

2

AD-A257 722



DARPA/ONR Quarterly Report

CONTRACT NUMBER: N00014-91-J-1985

CONTRACTOR: Duke University

ARPA ORDER NO: 4331714-01/4-6-88

PROGRAM NAME: N00014

CONTRACT AMOUNT: 696,899.00

EFFECTIVE DATE OF CONTRACT: July 1, 1991

EXPIRATION DATE OF CONTRACT: June 30, 1994

PRINCIPAL INVESTIGATOR: John H. Reif

TELEPHONE NUMBER: (919) 660-6568

SHORT TITLE OF WORK: Implementation of Parallel Algorithms

REPORTING PERIOD: 1 July 1992 — 30 Sept 1992

S DTIC ELECTE D
A DEC 02 1992

This document has been approved for public release and sale; its distribution is unlimited.

The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied of the Defense Advanced Research Projects Agency or the U.S. Government.

92 1 1992

92-30655

 1388

is given which is optimal to within an additive constant of 2 time-steps. The method has applications when the operands are matrices, spread over an array of PEUs, as well. Some lower bounds for this problem, in more general networks, are also proven.

Current work

Our current work in this area is to extend this parallel reduction operation to higher dimensional grids. This work is nearing its successful completion. Robert Wagner and Mike Landis have developed a method for performing reductions on multidimensional processor arrays that is within a few steps of a provably minimal time. They are currently finishing their paper, which is a follow-up to Robert Wagner's paper, "Evaluating Uniform Expressions Within Two Steps of Minimum Parallel Time." This paper solved the problem for two-dimensional arrays only.

Future work

In looking at the problems of distributing collection-oriented operations and collections across many MIMD processors, the question of data-communication cost comes up.

Suppose PEUs are connected in a 2-D grid, with the property that any PE can receive a datum from any one neighbor at a given time-step. (Other network models, and PE communication behavior schemes are also possible.) In this setting, ignoring operation costs, we've studied the problem of computing $\sum v_i$, where each v_i is originally located on a different PE. The goal is to minimize parallel communication time. We have written a TR that solves the problem within 2 steps of optimality regardless of the initial placement of the operands in the grid.

This sort of study opens an area of research, delimited by choosing different combinations of assumptions about PE behavior, network topology, and problems to be solved. A systematic study of these questions, oriented toward problems which are communication-intensive and of interest to people developing packages like LINPAK, seems in order.

(3) Peter Mills (Research Associate) with John Reif: Abstractions and Transformational Implementation for Loosely Synchronous and Asynchronous Parallel Algorithms

Recent research has focused on extension of high-level parallel computation models with abstractions for asynchronous tagged memory communication and distribution of computational resources; development of intermediate representations for these abstractions; and transformation techniques to realize these abstractions on practical parallel machines. We have developed a new parallel programming construct, the rate construct, which constrains relative rates of progress of parallel processes, and specified its semantics as well as evidenced its utility through application to several algorithms. The rate construct attempts to fill a gap in the expression of resource requirements such as computational progress which real-time constructs do not adequately address. The succinct expression of such requirements is needed to explore parallel algorithms at a suitable level of abstraction. These and other wide-spectrum constructs we are developing provide a rich vehicle for expressing parallel algorithms and serve as a concrete carrier for refinement techniques, forming a key component in practical implementation of the parallel algorithms.

This is a top-down approach; we are incorporating these abstractions into our architecture-independent parallel language Proteus and developing refinement techniques targeting such lower-level languages such as the C Vector Library (CMU). At the same time we are investigating extending an existing widely portable data-parallel language, CMU's NESL (supporting nested data parallelism) with a wrapper for asynchronous parallelism built on

shared state and result-parallelism (MultiLisp futures). The intent is to extend and thus capitalizing on their existing techniques for refining nested data parallelism to vector models.

Recent Accomplishments:

(1) Rate Control as a Resource Distribution Abstraction:

The rate construct specifies constraints on the relative rates of progress of tasks executing in parallel, where progress is the amount of computational work as measured by elapsed ticks on a local virtual clock. By prescribing expected work, rates abstractly specify the allocation of processors to tasks needed to achieve that work, effected for example by load balancing. We have developed a simple notion, fixed-rate, which apportions an unvarying percentage of processor time to each process with no assumptions of global clock synchronization between processors. The notion is generalized to encompass distributed systems and time-varying processor workloads by measuring computational work in terms of local virtual clocks updated in user-specified increments by process events.

