Recent Technical Report List
May 1990–April 1991

New York University
Courant Institute of Mathematical Sciences
Computer Science Department

504-R230 STABBING AND RAY SHOOTING IN 3-DIMENSIONAL SPACE
M. Pellegrini, May 1990

In this paper we consider the following problems: given a set T of triangles in 3-space, with $|T| = n$,

a) answer the query " given a line l, does l stab the set of triangles?" (*query problem*).

b) find whether a stabbing line exists for the set of triangles (*existence problem*).

c) Given a ray ρ, which is the first triangle in T hit by ρ?

The following results are shown.

1. There is an $\Omega(n^3)$ lower bound on the descriptive complexity of the set of all stabbers for a set of triangles.
2. The existence problem for triangles on a set of planes with g different plane inclinations can be solved in $O(g^2 n^2 \log n)$ time.
3. The query problem is solvable in quasi-quadratic $O(n^{2+\epsilon})$ preprocessing and storage and logarithmic $O(\log n)$ query time.
4. All stabbing results for triangles extend, with the same asymptotic bounds, to sets of convex polyhedra with total complexity n.
5. Using $O(n^{3+\epsilon})$ preprocessing time and storage we can determine the first triangle hit in a set of disjoint triangles by any given ray ρ in $O(\sqrt{n} \log n)$ worst case query time.
6. If we are given m rays we can answer ray shooting queries in

$$O(m^{5/6-\delta} n^{5/6+5\delta} \log^2 n + m \log^2 n + n \log n \log m)$$

   randomized expected time and $O(m+n)$ space.
7. In time $O((n+m)^{5/3+4\delta})$ it is possible to decide whether two non convex polyhedra of complexity m and n intersect.
8. Given m rays and n axis-oriented boxes we can answer ray shooting queries in randomized expected time $O(m^{3/4-\delta} n^{3/4+3\delta} \log^4 n + m \log^4 n + n \log n \log m)$ and $O(m+n)$ space.

505 LARGE-SCALE OPTIMIZATION OF EIGENVALUES
M. L. Overton, May 1990

Optimization problems involving eigenvalues arise in many applications. Let x be a vector of real parameters and let $A(x)$ be a differentiable symmetric matrix function of x. We consider a particular problem which occurs frequently: the minimization of the maximum eigenvalue of $A(x)$, subject to linear constraints and bounds on x. The eigenvalues of $A(x)$ are not differentiable at points x where they coalesce, so the optimization problem is said to be nonsmooth. Furthermore, it is typically the case that the optimization objective tends to make eigenvalues coalesce at a solution point.

There are three main purposes of the paper. The first is to present a clear and self-contained derivation of the Clarke generalized gradient of the max eigenvalue function in terms of a "dual matrix". The second purpose is to describe a new algorithm, based on the ideas of a previous paper by the author (SIAM J. Matrix Anal. Appl. 9 (1988) 256-268), which is suitable for solving large-scale eigenvalue optimization problems. The algorithm uses a "successive partial linear programming" formulation which should be useful for other large-scale structured nonsmooth optimization problems as well

as large-scale nonlinear programming with a relatively small number of nonlinear constraints. The third purpose is to report on our extensive numerical experience with the new algorithm, solving problems which arise in the following application areas: the optimal design of columns against buckling; the construction of optimal preconditioners for numerical linear equation solvers; the bounding of the Shannon capacity of a graph. We emphasize the role of the dual matrix, whose dimension is equal to the multiplicity of the minimal max eigenvalue. The dual matrix is computed by the optimization algorithm and used for verification of optimality and sensitivity analysis.

506 DOMAIN DECOMPOSITION ALGORITHMS FOR INDEFINITE ELLIPTIC PROBLEMS
X. Cai, O. Widlund, May 1990

Iterative methods for the linear systems of algebraic equations arising from elliptic finite element problems are considered. Methods previously known to work well for positive definite, symmetric problems are extended to certain nonsymmetric problems, which also can have some eigenvalues in the left half plane.

We first consider an additive Schwarz method applied to linear, second order, symmetric or nonsymmetric, indefinite elliptic boundary value problems in two and three dimensions. An alternative linear system, which has the same solution as the original problem, is derived and this system is then solved by using GMRES, an iterative method of conjugate gradient type. In each iteration step, a coarse mesh finite element problem and a number of local problems are solved on small, overlapping subregions into which the original region is subdivided. We show that the rate of convergence is independent of the number of degrees of freedom and the number of local problems if the coarse mesh is fine enough. The performance of the method is illustrated by results of several numerical experiments.

We also consider two other iterative method for solving the same class of elliptic problems in two dimensions. Using an observation of Dryja and Widlund, we show that the rate of convergence of certain iterative substructuring methods deteriorates only quite slowly when the local problems increase in size. A similar result is established for Yserentant's hierarchical basis method.

507 MULTILEVEL ADDITIVE METHODS FOR ELLIPTIC FINITE ELEMENT PROBLEMS
M. Dryja, O. Widlund, May 1990

An additive variant of the Schwarz alternating method is discussed. A general framework is developed, which is useful in the design and analysis of a variety of domain decomposition methods as well as certain multigrid methods. Three methods of this kind are then considered and estimates of their rates of convergence given. One is a Schwarz-type method using several levels of overlapping subregions. The others use multilevel, multigrid-like decompositions of finite element spaces and have previously been considered by Yserentant and Bramble, Pasciak and Xu. Throughout, we work with finite element approximations of linear, self-adjoint, elliptic problems.

508-R231 SELECTING DISTANCES IN THE PLANE
P. Agarwal, B. Aronov, M. Sharir, S. Suri, May 1990

We describe a randomized algorithm for computing the k^{th} smallest distance in a set of n points in the plane, based on the parametric search technique of Megiddo [Me1]. The expected running time of our algorithm is $O(n^{4/3}\log^{8/3} n)$. A deterministic version of our procedure runs in time $O(n^{3/2}\log^{5/2} n)$. Both versions improve the previously best known upper bound of $O(n^{9/5}\log^{4/5} n)$ by Chazelle [Ch]. A simple $O(n\log n)$ time algorithm for computing an approximation of the median distance is also presented.

509-U169 PERFORMANCE ANALYSIS OF CLOCK-REGULATED QUEUES WITH OUTPUT MULTIPLEXING IN THREE DIFFERENT TYPES OF 2 BY 2 CROSSBAR SWITCH ARCHITECTURES
O. Percus, S. Dickey, May 1990

Switches in interconnection networks for highly parallel shared memory computer systems may be implemented with different internal buffer structures. For a 2×2 synchronous switch, previous studies have often assumed a switch composed of two queues, one at each output, each of which has unbounded size and may accept two inputs in one clock cycle. Hardware implementations may actually use simpler queue designs and will have bounded size.

