Recent Technical Report List
July 1989 - April 1990

New York University
Courant Institute of Mathematical Sciences
Computer Science Department

457-R204 MULTIFINGER GRASPS FOR FINE MANIPULATION OF SMOOTH OBJECTS
G. Lafferriere, July 1989

We give rigorous proofs of the existence of two and three finger grasps in the presence of arbitrarily small friction for two and three dimensional smooth objects. No convexity of the objects is assumed. We also prove the existence of *finger gaits* for rotating a planar object using three and four fingers. Additional results for smooth convex objects are presented which describe two different finger gaits using four fingers.

458 THE BIDIAGONAL SINGULAR VALUE DECOMPOSITION AND HAMILTONIAN MECHANICS
P. Deift, J. Demmel, L.-C. Li, C. Tomei, July 1989

We consider computing the singular value decomposition of a bidiagonal matrix B. This problem arises in the singular value decomposition of a general matrix, and in the eigenproblem for a symmetric positive definite tridiagonal matrix. We show that if the entries of B are known with high relative accuracy, the singular values and singular vectors of B will be determined to much higher accuracy than the standard perturbation theory suggests. We also show that the algorithm in [Demmel and Kahan] computes the singular vectors as well as the singular values to this accuracy. We also give a Hamiltonian interpretation of the algorithm and use differential equation methods to prove many of the basic facts. The Hamiltonian approach suggests a way to use flows to predict the accumulation of error in other eigenvalue algorithms as well.

459-R205 IMPROVED BOUNDS ON THE NUMBER OF EDGES
B. Aronov, H. Edelsbrunner, L. J. Guibas, M. Sharir, July 1989

We show that the maximum number of edges bounding m distinct faces in an arrangement of n line segments in the plane is $O(m^{2/3}n^{2/3}\log^{1/3}(\frac{n^2}{m}) + n\alpha(n) + n\log m)$. This improves a weaker previous upper bound of Edelsbrunner, Guibas, and Sharir, and lies within a logarithmic factor of the known lower bound $\Omega(m^{2/3}n^{2/3} + n\alpha(n))$. In addition, we show that the number of edges bounding any m distinct faces in an arrangement of n line segments with a total of t intersecting pairs is $O(m^{2/3}t^{1/3}\log^{1/3}(\frac{t}{m}) + n\alpha(\frac{t}{m}) + n\min\{\log m, \log\frac{t}{m}\})$. This bound is within a logarithmic factor of the lower bound $\Omega(m^{2/3}t^{1/3} + n\alpha(\frac{t}{n}))$ demonstrated in this paper.

460-R210 QUANTITATIVE STEINITZ'S THEOREMS WITH APPLICATIONS TO MULTIFINGERED GRASPING
D. Kirkpatrick, S. Rao Kosaraju, B. Mishra, C. Yap, September 1989

We prove the following quantitative form of a classical theorem of Steinitz: Let m be sufficiently large. If the convex hull of a subset S of Euclidean d-space contains a unit ball then there is a subset of S with at most m points whose convex hull contains a solid ball with the same center and having *residual radius*

$$1 - 3d\left(\frac{2d^2}{m}\right)^{\frac{2}{d-1}}.$$

The case $m = 2d$ was first considered by Bárány, Katchalski and Pach (1982). We also show an upper bound on the achievable radius: the residual radius must be less than

$$1 - \frac{1}{13}\left(\frac{2d^2}{m}\right)^{\frac{2}{d-1}}.$$

These results have applications in the problem of computing the so-called *closure grasps* by an m-fingered robot hand. The above quantitative form of Steinitz's theorem gives a notion of *efficiency* for closure grasps. The theorem also gives rise to some new problems in computational geometry. We present some efficient algorithms for these problems, especially in the two dimensional case.