The utility of the rate construct has been evidenced for a variety of problems, including weighted parallel search for a goal, adaptive many-body simulation in which rates abstract the requirements for load-balancing, and variable time-stepped computations in which the use of fictitious rates can alter the frequency of asynchronous iterations. We are currently investigating means of transforming rate primitives to lower-level real-time and scheduling constructs.

(2) Abstractions for Tagged-memory:

We are integrating the tagged-memory model into the parallel language, Proteus, to serve as a vehicle for our refinement techniques. We have extended our set-theoretic parallel language Proteus with synchronization variables. Synchronization variables, common to coordination languages such as PCN, are a synchronization mechanism in which processes must wait for an referenced uninitialized variable to become defined. We extend this model with other shared-memory abstractions for specifying process topology and directing non-local references to processes. Synchronization variables, in combination with other Proteus features such as barriers for expressing loosely synchronous computations (SPMD and SIMD), prove particularly advantageous in that they can be used to map directly to tagged-memory machines such as the J-machine, or serve as a foundation for implementing primitives on other machines.

Ongoing work:

Development of an MSIMD language extending CMU's nested data-parallel language NESL with a wrapper for asynchronous process execution based on shared variables and MultiLisp futures.

Development of refinement techniques for transforming extended NESL to threads of vector code, targeting such machines as Cray YMP and CM-5.

Demonstration of viability of these techniques through concrete implementations of N-body algorithms, specifically clustering and Fast Multipole Methods, targeting asynchronous collections of SIMD machines.

Extending models of parallel computation with real-time properties, such as processor rates, in order to support timing analysis. Current models, variants of the PRAM such as the APRAM and HPRAM, only accommodate asynchrony of control or hierarchy of control and communication.

(4) Peter Su (postdoc) and John Reif: Implementations of Parallel Algorithms in Computational Geometry

We have been working on the implementational aspects of parallel algorithms. Specifically, we have been studying parallel algorithms for constructing Voronoi Diagrams and related problems. Our interest in this study is not only to build effective algorithms for these problems, but also to consider the kinds of tools that make such work easier and more effective.

Our work has been broken up into three stages:

- (1) Study the theory and practice of conventional algorithms for this problem.
- (2) Study the current body of theoretical work on parallel algorithms for this problem.
- (3) Using the knowledge gained in (a) and (b), design and implement parallel algorithms for this problem on several machines. Then study the performance of these algorithm and how well the theoretical results match the behavior of the implementation.

We have been actively working on stages (1) and (2) for the last few months and we are now ready to move on to stage (3). The study of practical sequential algorithms has been especially helpful in the pursuit of simple and efficient parallel algorithms, since they provide a good set of ideas to extend and refine in a parallel setting.

Using the experience that we gain from this work, we are also investigating and planning tools that could aid the programmer in implementing effective parallel algorithms. Since many parallel algorithms, especially in computational geometry, have similar structure, one could imagine a tool for reasoning about abstract classes of algorithms. In particular, such a system could aid the programmer in tuning performance parameters for specific machines based on architectural characteristics such as global memory bandwidth and latency, processor speed, local memory size, and so on. Also, more basic tools for doing visualization and performance analysis are needed to help the programmer to effective experimental analysis of his implementations. Tools for profiling, animation, simulation and data analysis would all be extremely useful in these settings. At this point, there are no such tools widely available to the research community.

Initial implementations of the ideas in this work has begun and has been successful. I presented a paper at the DAGS conference this summer that describes Cray algorithms for basic proximity problems. I have also begun to explore implementations of the other ideas on various machines, including the MasPar MP-1, the CM-5, and the KSR-1. This development work will make up a large part of my PhD thesis, which should be finished by this spring.

(5) Shenfeng Chen with John Reif: Parallel Sort Implementation

The fastest known sort is a parallel implementation of radix sort in a CRAY, due to CMU's Guy Blelloch. The current sorting algorithms on parallel machines like Cray and CM-2 use radix and bucket sort. But they are not taking advantage of possible distribution of the input keys. We are developing an algorithm using data compression to achieve a fast parallel algorithm which takes this advantage. We expect the new algorithm to beat the previous fastest sort by a few factors. We are working to implement this new parallel sorting algorithm on various parallel machines.

Details:

Radix sort is very efficient when the input keys can be viewed as bits. But the basic radix sort is not distribution based so it needs to look up all digits.