Two additional models for a 2×2 switch using queues that may accept only one input at a time are analyzed. The first uses four queues, one for each input/output pair. The second uses two queues, one at each input. In each case, a multiplexer blocks one queue if two queues desire the same output, making these models more difficult to analyze than the previous model. Maximum bandwidth, expected queue length, expected waiting time, and queue length distribution are presented for

both models, with unbounded queue size and with queue size equal to 1.

Keywords: Interconnection Network, Clock-regulated Queues, Crossbar Switch

510 ON THE OPTIMAL DESIGN OF COLUMNS AGAINST BUCKING

S. Cox, M. L. Overton, June 1990

We establish existence, derive necessary conditions, and construct and test an algorithm for the maximization of a column's Euler buckling load under a variety of boundary conditions over a general class of admissible designs. We prove that symmetric clamped–clamped columns possess a positive first eigenfunction and introduce a symmetric rearrangement that does not decrease the column's buckling load. Our necessary conditions, expressed in the language of Clarke's generalized gradient, subsume those proposed by Olhoff and Rasmussen, Masur, and Seiranian. The work of Olhoff and Rasmussen, Masur, and Seiranian sought to correct the necessary conditions of Tadjbakhsh and Keller who had not foreseen the presence of a multiple least eigenvalue. This remedy has been hampered by Tadjbakhsh and Keller's miscalculation of the buckling loads of their clamped-clamped and clamped–hinged columns. We resolve this issue in the appendix.

In our numerical treatment of the associated finite dimensional optimization problem we build on the work of Overton in devising an efficient means of extracting an ascent direction from the column's least eigenvalue. Owing to its possible multiplicity this is indeed a nonsmooth problem and again the ideas of Clarke are exploited.

511 PROGRAM TRANSFORMATION FOR EFFICIENT DERIVATION OF MULTIPLE SOLUTIONS IN CONCURRENT LOGIC LANGUAGES

N. Markantonatos, June 1990

Concurrent logic languages provide a flexible and powerful vehicle for expressing parallel programs using explicit processes. However, their drastic departure from conventional logic programming with respect to completeness renders them unsuitable for a variety of useful applications involving search. A multiple solution extension to concurrent logic languages appears to successfully obtain the effect of backtracking in a parallel environment, but has been impeded by inefficiency problems. Moreover, the multiple solution subset introduces a new language which is incoherent with the single solution base language.

We propose a multiple solution subset definition that adheres to the base language both syntactically and semantically. Subsequently, we advocate a source-to-source transformational approach for the efficient implementation of the subset. Multiple solution programs are converted at compile-time into equivalent single solution programs that derive all possible solutions into a single list. Alternative solutions are obtained in an eager or lazy fashion as specified by the program. A number of multiple solution program classes that are transformable into efficient single solution programs are identified and the corresponding transformation procedures are presented and further illustrated using a variety of examples.

The techniques employed for the various transformations include partial evaluation, abstract interpretation, continuation-based transformation, layered stream transformation and loop fusion. As a result of such a static transformational methodology, a broad range of multiple solution programs enjoy efficient execution. We believe that our approach forms a definite step towards an efficient multiple solution subset for concurrent logic languages.

512 TIGHT BOUNDS ON THE COMPLEXITY OF THR BOYER-MOORE PATTERN MATCHING ALGORITHM

R. Cole, June 1990

The problem of finding all occurrences of a pattern of length m in a text of length n is considered. It is shown that the Boyer-Moore pattern matching algorithm performs roughly $3n$ comparisons and that this bound is tight up to $O(n/m)$; more precisely, an upper bound of $3n - \frac{n}{m}$ comparisons is shown, as is a lower bound of $3n(1 - o(1))$ comparisons, as $\frac{n}{m} \rightarrow \infty$ and $m \rightarrow \infty$. While the upper bound is somewhat involved, its main elements provide a quite simple proof of a $4n$ upper bound for the same algorithm.

513-R232 IMPROVED COMBINATORIAL BOUNDS AND EFFICIENT TECHNIQUES FOR CERTAIN MOTION PLANNING PROBLEMS WITH THREE DEGREES OF FREEDOM

D. Halperin, M. Sharir, July 1990

We study motion planning problems for several systems with three degrees of freedom. These problems can be rephrased as the problem of analyzing the combinatorial complexity of a single cell in arrangements of certain types of surfaces (actually, surface patches) in 3-dimensional space. The combinatorial complexity of the entire arrangement in each case that we study can be $\Theta(n^3)$ in the worst case and for each such arrangement we obtain a subcubic bound on the total combinatorial complexity of all the 3D cells in the arrangement that contain a portion of the 1D boundary of a surface patch in their closure

(these are called the *interesting* cells); the bound is $O(n^{7/3})$ in the case of arrangements related to the motion planning problem of a so-called telescopic arm moving in the plane among polygonal obstacles with n corners, and $O(n^{5/2})$ in the case of arrangements resulting from the motion planning problem for an L-shaped object in the plane amidst n point obstacles. We also devise an algorithm to compute the interesting cells in the second type of arrangements, whose time complexity is $O(n^{5/2}\log^2 n)$, and an algorithm with running time $O(n^{7/3})$ for the case of a telescopic arm moving among point obstacles, in both cases improving over the best previously known algorithms for these problems, whose time complexity is $O(n^3 \log n)$. Our approach reduces each three-dimensional problem into a collection of problems involving two-dimensional arrangements. To solve these two-dimensional problems we obtain two combinatorial results of independent interest for arrangements in the plane: (i) a tight bound $\Theta(nm^{1/2})$ on the maximum joint combinatorial complexity of m "concave chains" in an arrangement of n pseudo lines, and (ii) an upper bound $O(m^{2/3}n^{2/3} + n\alpha(n))$ on the maximum number of edges of m distinct faces in certain types of arrangements of n pseudo segments, which is within an $\alpha(\cdot)$ factor off the lower bound for this quantity.

514 BEYOND FAIL-STOP: WAIT-FREE SERIALIZABILITY AND RESILIENCY IN THE PRESENCE OF SLOW-DOWN FAILURES

D. Shasha, J. Turek, September l990

Historically, database researchers have dealt with two kinds of process failures: fail-stop failures and malicious failures. Under the fail-stop assumption, processes fail by halting. Such failures are easily detectable. Under the malicious (or Byzantine) failure assumption, processes fail by behaving unpredictably, perhaps as adversaries. Such failures are not necessarily detectable. When system designers discuss fault tolerance, they typically restrict themselves to the problem of handling fail-stop failures only. This paper proposes an intermediate failure model and presents a practical algorithm for handling transactions under this model. The new failure model allows processes to fail by either slowing down or stopping; slow processes may later speed up, continue to proceed slowly, or, (eventually) stop. We call such failures *slow-down* failures. The model does not assume the ability to distinguish among these possibilities, say, by using a timeout mechanism, nor does it assume that it is possible to kill a slow process. Our algorithm, instead, allows for a new process to be dispatched to do the job that had been assigned to a slow process. The problem is that several processes may end up doing the same task and interfere with one another. Our algorithm controls such interference while guaranteeing both serializability and resiliency.