461 SOME DOMAIN DECOMPOSITION ALGORITHMS FOR NONSELFADJOINT ELLIPTIC AND PARABOLIC PARTIAL DIFFERENTIAL EQUATIONS
X. Cai, September 1989

In this dissertation, we consider the solution of linear systems of algebraic equations that arise from finite element problems. We study some additive Schwarz type domain decomposition algorithms for general, not necessarily selfadjoint, linear elliptic and parabolic PDEs. We use the GMRES method to solve the resulting linear systems of equations. In each step of the iteration, a number of smaller linear systems, which correspond to the restriction of the original problem to subregions, are solved instead of the large system. The number of subproblems can be potentially large and these methods are therefore promising for parallel computation. We estimate the rate of convergence of the algorithms. Numerical experiments are also reported.

462-R208 STORING LINE SEGMENTS IN PARTITION TREES
M.H. Overmars, H. Schipper, M. Sharir, August 1989

We design two variants of planar partition trees, called *segment partition trees* and *interval partition trees*, that can be used for storing arbitrarily oriented line segments in the plan in an efficient way. The raw structures use $O(nlogn)$ and $O(n)$ storage, respectively, and their construction time is $O(nlogn)$. In our applications we augment these structures by certain (simple) auxiliary structures, which may increase the storage and preprocessing time by a polylogarithmic factor. It is shown how to use these structures for solving line segment intersection queries, triangle stabbing queries and ray shooting queries in reasonably efficient ways. If we use the conjugation tree of [10] as the underlying partition tree, the query time for all problems is $O(n^{\gamma}logn)$, where $\gamma = log(1+\sqrt{5})-1$. The techniques are fairly simple and easy to understand.

463 DOMAIN DECOMPOSITION AND ITERATIVE REFINEMENT METHODS FOR MIXED FINITE ELEMENT DISCRETISATIONS OF ELLIPTIC PROBLEMS
T.P. Mathew, September 1989

In this thesis, we first study the classical Schwarz alternating method and an additive, more parallel variant of it, known as the additive Schwarz method, applied to solve saddle point linear systems obtained by discretising a saddle point formulation of elliptic Neumann problems. We assume that the discretisation is obtained by using a mixed finite element method, in particular the Raviart-Thomas elements. We prove convergence with a rate independent of the mesh parameter h. We also present results of numerical experiments using these algorithms.

Following that, we study two algorithms to solve problems on iteratively refined meshes, namely the fast adaptive composite grid method (FAC), and the asynchronous fast adaptive composite grid method (AFAC). We give a proof of convergence of both these methods in the mixed finite element case, with a rate of convergence independent of the mesh parameters h_i. For the FAC algorithm, we also present quantitative bounds for the rate of convergence in the case of special discretisations, which show a fast rate of convergence, and one which is also independent of the geometry of the refined meshes, providing they are shape regular.

Finally, we study a Dirichlet-Neumann type algorithm for the mixed finite element case involving two non-overlapping subdomains. We use this as a preconditioner for a reduced Schur complement system obtained by using an algorithm of Glowinski and Wheeler. We prove that the rate of convergence is independent of h, the mesh parameter. We also present quantitative bounds in special geometries, that show a fast rate of convergence.

464 A LINEAR TIME ALGORITHM FOR FINDING AN AMBITUS
B. Mishra, R.E. Tarjan, August 1989

We devise a linear time algorithm for finding an ambitus in an undirected graph. An ambitus is a cycle in a graph containing two distinguished vertices such that certain different groups of bridges (called B^P-, B^Q- and B^{PQ}-bridges) satisfy the property that a bridge in one group does not interlace with any bridge in the other groups. Thus, an ambitus allows the graph to be cut into pieces, where, in each piece, certain graph properties may be investigated independently and recursively, and then the pieces can be pasted together to yield information about these graph properties in the original graph. In order to achieve a good time-complexity for such an algorithm employing the divide-and-conquer paradigm, it is necessary to find an ambitus *quickly*. We also show that, using ambitus, linear-time algorithms can be devised for abiding-path-finding and nonseparable-induced-cycle-finding problems.