Our approach is to find the structure (distribution) of the input. This is achieved by sampling from the original set. Then a hash table is build from those sample keys. All keys are indexed to buckets separated by consecutive sample keys. A probability analysis shows that the largest set can be bounded within a constant of the average size.

The indexing step is made faster by binary searching the hash table for match. From previous result, each hash function computation needs only constant time.

Our algorithm needs $O(n \log \log n)$ time in sequential given that the compression ratio of the given input set is not too big. In parallel, our algorithm works well in chain-sorting. In list ranking sorting, the total work is also reduced.

We have implemented this algorithm on Sparc II and compared its performance with the system routine quicksort. It turns out that our algorithm outwins the quicksort() for sufficiently large number of keys (32M). Thus, it may find its place in sorting large database operations (e.g., required by joint operations). In these applications the keys are many words long so our algorithm is even more advantageous in this case where the cutoff is much lower.

(6) Deganit Armon (A.B.D.) with John Reif: Parallel Implementation of Nested Dissection

Summary:

We worked on parallel implementation of nested dissection, a numerical method for solving large sparse systems of linear equations, showing three improvements to the known algorithms and implementations. These include a reduction in the memory requirements of the algorithm, a widening of the class of problems solvable with parallel nested dissection (PND) on a mesh-connected processor array, and a reduction in the asymptotic time bounds of PND.

Details:

Nested dissection is a method for solving sparse linear systems of equations by exploiting the graph structure underlying the input matrix.

One of the problems with the known implementations of PND is the large storage requirements of the algorithm; these limit the size of problem for which PND is useful. We show that it is possible to significantly reduce the storage used by an implementation on a processor array to a constant factor of the size of the input matrix. This improvement can be added without affecting the time bounds of the algorithm. We are currently working on an actual implementation.

Using load balancing techniques, we show that PND can be used to solve a larger class of problems on a mesh-connected processor array. In particular, we can use PND to solve any system of equations whose underlying graph is of bounded degree.

We improve on the results of Pan and Reif, who showed that PND can be implemented in $O(\log^3 n)$ time. By taking into account the fact that processors are idle during later stages of the algorithm, several stages can be grouped and performed together to achieve an $O(\log^2 n)$ time algorithm for hypercubes.

An area of ongoing investigation is the implementation of PND on various parallel models. One possible such model that we are looking at is the hierarchical model of parallel memory proposed by Heywood, which may be closer to existing parallel architecture.

(7) Prokash Sinha with John Reif: Randomized Parallel Algorithms for Min Cost Paths

Summary:

We have completed our initial investigation to derive randomized parallel algorithms for Min Cost Paths in a Graph of High Diameter. Our present accomplishment is a randomized sequential algorithm with an order of magnitude performance gain for some dense graphs.

We also found a similar result for PRAM computational model which meets the work we proposed to do in our paper "A Randomized Algorithm for Min Cost Paths in a Graph of High Diameter: Extended Abstract" (J. Reif and P. Sinha). Currently we are in the process of submitting our findings to technical journals and conferences. Our next phase of work would include similar derivations of randomized parallel algorithms for a wide variety of discrete structures which arises naturally in the area of Graph Theory and Combinatorics. Our current research effort is to extend the techniques of Flajolet and Karp to Develop techniques and tools for timing analysis of algorithms. This effort is to derive tools for semiautomatic randomized analysis.

(8) Hongyan Wang with John Reif: Social Potential Fields: A Molecular Dynamics Approach for Distributed Control of Multiple Robots.

Summary:

Much of the early research in robotic planning and control has considered the case of only a single robot. There is now a number of robot systems which include a small number of autonomous robots and consequently there is a quickly growing literature on the planning and cooperative control of systems of small numbers of robots. Our work is concerned with Very Large Scale Robotic (VLSR) systems consisting of at least hundreds to perhaps tens of thousands or more autonomous robots. Our molecular dynamics approach is distributed and robust and flexible.

Details:

We view our VLSR systems as a molecular dynamics system, with predefined force laws between each ordered pair of components (robots, obstacles, objectives and other configurations). These force laws are similar to those found in molecular dynamics, incorporating both attraction and repulsion in the form of inverse power laws. However these laws may differ from molecular systems in that we allow the controller to arbitrarily define distinct laws of attraction and repulsion for separate pairs and groups of robots to reflect their social relations or to achieve some goals. For example, we define a pair-wise force law of attraction and repulsion for a group of identical robots. The repulsion will prevent collision among robots and the attraction will keep them in a cluster. This simulates the phenomena called "individual distance" in sociobiology.