515-R233 PLANAR GEOMETRIC LOCATION PROBLEMS

P. Agarwal, M. Sharir, September 1990

We present an $O(n^2 \log^3 n)$ algorithm for the two-center problem, in which we are given a set S of n points in the plane and wish to find two closed discs whose union contains S so that the larger of the two radii is as small as possible. We also give an $O(n^2 \log^5 n)$ algorithm for solving the two-line center problem, where we want to find two strips that cover S whose maximum width is as small as possible. The best previous solutions of both problems require $O(n^3)$ time.

516-R234 CIRCLE SHOOTING IN A SIMPLE POLYGON

P. Agarwal, M. Sharir, September 1990

Consider the following problem: Given a simple n-gon $\mathcal{P}$, preprocess it so that for a query circle π and a point s on π, one can quickly compute $\Phi(\mathcal{P}, \pi, s)$, the first intersection point between $\mathcal{P}$ and π as we follow π from s in clockwise direction. We show that $\mathcal{P}$ can be preprocessed, in time $O(n \log^3 n)$, into a data structure of size $O(n \log^3 n)$, so that, for a query circle π, $\Phi(\mathcal{P}, \pi, s)$ can be computed in $O(\log^4 n)$ time. We apply the circle shooting algorithm to report all K intersections between a set of m circular arcs and another collection of n circular arcs in time $O((m\sqrt{n}+n\sqrt{m}) \log^{2.5}(m+n)+(K+m+n) \log^4(m+n))$.

517 DOMAIN DECOMPOSITION ALGORITHMS FOR THE PARTIAL DIFFERENTIAL EQUATIONS OF LINEAR ELASTICITY

B. F. Smith, September 1990

The use of the finite element method for elasticity problems results in extremely large, sparse linear systems. Historically these have been solved using direct solvers like Choleski's method. These linear systems are often ill-conditioned and hence require good preconditioners if they are to be solved iteratively. We propose and analyze three new, parallel iterative domain decomposition algorithms for the solution of these linear systems. The algorithms are also useful for other elliptic partial differential equations.

Domain decomposition algorithms are designed to take advantage of a new generation of parallel computers. The domain is decomposed into overlapping or non-overlapping subdomains. The discrete approximation to a partial differential equation is then obtained iteratively by solving problems associated with each subdomain. The algorithms are often accelerated using the conjugate gradient method.

The first new algorithm presented here borrows heavily from multi-level type algorithms. It involves a local change of basis on the interfaces between the substructures to accelerate the convergence. It works well only in two dimensions.

The second algorithm is optimal in that the condition number of the iteration operator is bounded independently of the number of subdomains and unknowns. It uses non-overlapping subdomains, but overlapping regions of the interfaces between subdomains. This is an additive Schwarz algorithm, which works equally well in two or three dimensions.

The third algorithm is designed for problems in three dimensions. It includes a coarse problem associated with the unknowns on the wirebaskets of the subdomains. The new method offers more potential parallelism than previous algorithms proposed for three dimensional problems since it allows for the simultaneous solution of the coarse problem and the local problems.

518 DAG REPRESENTATION AND OPTIMIZATION OF REWRITING
K. Li, September 1990

In all applications of term rewriting systems, computing the normal forms of terms is the fundamental step. In this paper, we propose a directed acyclic graph (dag) representation for term and term rewriting. The rewriting on dags is much more efficient than the rewriting on trees. We design several efficient strategies for rewriting on dags. By dag representation, we can even obtain efficient rewriting strategies for non-left-linear rewriting systems.

519 A DOMAIN DECOMPOSITION ALGORITHM FOR ELLIPTIC PROBLEMS IN THREE DIMENSIONS
B. F. Smith, October 1990

Most domain decomposition algorithms have been developed for problems in two dimensions. One reason for this is the difficulty in devising a satisfactory, easy-to-implement, robust method of providing global communication of information for problems in three dimensions. Several methods that work well in two dimensions do not perform satisfactorily in three dimensions.

A new iterative substructuring algorithm for three dimensions is proposed. It is shown that the condition number of the resulting preconditioned problem is bounded independently of the number of subdomains and that the growth is quadratic in the logarithm of the number of degrees of freedom associated with a subdomain. The condition number is also bounded independently of the jumps in the coefficients of the differential equation between subdomains. The new algorithm also has more potential parallelism than the iterative substructuring methods previously proposed for problems in three dimensions.

520 STABLE PERTURBATIONS OF NONSYMMETRIC MATRICES
J. V. Burke, M. L. Overton, October 1990

A complex matrix is said to be stable if all its eigenvalues have negative real part. Let J be a Jordan block with zeros on the diagonal and ones on the superdiagonal, and consider analytic matrix perturbations of the form $A(\epsilon) = J + \epsilon B + O(\epsilon^2)$, where ϵ is real and positive. A necessary condition on B for the stability of $A(\epsilon)$ on an interval $(0, \epsilon_0)$, and a sufficient condition on B for the existence of such a family $A(\epsilon)$, is (i) Re tr $B \leq 0$; (ii) the sum of the elements on the first subdiagonal of B has nonpositive real part and zero imaginary part; (iii) the sum of the elements on each of the other subdiagonals of B is zero. This result is extended to matrices with any number of nonderogatory eigenvalues on the imaginary axis, and to a stability definition based on the spectral radius. A generalized necessary condition, though not a sufficient condition, applies to arbitrary Jordan structures. The proof of our results uses two important techniques: the Puiseux-Newton diagram and the normal form of Arnold. In the nonderogatory case our main results were obtained by Levantovskii in 1980 using a different proof. Practical implications are discussed.

521-R235 SAGE: A REAL-TIME OPERATING SYSTEM FOR ROBOTIC SUPERVISORY CONTROL
L. Salkind, October 1990

The next generation of robotic applications—computer integrated manufacturing, teleoperation, and mobile autonomous robots—will require far more computer systems support than currently available. In particular, real-time supervisory control systems will be needed to integrate an increasing number of sensors and actuators, as well as to communicate with other computers in a distributed environment.

This thesis describes the design and implementation of SAGE, an operating system built specifically for real-time robotic supervisory control. The SAGE kernel runs on off-the-shelf Motorola 68020 processor boards, and features lightweight processes, virtual memory support, extensible low-overhead synchronization primitives, and real-time communications capabilities.

Because SAGE is one of the first systems build for robotic supervisory control, the thesis focuses on the issues and design tradeoffs that arise in building a supervisory control operating system.

The thesis also describes how SAGE was used to control a number of intelligent devices, including a Utah/MIT hand and a PUMA robot arm. The robotic experiments performed demonstrate that the operating system can be used in real-time supervisory control applications.