465 AUTOMATED INVERSION OF A UNIFICATION PARSER INTO A UNIFICATION GENERATOR
T. Strzalkowski, September 1989

This paper describes an algorithm for an automated conversion of a unification parser for natural language into an efficient unification generator for the same language. The scope of applicability of the algorithm is discussed and possible avenues of extension are suggested. The algorithm is tested on an actual grammar derived from a larger string grammar for English.

466-R211 DECISION IN MOTION PLANNING IS EASIER THAN PATH CALCULATION
S. Sifrony, September 1989

We show in this note that for translating a rod B within a 3-D domain, amidst n horizontal plates, we can decide in roughly linear time whether there exists a path between two given placements $Psub1$ and $Psub2$ of B, whereas the length of a piecewise linear path (measured in term of the number of its linear segments) may be quadratic, in the worst-case situation. We thus resolve for the negative a long standing open question in motion planning asking whether the time complexity for deciding the existence of a path is roughly the same as the complexity of calculating the path. Our example is based on that appearing in [Si].

467-R213 ON THE ERROR ANALYSIS OF 'GEOMETRIC HASHING'
Y. Lamdan, H. J. Wolfson, October 1989

We present an experimental and theoretical analysis of the noise model in the *Geometric Hashing* technique. In particular the accuracy of the technique is discussed in the cases of affine transformations, similarities and rigid motions. The efficacy of the voting procedure in *Geometric Hashing* is discussed and it is shown that its introduction significantly reduces the expected burden of the verification stage. The discussion is illustrated by results of numerous simulations which have been performed to test the expected performance of the algorithm.

468 JACOBI'S METHOD IS MORE ACCURATE THAN QR
J. Demmel, K. Veselić, October 1989

We show that Jacobi's method (with a modified stopping criterion) computes small eigenvalues of symmetric positive definite matrices with uniformly higher relative accuracy than QR, divide and conquer, traditional bisection, or any algorithm which first involves tridiagonalizing the matrix. If fact, modulo an assumption based on extensive numerical tests, we show that Jacobi's method is optimally accurate in the following sense: if the matrix entries have small relative errors, so that the eigenvalues have certain intrinsic uncertainties, Jacobi will compute them with nearly this accuracy. In other words, as long as the initial matrix has small relative errors in each component, even using infinite precision will not improve on Jacobi (modulo factors of dimensionality). We do prove bisection is optimal in this sense. We also show the eigenvectors are computed more accurately by Jacobi (and inverse iteration) than previously thought possible. We prove similar results for using one-sided Jacobi for the singular value decomposition of a general matrix. We present a version of Jacobi with the property that the more its accuracy exceeds that of QR, the faster it converges.

469 ON THE CONDITIONING OF THE NONSYMMETRIC EIGENPROBLEM: THEORY AND SOFTWARE
Z. Bai, J. Demmel, A. McKenney, October 1989

This report reviews the theory and practical estimation of condition numbers for the nonsymmetric eigenvalue problem. The report provides a manual for using *LAPACK* subroutines `STRSNA` and `STRSEN` to estimate condition numbers for individual eigenvalues and eigenvectors, multiple (or clustered) eigenvalues, and invariant subspaces.

470 ON THE DYNAMIC FINGER CONJECTURE FOR SPLAY TREES. PART I: SPLAY SORTING $\log n$-BLOCK SEQUENCES
R. Cole, B. Mishra, J. Schmidt, A. Siegel, October 1989

A special case of the Dynamic Finger Conjecture is proved; this special case introduces a number of useful techniques.

471 **ON THE DYNAMIC FINGER CONJECTURE FOR SPLAY TREES. PART II: FINGER SEARCHING**
R. Cole, October 1989

The following result is shown: On an n-node splay tree, the amortized cost of a finger access which is at distance d from the previous access is $O(\log d)$. In addition, there is an $O(n \log\log n)$ initialization cost.