Once the force laws are set up (they can be modified by the global controller), each individual's movement is computed locally according to the local environment sensed by individual robots and the force laws. Thus the control is distributed and robust. Each robots obeys Newton's Law and makes movement complying to the total force on it from the other components.

We give concrete examples to show that this distributed autonomous control will have lots of applications in industry, military and other areas in the future when costs for individual robots drop and robots can be made much more compact and more capable and flexible.

We did computer simulations involving large numbers of robots. Some interesting and useful patterns can be achieved by defining proper force laws for the system, e.g. forming a more or less evenly distributed single cluster, forming a circle to guard a static point particle standing for castle. We are doing more simulations showing more complex patterns.

We also discuss about spring laws similar to molecular bondings to robotic control. Theories of graph rigidity support that we can design a VLSR system which has a rigid structure. This has also applications where assemblies are needed to finish some job efficiently.

(9) Hongyan Wang with John Reif: A Constant Time Algorithm for N-body Simulation with Smooth Distributions.

Summary:

N-body simulation problem is as follows: Given N points that have pair-wise interactions, compute the equilibrium configuration of the N points. This problem is central to a large body of work in theoretical physics, chemistry, and scientific computing, including: cosmology, plasma simulation, molecular dynamics, and fluid mechanics. The fastest N-body simulation algorithm due to Greengard has time complexity of $O(N^2)$ for one step simulation. We propose to use the concept of density function to describe the configuration of the large particle system and a method to compute the equilibrium density function iteratively when given the initial density function in constant time with the time complexity depending only on the potential function and the required precision.

Details:

In a system of large number, say millions of particles, we are interested more in the structure of the system, especially the structure under equilibrium conditions than in the exact positions of all particles. Observations from many fields, such as cosmology, plasma simulations, molecular dynamics, and fluid mechanics, suggest that the distribution of particles is homogeneous and can be described by smooth functions. Thus we propose to use density function to describe the configuration of particle systems.

Based on the fact that under equilibrium conditions, the total force on each particle should equal to 0, we derive an iterative procedure IMPROVE for improving the density function, which is of the form $\phi^{(n+1)}(x) = \text{IMPROVE}(\phi^{(n)}(x))$, where x is a position in the domain of interest. To compute the total force on one robot by summing up discretely all the forces from other robots will require $\Omega(N)$ time. Instead, we only sum up forces from a constant number of nearby particles. For particles far away, we do an integration of force function multiplied by density function to approximate the resultant force. Thus reduce the time complexity to constant. Thus each improvement procedure requires constant time and the number of iterations depends on the required precision thus can be constant.

Simulations showed that the iterative improvement procedures converges. The results showed that in 1-d the density function has a bell-shaped curve and in 2-d has a vault-shaped surface in the domain of interest and outside the domain has 0 value.

(10) Akitoshi Yoshida with John Reif: Image and Video Compression

We considered several compression techniques using optical systems. Optics can offer an alternative approach to overcome the limitations of current compression schemes. We gave a simple optical system for the cosine transform. We designed a new optical vector quantizer system using holographic associative matching and discussed the issues concerning the system.

Optical computing has recently become a very active research field. The advantage of optics is its capability of providing highly parallel operations in a three dimensional space. Image compression suffers from large computational requirements. We propose optical architectures to execute various image compression techniques, utilizing the inherent massive parallelism of optics.

In our paper[RY2], we optically implemented the following compression and corresponding decompression techniques:

- o transform coding
- o vector quantization
- o interframe coding for video

We showed many generally used transform coding methods, for example, the cosine transform, can be implemented by a simple optical system. The transform coding can be carried out in constant time.

Most of this paper is concerned with an innovative optical system for vector quantization using holographic associative matching. Limitations of conventional vector quantization schemes are caused by a large number of sequential searches through a large vector space. Holographic associative matching provided by multiple exposure holograms can offer advantageous techniques for vector quantization based compression schemes. Photo-refractive crystals, which provide high density recording in real time, are used as our holographic media. The reconstruction alphabet can be dynamically constructed through training or stored in the photorefractive crystal in advance. Encoding a new vector can be carried out by holographic associative matching in constant time.

We also discussed an extension of this optical system to interframe coding.

On going work:

We are investigating optical algorithms for video compression.