522 TOWARD A FULLY INTEGRATED VLSI CAD SYSTEM: FROM CUSTOM TO FULLY AUTOMATIC
Y. You, October 1990

This thesis describes an integrated CAD environment, which is intented to support almost all phases of the VLSI circuit design cycle, from high-level circuit description down to mask generation. Several VLSI CAD tools have been integrated together under the environment, including a multi-level simulator, automatic placement tools, a schematic layout editor, and a UC Berkeley-developed geometry layout editor. The multi-level simulator supports top-down design by allowing circuits whose components are described at different levels to be simulated together. The levels of circuit description currently supported include a variant of C programming language for circuit behavior descriptions, the schematic layout representation, and the Magic layout from which masks for wafer fabrication can be generated. The hardness of charge-sharing modeling problem is resolved, and a new model for it is given.

The schematic layout editor allows designers to specify interconnections among circuit components in a very efficient manner. It separates behavioral descriptions of a circuit from its geometric layout. Designers can have a graphical view of their design, and specify, within this graphical organization, the behavioral description of components at different levels of abstraction. These schematic layouts with different levels of representation can be simulated using the multi-level simulator.

The automatic placement tool presently performs bottom-up iterative improvement, with simulated annealing as its assistant when needed. Interactive graphics interface is provided which allows human intervention on intermediate as well as final layout.

523 EXECUTION OF REGULAR DO LOOPS ON ASYNCHRONOUS MULTIPROCESSORS
P. Ouyang, October 1990

This paper studies issues concerning parallel execution of regular Fortran DO loops on an asynchronous shared-memory multiprocessor, where each iteration is the basic unit to be executed by a single processing element. An iteration is a dependent predecessor of another iteration if execution of the latter iteration has to wait until execution of the former iteration has completed. During the execution of a DO loop, an iteration will pass through four states, namely, idle, pending, ready, and finished states. An iteration is idle if none of its dependent predecessors have completed; an iteration is pending if some of its dependent predecessors have completed, but not all; an iteration is ready if all its dependent predecessors have completed, but itself has not; otherwise, an iteration is finished. In addition, an iteration without any dependent predecessors is called an initial iteration, which can only have ready and finished states. Via describing an execution scheme, this paper studies the characteristics of Fortran DO loops which are related to the efficiency of the execution. Specifically, this paper investigates (1) the number of initial iterations, (2) the maximum number of ready iterations at any instances during the execution, (3) the maximum number of pending iterations at any instances during the execution, (4) a hash function to disperse different pending iterations, and (5) the parallel execution time.

524-R236 DEXTEROUS MANIPULATION: A GEOMETRIC APPROACH
B. Mishra, October 1990

These notes present in a self-contained manner the basic facts about the theory of dexterous manipulation. In particular, we discuss the models of robot hands and analysis and synthesis of multifingered grasps for objects of arbitrary geometry. The techniques exploit an interesting connection between the problems of dexterous manipulation and some classical results in combinatorial geometry.

525 IMPLEMENTATION OF THE SCHNABEL & ESKOW MODIFIED CHOLESKY FACTORIZATION IN THE CONTEXT OF A TRUNCATED-NEWTON OPTIMIZATION METHOD
T. Schlick, October 1990

We report our experience with an application of the new modified Cholesky Factorization of Schnabel & Eskow in the context of nonlinear optimization. We have implemented a modified version of the new factorization for sparse symmetric linear systems, without pivoting, and have incorporated it into the framework of a large-scale truncated Newton minimization package. In our truncated Newton algorithm, we use the preconditioned linear conjugate gradient method to solve iteratively and approximately for the Newton search direction. The MCF is used here to guarantee that the effective preconditioner is positive definite. Details of the implementation are provided, and performance comparisons with the Gill, Murray and

Wright modified Cholesky factorization for sample computational chemistry problems are reported. Our preliminary results demonstrate that the two modified Cholesky factorizations perform quite differently in practice in our truncated Newton context. In particular, a clear difference in timing of the modification (i.e., first nonzero increment) in addition to differences numerical values suggests that the Schnabel & Eskow strategy may work better for more ill-conditioned and indefinite systems while the Gill, Murray and Wright strategy may be more effective for nearly positive-definite systems. As a consequence, the new factorization may be especially useful for minimization of highly nonlinear functions by Newton methods.

526 DETECTING NONDETERMINISM IN SHARED MEMORY PARALLEL PROGRAMS
A. Dinning, November 1990

This thesis addresses the problem of detecting of a specific type of nondeterminism in shared memory parallel programs known as *access anomalies*. An access anomaly occurs when an update to a shared variable X is concurrent with either a read of X or another update of X.

The first part of the work considers dynamic detection of access anomalies. We introduce a new technique called *task recycling* that detects access anomalies "on the fly" by monitoring the program execution. This technique is designed with two goals in mind. The first goal is minimal monitoring overhead. Costs are incurred only at thread create, terminate, and coordinate operations and every time a monitored variable is accessed. Because variable accesses are generally the most frequent operation, the task recycling technique reduces the overhead per variable access to a small constant. The second goal is generality. The task recycling technique is applicable to a wide variety of parallel constructs and all common synchronous and asynchronous coordination primitives. Combined with a protocol for specifying ordering constraints, the method of representing concurrency relationships in task recycling can be extended to detect general race conditions in parallel programs.

The second part of the thesis involves static detection of several types of nondeterminism that makes dynamic anomaly detection inefficient. In particular, the notion of nondeterminism arising from critical section coordination is refined by distinguishing between three types of nondeterminism — *parallel*, *sequential*, and *reference* nondeterminism. The presence of these types of nondeterminism in a program impacts access anomaly detection in two significant ways: (i) how critical section coordination is modeled during anomaly detection, and (ii) the confidence level and complexity of guaranteeing that a program has no access anomalies. In particular, it is shown that access anomalies can be detected efficiently only if a program is parallel, sequential and reference deterministic. Heuristics are presented that make access anomaly detection tractable in the presence of other nondeterminism through a better classification and semantic understanding of a coordination protocol.

527-R237 CIRCULAR VISIBILITY IN A SIMPLE POLYGON FROM A POINT
P. Agarwal, M. Sharir, November 1990

In this paper we define the notion of circular visibility, and present an $O(n \log n)$ algorithm to compute the portion of a given simple polygon $\mathcal{P}$ circularly visible from a fixed point lying inside $\mathcal{P}$. The second problem that we study is the *circle shooting from a fixed-point*, which can be formulated as follows: Given a simple polygon $\mathcal{P}$ and a point s inside $\mathcal{P}$, preprocess $\mathcal{P}$ so that given a circle π passing through s, one can quickly compute the first point of $\mathcal{P}$ intersected by π as we follow π from s in clockwise direction. We show that $\mathcal{P}$ can be preprocessed, in time $O(n \log n)$, into a data structure of linear size, so that a query can be answered in logarithmic time.