472 **EFFICIENT ROBUST PARALLEL COMPUTATIONS**
Z. Kedem, K. Palem, P. Spirakis, November 1989

A parallel computing system becomes increasingly prone to failure as the number of processing elements in it increases. In this paper, we describe a completely general strategy that takes an *arbitrary* step of an ideal CRCW PRAM and *automatically* translates it to run efficiently and robustly on a PRAM in which processors are prone to failure. The strategy relies on efficient robust algorithms for solving a core problem, the *Write-All Problem*. This problem is critical, because, as we show, its complexity is equal to that of any general strategy for realizing robustness in our model. We analyze the expected parallel time and work of various algorithms for solving this problem and prove a lower bound. Our general strategy is also applicable in the context of asynchronous PRAMs given asynchronous algorithms for the Write-All problems.

473-U165 **A DOMAIN DECOMPOSITION ALGORITHM BASED ON A CHANGE TO A HIERARCHICAL BASIS**
B. Smith, O. Widlund, November 1989

Over the last five years, several new fast algorithms have been developed for the solution of elliptic finite element problems with many degrees of freedom. Among them are Yserentant's hierarchical basis multigrid method and a number of domain decomposition algorithms for problems divided into many subproblems. The condition number of the relevant operators, for the best of these preconditioned conjugate gradient methods, grows only slowly with the size of the problems; the growth typically is quadratic in the logarithm of the number of degrees of freedom associated with a subregion or an element of the coarsest mesh.

Important components of many domain decomposition preconditioners are subproblems related to the sets of variables on the interfaces between neighboring subregions. In this paper, we demonstrate that, for problems in two dimensions, a very simple change of basis for these one-dimensional problems leads to a preconditioner which is as effective as those previously considered. In our proof, we use only tools of linear algebra and a result of Yserentant's. Our numerical experiments confirm that the new algorithm is better conditioned than the original hierarchical basis multigrid method.

474 **COMPLEXITY OF TERM REWRITING**
K. Li, November 1989

Term rewriting systems (TRSs) are efficient reduction systems. They find many applications in logic, languages, specifications, proof systems, etc. Rewriting a term to its normal form is a basic procedure. We show that the optimal normal form rewriting problem for terminating and confluent TRSs are NP-complete. If commutative and associative functions are allowed for a terminating and confluent TRS, the optimal normal form rewriting problem is also NP-complete.

475 **APPROXIMATE AND OPTIMAL STRATEGIES FOR NON-AMBIGUOUS AND LEFT LINEAR TERM REWRITING SYSTEMS**
K. Li, November 1989

Term rewriting systems find many applications in logic, languages, specifications, proof systems, etc. Rewriting a term to its normal form is a basic procedure. We classify non-ambiguous and left linear term rewriting systems into three categories: *variable less*, *variable more*, and *variable equal*. We propose and show approximate or optimal strategies for the normal form rewriting for the three categories.

476-R215 **A REAL NEARLY LINEAR ALGORITHM FOR TRANSLATING A CONVEX POLYGON**
S. Sifrony, November 1989

We show in this paper how to translate a convex k-gon in the plane amidst polygonal obstacles having n corners, in *nearly linear* time (and linear space). Specifically, our algorithm takes $O(n(\log^2 k + \log n)\log n + k \log k)$ time. Previous results on this problem ([KS], [KLPS] and [LS]) require at least $\Omega(nk)$ time (and space), so when k is unbounded they are *quadratically*

lower bounded. Our algorithm constructs a "skeletonized" version of the free configuration space that enables to answer any specific motion query in $O(n)$ time (and space). We also show how to translate a *regular* k-gon in time $O(n \log^2 n)$, eliminating thus every k and $\log k$ factor from the time complexity in this special case.

477 LEARNING AS THE EVOLUTION OF REPRESENTATION
P. Caianiello, November 1989

A widely accepted approach to the formal modelling of learning is to view it as a process of constructing high-level structures that allow compression of raw data. The main results of this thesis are a number of precise formulation of this process, and a few algorithms that carry it out.

Our rigorous definitions allow us to identify of a novel concept of structure: the *alphabet.* We propose that a learning mechanism is one which constructs such structures as representations of the given data. The alphabet allows to draw interesting links between theories in pattern recognition and pattern classification.

