A. T. Phillips, J. B. Rosen, and V. H. Walke. Molecular structure determination by convex global underestimation of local energy minima. Technical Report Tech Report UMSI 94/126, University of Minnesota Supercomputer Institute, 1994.

[Ric95] John A. Rice. *Mathematical Statistics and Data Analysis, Second Edition*, page 134. Duxbury Press, 1995.

[RSR97] Randy L. Ribler, Huseyin Simitci, and Daniel A. Reed. The autopilot performance-directed adaptive control system. *Future Generation Computer Systems, special issue (Performance Data Mining)*, 1997.

[RSSS98] Emilia Rosti, Giuseppe Seazzi, Evgenia Smirni, and Mark S. Squillante. the impact of i/o on program behavior and parallel scheduling. In *Proceedings of SIGMETRICS*, 1998.

[SB97a] Jennifer Schopf and Francine Berman. Performance prediction using intervals. Technical Report #CS97-541, UCSD, CSE Dept., 1997.

[SB97b] Jennifer M. Schopf and Francine Berman. Performance prediction in production environments. Technical Report #CS97-558, UCSD, CSE Dept., 1997.

[SBSM89] R. Saavedra-Barrera, A. J. Smith, and E. Miya. Machine characterization based on an abstract high-level language machine. *IEEE Transactions on Computers*, 38:1659–1679, December 1989.

[Sch97] Jennifer M. Schopf. Structural prediction models for high-performance distributed applications. In *CCC '97*, 1997. Also available as `www.cs.ucsd.edu/ users/jenny/ CCC97/index.html`.

[Sco79] D.W. Scott. On optimal and data-based histograms. *Biometrika*, 66, 1979.

[Sev89] Kenneth C. Sevcik. Characterizations of parallelism in applications and their use in scheduling. *Performance Evaluation and Review*, 17(1):171–180, May 1989.

[Sil] Silicon Graphics. *Pixie Man Page*. also `http://web.cnam.fr/Docs/man/Ultrix-4.3/pixie.1.html`.

[SK95] Steffen Schulze-Kremer. *Molecular Bioinformatics - Algorithms and Applications*. Walter de Gruyter, Berlin, 1995. Also available at `http://www.techfak.uni-bielefeld.de/bcd/Curric/ProtEn/proten.html`.

[Ske74] Stig Skelboe. Computation of rational integer functions. *BIT*, 14, 1974.

[SL90a] Gilbert C. Sih and Edward A. Lee. Dynamic-level scheduling for heterogeneous processor networks. In *Proceedings of the Second IEEE Symposium on Parallel and Distributed Systems*, 1990.

[SL90b] Gilbert C. Sih and Edward A. Lee. Scheduling to account for inter-processor communication within interconnection-constrained processor networks. In *Proceedings of the 1990 International Conference on Parallel Processing*, 1990.

[SL93] Gilbert C. Sih and Edward A. Lee. A compile-time scheduling heuristic for interconnection-constrained heterogeneous processor architectures. *IEEE Transactions on Parallel and Distributed Systems*, 4(2), 1993.

[SRar] Evgenia Smirni and Daniel A. Reed. workload characterization of input/output intensive parallel applications. *Performance Evaluation*, to appear.

[SSLP93] Jonathan Schaeffer, Duane Szafron, Greg Lobe, and Ian Parsons. The enterprise model of developing distributed applications. *IEEE Parallel and Distributed Technology*, 1(3), August 1993.

[Stu26] H. A. Sturges. *Journal of American Statistical Ass.*, 21, 1926.

[SW96a] Jens Simon and Jens-Michael Wierum. Accurate performance prediction for massively parallel systems and its applications. In *Proceedings of Euro-Par '96 (Second International European Conference on Parallel Processing)*, volume 2, pages 675–88, 1996.