528-R238 OFF-LINE DYNAMIC MAINTENANCE OF THE WIDTH OF A PLANAR POINT SET
P. Agarwal, M. Sharir, November 1990

In this paper we present an efficient algorithm for the off-line dynamic maintenance of the width of a planar point set in the following estricted case: We are given a real parameter W and a sequence $\Sigma = (\sigma_1, \ldots, \sigma_n)$ of n insert and delete operations on a set S of points in $\mathbb{R}^2$, initially consisting of n points, and we want to determine whether there is an i such that the width of S after the i^{th} operation is less than or equal to W. Our algorithm runs in time $O(n \log^3 n)$.

529 SUBSTRUCTURING METHODS FOR PARABOLIC PROBLEMS
M. Dryja, November 1990

Domain decomposition methods without overlapping for the approximation of parabolic problems are considered. Two kinds of methods are discussed. In the first method systems of algebraic equations resulting from the approximation on each time level are solved iteratively with a Neumann-Dirichlet preconditioner. The second method is direct and similar to certain iterative methods with a Neumann-Neumann preconditioner. An analysis of convergence of the methods is presented.

530-R239 NEW YORK UNIVERSITY UNDERGRADUATE ROBOTICS, CAM, AI AND VISION LABORATORY

Z. Li, B. Mishra, J. Schwartz, November 1990

Under a newly funded NSF grant, we have initiated a project at NYU, whose primary goal is to create a disseminable, multi-functional and inexpensive laboratory course sequence, aimed at improving the practical skills of undergraduate students specializing in robotics, vision, AI and manufacturing disciplines within a computer science and/or engineering program. The prototype laboratory will be made available to other universities interested in creating a similar educational infrastructure.

We invite response and participation by colleagues in other institutions having similar interest.

531 CUTTING A POLYTOPE

W. Jockusch, N. Prabhu, November 1990

We show that given two vertices of a polytope one cannot in general find a hyperplane containing the vertices, that has two or more facets of the polytope in one closed half-space. Our result refutes a long-standing conjecture.

We prove the result by constructing a 4-dimensional polytope that provides the counter-example. Also, we show that such a cutting hyperplane can be found for each pair of vertices, if the polytope is either simplicial or 3-dimensional.

532 ON TRIANGULATIONS OF THE 3-BALL AND THE SOLID TORUS

G. Bohus, W. Jockusch, C. Lee, N. Prabhu, November 1990

We show that neither the 3-ball nor the solid torus admits a triangulation in which (i) every vertex is on the boundary, and (ii) every tetrahedron has exactly one triangle on the boundary. (Such triangulations are relevant to an unresolved conjecture of Perles.) Our result settles a question posed at the *DIMACS Workshop on Polytopes and Convex Sets.*

533 ON A CONJECTURE OF MICHA PERLES

N. Prabhu, November 1990

We prove a conjecture of Micha Perles concerning simple polytopes, for a subclass that properly contains the duals of stacked and crosspolytopes. As a consequence of a special property of this subclass it also follows that, the entire combinatorial structure of a polytope in the subclass can be recovered from its graph, by applying our results recursively.

534 PHYSICAL IDEALIZATION AS PLAUSIBLE INFERENCE

E. Davis, November 1990

The analysis of physical systems almost always relies on idealized models of the objects involved. Any idealization, however, will be incorrect or insufficiently accurate some of the time. It seems reasonable, therefore, to view a physical idealization as a defeasible inference which can be withdrawn in the presence of contrary evidence. This talk discusses the consequences of such a view.

We focus on examples where a system may or may not go into a state where idealizations are violated, such as dropping a ball near an open switch connected across a battery. We show that:

1. Non-monotonic logics will try to enforce the idealization by supposing that the ball will miss the switch. This anomaly does not seem to be solvable by the kinds of techniques that have been applied to the Yale Shooting Problem, which it superficially resembles. We show that this problem is analogous to anomalies in non-monotonic logic that are time-independent.

2. A probabilistic analysis is possible, but it relies on independence assumptions that are hard to justify in general.

3. For completely specified systems, the rule "If the idealization gives solvable equations, then assumes that it holds" is, in fact, a monotonic system of inferences. It should therefore be possible to characterize this in a purely deductive theory. We show that this is, indeed, possible for simple cases, but can get messy in complex systems.

4. Programs that make physical predictions can avoid these problems by simply avoiding reasoning from the future to the past. Though most current programs observe this restriction, it seems likely that more powerful and general systems will have to violate it, and thus deal with this issue.

5. Finally, we look at dynamic systems where the idealization can be observed at any single instant, but it is inconsistent over extended intervals.

535-R240 COUNTING AND CUTTING CYCLES OF LINES AND RODS IN SPACE

B. Chazelle, H. Edelsbrunner, L. J. Guibas, R. Pollack, R. Seidel, M. Sharir, J. Snoeyink, December 1990

A number of rendering algorithms in computer graphics sort three-dimensional objects by depth and assume that there is no cycle that makes the sorting impossible. One way to resolve the problem caused by cycles is to cut the objects into smaller pieces. In this paper we address the problem of estimating how many such cuts are always sufficient. We also consider a few related algorithmic and combinatorial geometry problems. For example, we demonstrate that n lines in space can be sorted in randomized expected time $O(n^{4/3+\epsilon})$, provided they define no cycle. We also prove an $O(n^{7/4})$ upper bound on the number of points in space so that there are n lines with the property that for each point there are at least three non-coplanar lines that contain it.

536-R241 COMBINATORIAL AND ALGORITHMIC ANALYSIS OF STABBING AND VISIBILITY PROBLEMS IN 3-DIMENSIONAL SPACE

M. Pellegrini, December 1990

Given a set T of triangles in 3-space, with $|T| = n$, let $\mathcal{S}(T)$ be the set of all lines stabbing the set T. The combinatorial descriptive complexity of $\mathcal{S}(T)$ is denoted by $\#\mathcal{S}(T)$. The following questions about $\mathcal{S}(T)$ are considered in this thesis:

a) answer the query " given a line l, is $l \in \mathcal{S}(T)$?" (*query problem*).

b) decide whether $\mathcal{S}(T) \neq \emptyset$ (*existence problem*).

c) Give upper and lower bounds on $\#\mathcal{S}(T)$

The following results are shown in this thesis:

1. There is an $\Omega(n^3)$ lower bound for $\#\mathcal{S}(T)$. Also $\mathcal{S}(T)$ may have $\Omega(n^2)$ connected components.
2. There is an $O(n^{3+\epsilon})$ upper bound on $\#\mathcal{S}(T)$. Within the same time bound it is possible to solve the existence problem.
3. The existence problem for triangles on a set of planes with g different plane inclinations can be solved in $O(g^2 n^2 \log n)$ time.
4. The query problem is solvable in $O(n^{2+\epsilon})$ preprocessing and storage and logarithmic $O(\log n)$ query time.
5. The results 1), 2), 3) and 4) extend, with the same asymptotic bounds, to sets of convex polyhedra with total complexity n.