The structural descriptions constructed by the algorithms are based on the two concepts of *code* and *classification* used, as data compression tools. We provide means for evaluating the efficiency of the encodings created using ideas from information theory. We reduce

The applications discussed include an architecture and a learning rule for neural nets as well as a procedure for grammatical inference. Some of the ideas proposed have been tested on meaningful examples.

478-R216 EFFICIENT ALGORITHMS AND BOUNDS FOR WU-RITT CHARACTERISTIC SETS
G. Gallo, B. Mishra, December 1989

The concept of a characteristic set of an ideal was originally introduced by J.F. Ritt, in the late forties, and later, independently rediscovered by Wu Wen-Tsun, in the late seventies. Since then Wu-Ritt Characteristic Sets have found wide applications in Symbolic Computational Algebra, Automated Theorem Proving in Elementary Geometries and Computer Vision. The original algorithm of Ritt, and subsequent modifications by Wu, has a non-elementary worst-case time complexity. In this paper, we present optimal algorithms for computing a characteristic set with simple-exponential sequential and polynomial parallel time complexities. These algorithms are derived, via linear algebra, from simple-exponential degree bounds for a characteristic set. The degree bounds are obtained by using the recent effective version of Hilbert's Nullstellensatz, due to D. Brownawell and J. Kollar, and a version of Bezout's Inequality, due to J. Heintz.

479-R217 A SIMPLE OUTPUT-SENSITIVE ALGORITHM FOR HIDDEN SURFACE REMOVAL
M. Sharir, M.H. Overmars, December 1989

We derive a simple output-sensitive algorithm for hidden surface removal in a collection of n triangles. If k is the combinatorial complexity of the output *visibility map*, the method runs in time $O(n\sqrt{k}\log n)$. The method is extended to work for other classes of objects as well, sometimes with even improved time bounds. For example, we obtain an algorithm that performs hidden surface removal for n (non-intersecting) balls in time $O(n^{\frac{3}{2}}\log n + k)$.

480 OPTIMAL BEST-MATCH RETRIEVAL
D. Shasha, T. Wang, December 1989

This paper describes a scheme to answer best-match queries from a file containing a collection of objects. A best-match query is to find the objects in the file which are closest (according to some similarity measure) to a given target.

Previous work [4, 22] suggested that one can reduce the number of comparisons required to achieve the desired results using the triangle inequality, starting with a well-formed data structure for the file which reflects some precomputed intra-file distances. We generalize this technique to allow data structures with any kind of topology, that is, we allow an arbitrary set of object pairs in the file with absent distances. A Floyd-Warshall style algorithm is used to achieve the optimal approximation of those unknown distances. A heuristic for choosing objects out of the file to compare against the target is presented. Artificial data and actual protein sequences are used to illustrate our scheme, to demonstrate some weaknesses of the technique, and to show the better performance of our scheme over others when given the same data structure.

Keywords: Matching, file searching, best match, distance metrics, heuristics, lower bounds, upper bounds, topology

481-R218 RANDOMIZED INCREMENTAL CONSTRUCTION OF DELAUNAY AND VORONOI DIAGRAMS
L. Guibas, D. Knuth, M. Sharir, December 1989

We design and analyze a simple randomized incremental algorithm for calculating the Delaunay triangulation of a set of n

points in the plane. The points are inserted one by one in random order, and the triangulation is updated after each insertion. All versions of the triangulation are maintained "on top" of each other, and when a new point is inserted, it traces itself through all past versions of the triangulation to locate its position in the current one. We show that the expected storage of the algorithm is $O(n)$ and the expected running time is $O(n \log n)$. The analysis is fairly simple too.