(1) Computational Geometry by Optical Computers

Some problems require inherently high degrees of interconnections which may not be provided by any conventional electrical computers. The advantage of optical computers is their apparent parallelism in a three dimensional space. Several computational models have been proposed and constructed. As the progress of optical computers continues, there is a great demand in designing and investigating various algorithms that are efficient and appropriate for the proposed models. This situation resembles to the one a decade ago, when various algorithms were investigated for the theoretical VLSI model. Thus, we understand that the investigation on optical computing algorithms will be essential to the development of optical or hybrid massively parallel computers.

Optical techniques are particularly suited for processing images. This leads us to believe that many problems found in computational geometry may be efficiently solved by optical computers. Some researchers have recently started to investigate some basic problems. We have been investigating these and some other problems. We have obtained some new results.

(2) Optical Interconnection

Among processing units placed on a plane, various space-invariant interconnections can be holographically established in constant time. We are investigating appropriate interconnections and efficient algorithms for several problems.

(3) Efficient computation for optical scattering

An efficient algorithm to solve the Helmholtz equations was developed by Rokhlin at Yale. We have been studying his algorithm.

Other areas of research

Area of interest display:

Designing optical methods for eye tracking.

Simulation of Holographic elements

(11) Researchers supported (other than PI):

Salman Azhar, graduate student
Mike Landis, graduate student
Peter Mills, post-doc
Peter Su, visiting graduate student
Robert Wagner, professor
Akitoshi Yoshida, graduate student

(12) Degrees awarded:

Sandeep Sen received his Ph.D. from Duke and was here as a post-doc. Steve Tate and Lars Nyland received their Ph.D.'s in January 1991 under Reif. Tate is remaining at Duke as a post-doc.

(13) Papers

- (1) Optimal Parallel Algorithms for 3 Dimensional Convex Hulls and Related Problems (with S. Sen). *SIAM Journal on Computing*, vol. 21, no. 3, June 1992, pp. 466-485.
- (2) Expected Parallel Time and Sequential Space Complexity of Graph and Digraph Problems (with P. Spirakis). *Algorithmica*, Vol. 7, pp. 597-630, 1992.
- (3) Algebraic Methods for Testing the k-Vertex Connectivity of Directed Graphs, (with J. Cheriyan), *3rd Annual ACM-SIAM Symposium on Discrete Algorithms*, 1992.
- (4) Searching in an Unknown Environment (with M. Kao and S. Tate), *4th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA92)*, San Deigo, CA, 1992.
- (5) An Optimal Space and Efficient Parallel Nested Dissection Algorithm (with D. Armon). *4th Annual ACM Symposium on Parallel Algorithms and Architectures*, San Diego, CA, July 1992.
- (6) Implementations of Randomized Sorting on Large Parallel Machines (with W.L. Hightower and J.F. Prins). *4th Annual ACM Symposium on Parallel Algorithms and Architectures*, San Diego, CA, July 1992.
- (7) Efficient Parallel Algorithms for Computing All Pair Shortest Paths in Directed Graphs (with Y. Han and V. Pan). University of Kentucky Technical Report 204-92. *4th Annual ACM Symposium on Parallel Algorithms and Architectures*, San Diego, CA, July 1992.
- (8) Linear Time Approximate Evaluation of a Polynomial at Real Points (with V. Pan). *33rd Annual IEEE Symposium on Foundations of Computer Science (FOCS92)*, Pittsburgh,PA., Oct 1992.
- (9) The Power of Combining the Techniques of Algebraic and Numerical Computing: Improved Approximate Multipoint Polynomial Evaluation and Improved Multipole Algorithms (with V.Y. Pan and S.R. Tate), *33rd Symposium on Foundations of Computer Science*, October 1992. Also submitted for journal publication as "The Complexity of Trummer's Problem, Zeta Function Evaluation, and N-body Simulation" (with S. Tate).
- (10) Nested Annealing: A Provable Improvement to Simulated Annealing (with S. Rajasekaran). Accepted for publication in *Journal of Theoretical Computer Science*, November 1992.
- (11) On Threshold Circuits and Polynomial Computation. Accepted for publication in *SIAM Journal of Computing*, 1992
- (12) Optical Expanders with Applications in Optical Computing (with A. Yoshida). To appear in *Journal of Applied Optics*, 1992.
- (13) The Computability and Complexity of Ray Tracing (with D. Tygar and A. Yoshida). To appear in *Discrete & Computational Geometry*, 1992.
- (14) Randomized Algorithms for Binary Search and Load Balancing on Fixed Connection Networks with Geometric Applications (with S. Sen). *2nd Annual ACM Symposium on Parallel Algorithms and Architectures*, Crete, Greece, July 1990, pp. 327-337. To appear in *SIAM Journal of Computing*, 1992.