Given a set T of n *disjoint* triangles, the ray shooting problem for T is the following: preprocess T so to be able to answer queries of the form "Given a ray ρ, does ρ hit any triangle in T?". The following results are shown in this thesis:

1. Using $O(n^{3+\epsilon})$ randomized preprocessing time and storage we can solve ray-shooting queries in $O(\sqrt{n}\log^2 n)$ worst case query time.
2. If we are given $m > n^{7/5}$ rays and n disjoint triangles, we can answer *all* the ray shooting queries in

$$O(m^{5/6-\delta} n^{5/6+5\delta} \log n + m \log n + n \log m)$$

randomized expected time and $O(m+n)$ space, for every $\delta > 0$. The multiplicative constants depend on δ.

3. Given m rays and n axis-oriented boxes we can answer ray shooting queries in randomized expected time $O(m^{3/4-\delta} n^{3/4+3\delta} \log^3 n + m \log^3 n + n \log m)$ and $O(m+n)$ space, for $1/28 < \delta < 1/9$. The multiplicative constants depend on δ.

537-R242 WEAVING PATTERNS OF LINES AND LINE SEGMENTS IN SPACE

J. Pach, R. Pollack, E. Welzl, December 1990

A *weaving* W is a simple arrangement of lines (or line segments) in the plane together with a binary relation specifying which line is "above" the other. A system of lines (or line segments) in 3-space is called a *realization* of W, if its projection into

the plane is W and the "above-below" relations between the lines respect the specifications. Two weavings are equivalent if the underlying arrangements of lines are combinatorially equivalent and the "above-below" relations are the same. An equivalence class of weavings is said to be a *weaving pattern*. A weaving pattern is *realizable* if at least one element of the equivalence class has a 3-dimensional realization. A weaving (pattern) W is called *perfect*, if along each line (line segment) of W, the lines intersecting it are alternately "above" and "below". We prove that (i) a perfect weaving pattern of n lines is realizable if and only if $n \leq 3$, (ii) a perfect m by n weaving pattern of line segments (in a grid-like fashion) is realizable if and only if $\min(m, n) \leq 3$, (iii) if n is sufficiently large then almost all weaving patterns of n lines are nonrealizable.

538 PREDICTING THE BEHAVIOR OF FINITE PRECISION LANCZOS AND CONJUGATE GRADIENT COMPUTATIONS
A. Greenbaum, Z. Strakos, January 1991

It is demonstrated that finite precision Lanczos and conjugate gradient computations for solving a symmetric positive definite linear system $Ax = b$ or computing the eigenvalues of A behave very similarly to the exact algorithms applied to any of a certain class of larger matrices. This class consists of matrices $\hat{A}$ which have many eigenvalues spread throughout tiny intervals about the eigenvalues of A. The width of these intervals is a modest multiple of the machine precision times the norm of A. This analogy appears to hold, provided only that the algorithms are not run for huge numbers of steps. Numerical examples are given to show that many of the phenomena observed in finite precision computations with A can also be observed in the exact algorithms applied to such a matrix $\hat{A}$.

539 The APRAM – THE ROUNDS COMPLEXITY MEASURE AND THE EXPLICIT COSTS OF SYNCHRONIZATION
R. Cole, O. Zajicek, January 1991

This paper studies the explicit costs of synchronization by examining an asynchronous generalization of the PRAM model called the APRAM model. The APRAM model and its associated complexity measure, the rounds complexity, are defined and then illustrated by designing and analyzing two algorithms: a parallel summation algorithm which proceeds along an implicit complete binary tree and a recursive doubling algorithm which proceeds along a linked list. In both cases replacing global synchronization with local synchronization yields algorithms with reduced complexity.

540 FAST PARALLEL ITERATIVE SOLUTION OF POISSON'S AND THE BIHARMONIC EQUATIONS ON IRREGULAR REGIONS
A. Mayo, A. Greenbaum, January 1991

In [Mayo, SIAM Jour. Num. Anal. 21, 1984, 285:299] a method was introduced for solving Poisson's or the biharmonic equation on an irregular region by making use of an integral equation formulation. Because fast solvers were used to extend the solution to an enclosing rectangle, this method avoided many of the standard problems associated with integral equations. The equations that arose were Fredholm integral equations of the second kind with bounded kernels. In this paper we use iterative methods to solve the dense nonsymmetric linear systems arising from the integral equations. Because the matrices are very well-conditioned, conjugate gradient-like methods can be used and will converge very rapidly. The methods are very amenable to vectorization and parallelization, and we describe parallel and vector implementations on shared memory multiprocessors. Numerical experiments are described and results presented for a 3-dimensional interface problem for the Laplacian on a recording head geometry.

541 THE KINEMATICS OF CUTTING SOLID OBJECTS
E. Davis, January 1991

This paper studies how the cutting of one solid object by another can be described in a formal theory. We present two alternative first-order representations for this domain. The first views an object as gradually changing its shape until it is split, at which time the original object ceases to exist and two (or more) new objects come into existence. The second focusses instead on chunks of material which are part of the overall object. A chunk persists with constant shape until some piece of it is cut away, when the chunk ceases to exist. We prove that the two theories are equivalent under ordinary circumstances, and we show that they are sufficient to support some simple commonsense inferences and algorithms.

542-U170 A COMPILER-ASSISTED APPROACH TO SPMD EXECUTION
R. Cytron, J. Lipkis, E. Schonberg, July 1990

Today, two styles of scientific parallel programming prevail. In the SPMD style, all processors execute the same program,

with sequential code executed redundantly and parallel code executed cooperatively. In the fork-join style, a sequential thread of control spawns multiple threads to execute a portion of the code concurrently. In this paper we describe an automatic method for approaching the efficiency of SPMD-style execution for programs written in the more structured fork-join style.

543-U171 ARCHITECTURE AND PERFORMANCE OF PROCESSOR-MEMORY INTERCONNECTION NETWORKS FOR MIMD SHARED MEMORY PARALLEL PROCESSING SYSTEMS
Y-S. Liu, September 1990

In this thesis we study interconnection networks and their switch architectures, the performance of different architectures, under the MIMD shared memory environment, using both simulation and analytical methods.

The networks we study are constructed from a single basic building block, a switch element. A number of alternatives are proposed in designing the switch element. We study the performance of different switches with both analytical methods and extensive simulations. We also propose a switch architecture with multiple hand shaking signals (which is also referred as 2-CTS single switch for a 2 by 2 switch) that gives improved performance.

Various interconnection networks and interested issues related to them are reviewed in the thesis. We propose a new class of interconnection networks called F networks. In comparison to traditional multi-stage network, F networks provide faster communications among nodes within a cluster. Also extra routes available in the F network make the network fault-tolerant.