482 AN OPTIMAL DOMAIN DECOMPOSITION PRECONDITIONER FOR THE FINITE ELEMENT SOLUTION OF LINEAR ELASTICITY PROBLEMS
B. F. Smith, December 1989

For linear elasticity problems the finite element method is an extremely successful method to model complicated structures. The successful implementation requires the solution of very large, sparse, positive definite linear systems of algebraic equations. A new technique for solving these systems using preconditioned conjugate gradients is proposed. Using ideas from both additive Schwarz methods and iterative substructuring methods, we prove that the condition number of the resulting system does not grow as the substructures are made smaller and the mesh is refined. This result holds for two and three dimensions. Numerical experiments have been performed to demonstrate the power of this method. For linear elasticity problems in two dimensions the condition numbers are observed numerically to be less than four when using a regular mesh.

483-R219 COMPLETE SIGNAL REPRESENTATION WITH MULTISCALE EDGES
S.G. Mallat, S. Zhong, December 1989

The sharper variation points of a signal, measured at different scales, can be detected from the local maxima of its wavelet transform. We describe an algorithm that reconstructs one-dimensional signals and images from their sharper variation points at the scales $(2^j)_{j \in Z}$. This algorithm reconstructs exactly images from their multiscale edges. It provides a verification of Marr's conjecture concerning the completeness and stability of multiscale edges. We also prove that the evolution across scales of the wavelet maxima characterizes the local shape of the signal sharp variations. We can thus not only detect edges but also classify them. The wavelet maxima representation is a new reorganization of the image information that enables us to develop algorithms uniquely based on edges for solving image processing problems. We discuss more particularly compact image coding, texture discrimination and pattern recognition. Although we emphasize image processing, this representation also opens new approaches to the processing of many types of one-dimensional and two-dimensional transient signals.

484 USING TEMPORAL LOGIC AND DATALOG TO QUERY DATABASES EVOLVING IN TIME
A. Tuzhilin, Z. Kedem, December 1989

In this paper, we study a query language about databases evolving in (infinite) time. The syntax of the query language is based on a predicate temporal logic. The semantics of the language is defined with infinite sequences of database states, which in this paper are determined either by pure Datalog programs or by negated Datalog programs with inflationary semantics. In general, other mechanisms for defining the semantics, such as production systems, can be used. We analyze the relative expressive power of such a query language and the standard Datalog queries for both pure and negated Datalog programs. We show that our query language has more expressive power than Datalog queries for both pure and negated Datalog programs in general. However, we also prove a surprising technical result that the existential fragment of temporal logic has the same expressive power as Datalog queries for negated Datalog programs with inflationary semantics. This result implies the collapse of the existential fragment of temporal logic for such programs: any temporal logic formula from that fragment can be reduced to an equivalent formula with a single possibility operator.

485 VARIATIONAL ANALYSIS OF SOME CONJUGATE GRADIENT METHODS
D. B. Szyld, O. B. Widlund, December 1989

A number of conjugate gradient methods are considered for a class of linear systems of real algebraic equations. This class includes all symmetric and certain special nonsymmetric problems which give rise to three-term recursions. All the algorithms are characterized variationally. This makes it possible to derive error estimates systematically in terms of certain polynomial approximation problems. Bounds are obtained, which are functions of the extreme eigenvalues of the basic iteration operator. (*Also Duke University Department of Computer Science Report CS-1989-28.*)

486-U167 TOWARDS A UNIFIED THEORY OF DOMAIN DECOMPOSITION ALGORITHMS FOR ELLIPTIC PROBLEMS
M. Dryja, O. B. Widlund, December 1989

A distinction is often made between domain decomposition algorithms for elliptic partial differential equations, which use

overlapping subregions and those which do not. Schwarz's alternating method, the oldest of them all, belongs to the first category. It has recently been discovered that many of the algorithms that belong to the second category can also be regarded as generalizations of the classical Schwarz method or an additive variant thereof. This new approach provides new tools for the analysis and development of domain decomposition algorithms. In this paper, we introduce an abstract additive Schwarz method and develop a simple framework for the analysis of its rate of convergence. We show that it is possible and convenient to specify algorithms in terms of a set of subspaces and related orthogonal projections. This family of algorithms is further expanded by replacing the linear systems of equations, which correspond to the projections, by suitable preconditioners.