[SW96b] Jens Simon and Jens-Michael Wierum. Performance prediction of benchmark programs for massively parallel architectures. In *Proceedings of the 10th Annual Intl. Conf. on High-Performance Computers, (HPCS 96)*, 1996.

[SW98a] J. Simon and J. M. Wierum. The latency-of-data-access model for analyzing parallel computation. *information processing letters*, 66(5), june 1998.

[SW98b] Neil Spring and Rich Wolski. Application level scheduling of gene sequence comparison on metacomputers. In *Proceedings of the 12th ACM International Conference on Supercomputing*, 1998.

[TB86] A. Thomasian and P. F. Bay. Analysis queuing network models for parallel processing of task systems. *IEEE Transactions on Computers c-35*, 12:1045–1054, December 1986.

[TK93] Howard M. Taylor and Samuel Karlin. *An Introduction to Stochastic Modeling*. Academic Press, 1993.

[Ver] A. Verstraete. Business information systems. `www.smeal.psu.edu/misweb/ infosys/ibismain.html`.

[Wal90] David W. Wall. Predicting program behavior using real or estimated profiles. Technical Report WRL TN-18, Digital Western Research Laboratory, 1990.

[WH94] H. Wabnig and G. Haring. Paps - the parallel program performance prediction toolset. In *Proceedings of the 7th international conference on computer performance evaluation: modeling techniques and tools*, 1994.

[Wil95] Gregory V. Wilson. *Practical Parallel Programming.* The MIT Press, 1995.

[WK93] Yi-Shuen Mark Wu and Aron Kupermann. Prediction of the effect of the geometric phase on product rotational state distributions and integral cross sections. *Chemical Physics Letters*, 201:178–86, January 1993.

[Wol96] Rich Wolski. Dynamically forecasting network performance using the network weather service(to appear in the journal of cluster computing). Technical Report TR-CS96-494, UCSD, CSE Dept., 1996.

[Wol97] R. Wolski. Dynamically forecasting network performance to support dynamic scheduling using the network weather service. In *Proc. 6th IEEE Symp. on High Performance Distributed Computing*, August 1997.

[Wol98] Rich Wolski. Personal communication, 1998.

[WSF89] D. Whitley, T. Starkweather, and D'Ann Fuquay. Scheduling problems and traveling salesman: The genetic edge recombination operator. In *Proceedings of International Conference on Genetic Algorithms*, 1989.

[WSH98] Rich Wolski, Neil Spring, and Jim Hayes. Predicting the cpu availability of time-shared unix systems. In *submitted to SIGMETRICS '99 (also available as UCSD Technical Report Number CS98-602)*, 1998.

[WSP97] R. Wolski, N. Spring, and C. Peterson. Implementing a performance forecasting system for metacomputing: The network weather service. In *SC97*, November 1997.

[WW94] Carl A. Waldspurger and William E. Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proceedings of the First Symposium on Operating System Design and Implementation*, November 1994.

[WZ98] Jon B. Weissman and Xin Zhao. Scheduling parallel applications in distributed networks. *Journal of Cluster Computing*, 1998.

[YB95] W. Young and C. L. Brooks. Dynamic load balancing algorithms for replicated data molecular dynamics. *Journal of Computational Chemistry*, 16:715–722, 1995.

[YZS96] Yong Yan, Xiaodong Zhang, and Yongsheng Song. An effective and practical performance prediction model for parallel computing on non-dedicated heterogeneous now. *Journal of Parallel and Distributed Computing*, October 1996. Also University of Texas, San Antonio, Technical Report # TR-96-0401.

[ZLP96] Mohammed J. Zaki, Wei Li, and Srinivasan Parthasarathy. Customized dynamic load balancing for network of workstations. In *Proceedings of HPDC '96*, 1996.

[ZY95] X. Zhang and Y. Yan. A framework of performance prediction of parallel computing on non-dedicated heterogeneous networks of workstations. In *Proceedings of 1995 International Conference of Parallel Processing*, 1995.