- (15) Strong k -connectivity in Digraphs and Random Digraphs (with P. Spirakis). Accepted for publication in *Algorithmica*, 1992.
- (16) Probabilistic Parallel Prefix Computation. Accepted for publication in *Computers and Mathematics with Applications*, 1992.
- (17) Efficient VLSI Fault Simulation. To appear in *Computers and Mathematics with Applications*, 1992.
- (18) Continuous Alternation: The Complexity of Pursuit in Continuous Domains (with S. Tate). To appear in a special issue of *Algorithmica*, edited by B. Donald, 1992.
- (19) Quad Tree Structures for Image Compression Applications (with T. Markas). Special issue of *Journal of Information Processing and Management*, 1992.
- (20) Derivation of Parallel Graph Connectivity Algorithms via Stream Contraction (with D. Yen), Duke University Technical Report, 1989. Submitted for journal publication.
- (21) Error Resilient One-way Dynamic Communication (with J. A. Storer), October 1990. Submitted for journal publication.
- (22) Derivation of Linear Programming Algorithms (with M. Karr and S. Krishnan), 1990. Submitted for publication. Also as "Derivation of the Ellipsoid Algorithm", Duke University Technical Report CS-1991-17.
- (23) Towards Randomized Strongly Polynomial Algorithms for Linear Programming (with S. Krishnan). Duke University Technical Report CS-1991-18. Submitted for journal publication.
- (24) Fast Computations of Vector Quantization Algorithms (with T. Markas), NASA Technical Report TR-91-58, 1991.
- (25) A Method for Deriving Systolic Algorithms (by R.A. Wagner and M.D. Landis), 1992. Submitted for journal publication.
- (26) Evaluating Uniform Expressions Within Two Steps of Minimum Parallel Time (by R.A. Wagner), 1992. Submitted for journal publication.
- (27) Crypto-Complexity Based Models of Efficiency in Capital Markets (with S. Azhar). Submitted for publication. Also as "A Crypto-Complexity Based Proof of Efficiency in Capital Markets" (with S. Azhar), 1992. Submitted for journal publication.
- (28) Fully Dynamic Graph Connectivity in Logarithmic Expected Time (with P. Spirakis and M. Yung), 1992. Submitted for journal publication.
- (29) On Parallel Implementations and Experimentations of Lossless Data Compression Algorithms (with T. Markas), 1992. Submitted for publication.
- (30) Memory-Shared Parallel Architectures for Vector Quantization Algorithms (with T. Markas), 1992. Submitted for journal publication.
- (31) A Randomized Algorithm for Min Cost Paths in a Graph of High Diameter: Extended Abstract (with P. Sinha), 1992. Submitted for publication.

- (32) Fast Algorithms for Closest Point Problems: Practice and Theory (by Peter Su), 1992. Submitted for publication.
- (33) A Fast Sort and Priority Queue for Entropy Bounded Inputs (with Shenfeng Chen), 1992. Submitted for publication.
- (34) Rate Control as a Language Construct for Parallel and Distributed Programming (with Peter Mills), October 1992. Submitted for publication.
- (35) Social Potential Fields: A Molecular Dynamics Approach for Distributed Control of Social Behavior in Robots (with Hongyan Wang), 1992.
- (36) Dynamic Algebraic Algorithms (with S. R. Tate), 1992. Submitted for publication.
- (37) Dynamic Parallel Tree Contraction (with S. R. Tate), 1992. Submitted for publication.
- (38) A Dynamic Separator Algorithms with Applications to Computational Geometry and Nested Dissection" (with D. Armon), 1992. Submitted for publication.
- (39) Using Learning and Difficulty of Prediction to Decrease Computation: A Fast Sort and Priority Queue on Entropy Bounded Inputs (with S. Chen), 1992. Submitted for publication.
- (40) Parallel and Output Sensitive Algorithms for Combinatorial and Linear Algebra Problems (with J. Cheriyan), 1992. Submitted for publication.
- (41) Re-Randomization and Average Case Analysis of Fully Dynamic Graph Algorithms (with P.G. Spirakis and M. Yung), 1992. Submitted for publication.
- (42) The Complexity of N-body Simulation (with S.R. Tate), 1992. Submitted for publication.