We illustrate the usefulness of this approach by considering additive Schwarz methods of a relatively conventional type, iterative substructuring methods, domain decomposition methods defined on the curves or surfaces which subdivide the region, and iterative refinement methods. Throughout we work in a framework of conforming finite elements and self adjoint problems, but we also mention some new results for more general elliptic finite element problems.

487-R220 ARITHMETIC OF REAL ALGEBRAIC NUMBERS IS IN NC
P. Pedersen, January 1990

We describe NC algorithms for doing exact arithmetic with real algebraic numbers in the sign-coded representation introduced by Coste and Roy [CoR 1988]. We present polynomial sized circuits of depth $O(\log^3 N)$ for the monadic operations $-\alpha$, $1/\alpha$, as well as $\alpha + r$, $\alpha \cdot r$, and $sgn(\alpha - r)$, where r is rational and α is real algebraic. We also present polynomial sized circuits of depth $O(\log^7 N)$ for the dyadic operations $\alpha + \beta$, $\alpha \cdot \beta$, and $\mathrm{sgn}(\alpha - \beta)$, where α and β are both real algebraic. Our algorithms employ a strengthened form of the NC polynomial-consistency algorithm of Ben-Or, Kozen, and Reif [BKR 1986].

488 DIAGONAL SCALINGS OF THE LAPLACIAN AS PRECONDITIONERS FOR OTHER ELLIPTIC DIFFERENTIAL OPERATORS
A. Greenbaum, January 1990

We consider the use of diagonal scalings of the Laplacian matrix as preconditioners for matrices arising from other second order self-adjoint elliptic differential operators. It is proved that if a diffusion operator with a piecewise constant but discontinuous diffusion coefficient is preconditioned by a diagonal scaling of the Laplacian, then, in the limit as the mesh size goes to zero, the optimal diagonal scaling is just the identity. This is in contrast to the case in which the diffusion coefficient is smoothly varying, in which case numerical evidence suggests that the optimal diagonal scaling is approximately equal to the square root of the diagonal of the matrix.

489 A PROGRAM ANALYSIS TOOL FOR EVALUATING THE ADA COMPILER VALIDATION SUITE
D. Rennels, E. Schonberg, January 1990

The Ada Compiler Validation Capability (ACVC) is a large collection of programs used to verify that compilers conform to the Ada language standard. The maintenance of the ACVC is complicated by ongoing decisions of the Ada Rapporteur Group concerning possible gaps or ambiguities in the language definition. We describe a tool that is intended to simplify the maintenance and upgrade of the ACVC, by providing a specialized data-base front-end to the test suite. This front-end consists of a pattern-matching system, based on the translator of the Ada/Ed system, and index files into the ACVC, built on primary syntactic features of Ada and on previously defined patterns. The patterns used by the system are compact fragments of Ada programs that contain the features of interest. We describe in detail the pattern language, the index files, and the user interface of the system. We expect the system to be particularly useful when the modifications to Ada that will result from the Ada9X project are used to upgrade the ACVC suite.

490-R221 THE SETL2 PROGRAMMING LANGUAGE
W.K. Snyder, January 1990

The SETL2 programming language is a very high level language based on the theory and notation of finite sets. It is evolved from SETL, developed at New York University by J. T. Schwartz. SETL2 adds to SETL a syntax and name scoping closer to more recent imperative languages, full block structure, and procedures as first class objects.

This document is divided into three parts: First, we give an introduction to SETL2, highlighting its differences from SETL. A description of programming with sets along with a detailed description of SETL is given in Programming With Sets, An Introduction to SETL by Schwartz, Dewar, Dubinsky, and Schonberg, and we strongly encourage the reader to refer to that book if he has no experience with SETL. The second part of this document provides a detailed reference manual of the SETL2 language. The final section is a guide to the operation of the current implementations of SETL2.