Under the uniform traffic model, Kruskal & Snir discovered that different stages of buffered omega network experience about the same queueing delay for moderate load. Based on simulations and analysis, Kruskal, Snir & Weiss established a formula for calculating network delays under moderate traffic. From our simulations we discovered that the Kruskal-Snir-Weiss formula holds only for the forward path delay of the networks; the return path delay is actually substantially less than the delay of the forward path. We complete the network performance formula by extending the Kruskal-Snir-Weiss formula to include the return path. We also analyze the network performance under "hot spot" traffic and obtain analytic results on the performance attained.

544-U172 INTRINSIC RELATIONS IN THE STRUCTURE OF LINEAR CONGRUENTIAL GENERATORS MODULO 2^{β}
O. E. Percus, J. K. Percus, January 1991

We show two different general relationships that exist among terms of the linear congruential generator

$$x_{(i}+1) = ax_i + b \; mod \; 2^{\beta}$$

that are separated by powers of two. These two relationships are not equivalent except in one special case which turns out to be Marsaglia's result.

545-R243 COUNTING REAL ZEROS
P. Pedersen, February 1991

This paper presents an n-dimensional generalization of Hermite's theorem for counting real roots of a polynomial using quadratic forms. We solve the problem of counting the number of real solutions of a system of polynomial equations within an algebraic polyhedron in n-dimensional space, where the polynomials are taken to have rational coefficients.

Our algorithm is purely symbolic, which means that it may be used to implement infinite-precision algorithms for arithmetic in the real-algebraic subset of the real numbers. We present algorithms for doing this as an application of the general theory.

Our algorithms are based on resultant theory, both because this theory provides insights into the algorithms, and because it makes possible a comparatively clear complexity analysis which shows the algorithms to be worst-case optimal, i.e., singly exponential in the degree of the polynomials.

546 AN ASYNCHRONOUS PARALLEL ALGORITHM FOR UNDIRECTED GRAPH CONNECTIVITY
R. Cole, O. Zajicek, February 1991

An algorithm for computing the components of an undirected graph in the (asynchronous) APRAM model is given; the algorithm uses $O(n+e)$ processes and $O(\log n)$ rounds.

547-R244 SOME CONSTRUCTIONS IN RINGS OF DIFFERENTIAL POLYNOMIALS
G. Gallo, B. Mishra, March 1991

In the recent years, motivated by the success of various algorithms and techniques in computational algebraic ring theory, the researchers in the Computer Algebra community have turned their attention to the task of extending these ideas to differential ring. Appropriate analogs and generalizations of standard bases, Buchberger's completion algorithm, etc. are however yet to be found.

In this paper, we study the cause of this problem. First, we show that not all differential ideals are recursive. This opens the question of membership problem in a recursively generated differential ideal. However, the differential ideals still permit a notion of H-bases with certain peculiarities. For the cases where such H-bases can be effectively constructed, we devise a particularly simple procedure to test the membership of a differential polynomial in a differential ideal.

548 RANDOMIZED PARALLEL ALGORITHMS FOR TRAPEZOIDAL DIAGRAMS
K. L. Clarkson, R. Cole, R. E. Tarjan, March 1991

We describe randomized parallel CREW PRAM algorithms for building trapezoidal diagrams of line segments in the plane. For general segments, we give an algorithm requiring optimal $O(A+n\log n)$ expected work and optimal $O(\log n)$ time, where A is the number of intersecting pairs of segments. If the segments form a simple chain, we give an algorithm requiring optimal $O(n)$ expected work and $O(\log n \log\log n \log^* n)$ expected time, and a simpler algorithm requiring $O(n\log^* n)$ expected work. The serial algorithm corresponding to the latter is the simplest known algorithm requiring $O(n\log^* n)$ expected operations. For a set of segments forming K chains, we give an algorithm requiring $O(A + n\log^* n + K\log n)$ expected work and $O(\log n \log\log n \log^* n)$ expected time. The parallel time bounds require the assumption that enough processors are available, with processor allocations every $\log n$ steps.

549-R245 SINGULARITY DETECTION AND PROCESSING WITH WAVELETS
S. Mallat, W. L. Hwang, March 1991

Most of a signal information is often found in irregular structures and transient phenomena. We review the mathematical characterization of singularities with Lipschitz exponents. The main theorems that estimate local Lipschitz exponents of functions, from the evolution across scales of their wavelet transform are explained. We then prove that the local maxima of a wavelet transform detect the location of irregular structures and provide numerical procedures to compute their Lipschitz exponents. The wavelet transform of singularities with fast oscillations have a different behavior that we study separately. We show that the size of the oscillations can be measured from the wavelet transform local maxima. It has been shown that one and two-dimensional signals can be reconstructed from the local maxima of their wavelet transform. As an application, we develop an algorithm that removes white noises by discriminating the noise and the signal singularities through an analysis of their wavelet transform maxima. In two-dimensions, the wavelet transform maxima indicate the location of edges in images. We extend the denoising algorithm for image enhancement applications.

550 ON THE SUM OF THE LARGEST EIGENVALUES OF A SYMMETRIC MATRIX
M. L. Overton, R. Womersley, March 1991

The sum of the largest k eigenvalues of a symmetric matrix has a well known extremal property which was given by Ky Fan in 1949. We discuss a simple proof of this property which seems to have been overlooked in the vast literature on the subject and its many generalizations. The key step is the observation, which is neither new nor well known, that the convex hull of the set of projection matrices of rank k is the set of symmetric matrices with eigenvalues between 0 and 1 and summing to k. The connection with the well known Birkhoff theorem on doubly stochastic matrices is also discussed. Our approach provides a very convenient characterization for the subdifferential of the eigenvalue sum, to be described in a subsequent paper.

551 REFERENCE MANUAL FOR ATBE — A TOOL FOR APPROXIMATE TREE PATTERN MATCHING *Alpha Version 1.0*
J. T-L. Wang, K. Jeong, K. Zhang, D. Shasha, March 1991

Approximate-Tree-By-Example (ATBE) is a system designed to support constructing, comparing and querying sets of ordered labeled trees. In such trees, the nodes are labeled and the left-to-right order among siblings is significant. Ordered labeled trees have many applications in vision, pattern recognition, molecular biology, programming compilation and natural language processing.

The system supports two query paradigms:

1. A pattern tree (that may contain variables and don't care symbols) is matched against a set of data trees (containing neither variables nor don't care symbols). This is a generalization of (a) approximate regular expression matching in strings and (b) of exact tree matching with variables.

2. Given two sets of data trees S and T, find (a) the pair of trees s from S and t from T that are closest (or farthest) with respect to a user-definable editing distance metric, and (b) pairs of trees from the two sets that are within some distance. This is a generalization of the relational join operations.

This manual documents the ATBE system and illustrates its use. Several examples taken directly from the complete implementation are discussed in detail.

552 A PRACTICAL METHOD FOR CONSTRUCTING EFFICIENT LALR(k) PARSERS WITH AUTOMATIC ERROR RECOVERY

P. Charles, March 1991

LR parsing is used for a wide range of applications, including compiler construction, automatic code generation, language-specific editors and natural language processing. Currently, however, solutions have not been developed for practical multiple-lookahead parsing, fully-automatic error recovery, and space and time-efficient LR parsing across this wide-range of applications.

We present a practical framework for LR(k) parsing, for $k > 1$. We give an efficient algorithm that incrementally constructs an LALR(k) parser with varying-length lookahead strings, and whose symbols are consulted during parsing only when necessary.

Currently, effective LR error recovery systems require some user intervention. We describe an effective and fully automated syntactic error recovery method for LR(k) parsers. Finally, we present a generally effective method for compressing LR(k) parsing tables.

We have incorporated these innovations into a parser generator system that automatically constructs a production-quality parser with built-in error diagnostics and recovery. We will show examples of its performance on several programming languages.

553 SCALABLE PARALLEL GEOMETRIC HASHING FOR HYPERCUBE SIMD ARCHITECTURES

I. Rigoutsos, R. Hummel, January 1991

Geometric hashing was recently introduced as a new paradigm for model based object recognition [Robotics Lab. Report #206, August 1989]. The geometric hashing algorithm allows one to find instances of model point patterns in a scene, subject to noise, obscuration, and transformation. In this paper we concentrate on the cases of similarity and rigid transformations and examine the parallelizability of the algorithm. We describe two scalable algorithms for hypercube SIMD architectures. A number of important building block algorithms and several variations to the basic approach are discussed. For a single probe with a two point basis, the algorithms have time complexities that are $\mathcal{O}(\log^2 M + \log S \log M)$ and $\mathcal{O}(S + \log M)$ respectively, using Mn^3-processors; M is the number of models in the database, n is the number of points per model, S is the number of feature points in the scene, and the model of computation is a SIMD hypercube. In order to locate an instance of a model, it is necessary to probe over $\mathcal{O}(S^2)$ basis pairs until a basis pair is found where both points belong to a model embedded in the scene. We have implemented both algorithms on the Connection Machine: using the resulting implementations, it is possible to recognize models consisting of patterns of points embedded in scenes, independent of translation, rotation, and scale changes, when there are thousands of models containing approximately 16 points each, with scenes consisting of hundreds of points, where most of the scene points are spurious noise points, and where embedded model points in the scene may be obscured or misplaced. With 1024 models and a scene of 200 points, we can achieve a probe time of 70 milliseconds on a 64K-processor CM-2 model.

Keywords: Hypercubes, Data Parallelism, SIMD Architectures, Histogramming, Model Based Image Understanding, Pattern Recognition, Rotation/Translation/Scale Independence, Geometric Hashing, Inverse Indexing.

554 ON A PARALLEL IMPLEMENTATION OF GEOMETRIC HASHING ON THE CONNECTION MACHINE

I. Rigoutsos, R. Hummel, April 1991

We report on a scalable parallel implementation of geometric hashing on a Connection Machine. The algorithm that is employed was described in [CIMS Report #553, April 1991]. Using the resulting implementation, it is possible to recognize

models consisting of patterns of points embedded in scenes, independent of translation, rotation, and scale changes, when there are thousands of models containing approximately 16 points each, with scenes consisting of hundreds of points, where most of the scene points are spurious noise points, and where embedded model points in the scene may be obscured or misplaced. With 1024 models and a scene of 200 points, the implementation yields an execution time of 70 milliseconds per probe on a 64K-processor CM-2 parallel computer. Most of the execution time is taken in performing histograms of (model, basis-set) records. The algorithm is scalable, yielding an expected execution time that is $\mathcal{O}(\log^2 M + \log S \log M)$ on a Mn^3-processor hypercube-connected SIMD machine such as the Connection Machine. M is the number of models, n is the number of points per model, and S is the number of scene points. We also describe and report on a series of experiments for both the similarity and rigid transformation cases; these experiments provide information about detection and false alarm rates for varying amounts of noise in the input.

555 COMPARING THREE APPROACHES TO TRANSFORMATIONAL PROGRAMMING

K. Läufer, April 1991

Transformational programming is a methodology that intends to formalize the development of programs from problem specifications. Given the recent effort towards the design of a common prototyping system (CPS) for the Ada programming language, transformation systems may be reconsidered as possible components of prototyping systems. This paper examines and evaluates three approaches to transformational programming:

- The Munich CIP project (Computer-aided, Intuition-guided Programming) consists of a strongly typed, wide-spectrum language with user-defined algebraic types and a semi-automatic transformation system that requires user guidance.
- By contrast, "Algorithmics," the work on algebraic specification originating from IFIP WG 2.1, is a pure pencil-and-paper approach to transformational programming. It provides a concise, uniform mathematical notation and includes work on nondeterminism.
- RAPTS (Robert A. Paige's Transformation System) is a fully mechanical system that transforms high-level specifications to C code. The specifications are given in a functional subset of SETL augmented with fixed-point operations.

First, we describe each system in detail and highlight interesting features. Next, we establish a framework of common criteria by which such different transformational systems can be evaluated. Finally, we point out common features and differences of the three systems and compare them.

556 PROGRAMMING WITH STRUCTURES, FUNCTIONS, AND OBJECTS

F. Henglein, K. Läufer, April 1991

We describe program structuring mechanisms for integrating algebraic, functional and object-oriented programming in a single framework. Our language is a statically typed higher-order language with specifications, structures, types, and values, and with universal and existential abstraction over structures, types, and values.

We show that existential types over structures generalize both the necessarily homogeneous type classes of Haskell and the necessarily heterogeneous object classes of object-oriented programming languages such as C++ or Eiffel. Following recent work on ML, we provide separate linguistic mechanisms for reusing specifications and structures. Subtyping is provided in the form of explicit type conversions.

The language mechanisms are introduced by examples to emphasize their pragmatic aspects. We compare them with the mechanisms of XML+, Haskell and Eiffel and give a type-theoretic perspective. These mechanisms have been developed within a larger, ongoing prototyping language design project.

557 ON THE DETECTION OF ROBUST CURVES

R. Cole, U. Vishkin, April 1991

Given m points in the plane and a threshold t, a curve is defined to be robust if at least t points lie on it. Efficient algorithms for detecting robust curves are given; the key contribution is to use randomized sampling. In addition, an approximate version of the problem is introduced. A geometric solution to this problem is given; it too can be enhanced by randomization.

These algorithms are readily generalized to solve the problem of robust curve detection in a scene of curve fragments: given a set of curve segments, a curve σ is defined to be robust if curve segments of total length at least l lie on σ. Again, both an exact and an approximate version of the problem are considered.

The problems and solutions are closely related to the well-investigated Hough Transform technique.