491-R222 LINES IN SPACE: COMBINATORICS AND ALGORITHMS
B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir, J. Stolfi, February 1990

Questions about lines in space arise frequently as subproblems in 3-dimensional computational geometry. In this paper we study a number of fundamental combinatorial and algorithmic problems involving arrangements on n lines in 3-dimensional space. Our main results include:

1. A tight $\Theta(n^2)$ bound on the maximum combinatorial description complexity of the set of all oriented lines that have specified orientations relative to the n given lines.
2. A similar bound of $\Theta(n^3)$ for the complexity of the set of all lines passing above the n given lines.
3. A randomized preprocessing procedure using $O(n^{2+\epsilon})$ time and storage, for any $\epsilon > 0$, that builds a structure supporting $O(\log n)$-time queries for testing if a line lies above all the given lines.
4. An $O(n^{4/3+\epsilon})$ randomized expected time algorithm, for any $\epsilon > 0$, that tests the "towering property": do n given red lines lie all above n given blue lines?

The tools used to obtain these results include Plücker coordinates for lines in space and random sampling in geometric problems.

492-R223 REPEATED ANGLES IN THE PLANE AND RELATED PROBLEMS
J. Pach, M. Sharir, February 1990

We show that a set of n points in the plane determine $O(n^2 \log n)$ triples that define the same angle α, and that for many angles α (including $\frac{\pi}{2}$) this bound is tight in the worst case. We also show that, for a broad family of properties $\mathcal{P}$, the number of triangles spanned by the given points and having property $\mathcal{P}$ is $O(n^{7/3})$. Typical such properties are: having a specified area, a specified perimeter, being isosceles, etc.

493-R224 AN IMPROVED TECHNIQUE FOR OUTPUT-SENSITIVE HIDDEN SURFACE REMOVAL
M. H. Overmars, M. Sharir, February 1990

We derive a new output-sensitive algorithm for hidden surface removal in a collection of n triangles, viewed from a point z such that they can be ordered in an acycliv fashion according to their nearness to z. If k is the combinatorial complexity of the output *visibility map*, then we obtain a sophisticated randomized algorithm that runs in time $O(n^{4/3} \log^{2.89} n + k^{3/5} n^{4/5+\delta})$, for any $\delta > 0$. The method is based on a new technique for tracing the visible contours using ray shooting.

494-R225 POINTS AND TRIANGLES IN THE PLANE AND HALVING PLANES IN SPACE
B. Aronov, B. Chazelle, H. Edelsbrunner, L.J. Guibas, M. Sharir, R. Wenger, February 1990

We prove that for any set S of n points in the plane and $n^{3-\alpha}$ triangles spanned by the points of S there exists a point (not necessarily of S) contained in at least $n^{3-3\alpha}/(512 \log^5 n)$ of the triangles. This implies that any set of n points in three–dimensional space defines at most $6.4n^{8/3} \log^{5/3} n$ halving planes.

495 SMARTS – SHARED-MEMORY MULTIPROCESSOR ADA RUN TIME SUPERVISOR
S. Flynn Hummel, February 1990

The programming language Ada is primarily intended for the construction of large-scale and real-time systems. Although the tasking model was mainly aimed at embedded systems, its rich set of tasking features together with its support for programming in the large make Ada increasingly attractive for writing parallel programs for computationally intensive numeric and symbolic applications.

Highly parallel shared-memory machines such as the NYU Ultracomputer have traditionally been regarded as suitable for large-scale scientific code, and not for more symbolic or heterogeneous concurrent applications such as are found in artificial intelligence or real-time programming. However, these applications would benefit greatly from (and even require) the computational power provided by machines with hundreds of processors.

It is therefore desirable to develop Ada programs for highly-parallel machines. The remaining concern is that the cost of managing large numbers of Ada tasks will negate the speedup obtainable from their highly parallel execution. Indeed, a

run-time supervisor for Ada must contend with many potentially expensive serialization points, that is, constructs that may take time proportional to the number of tasks involved.

In this thesis, we describe the design of an Ada run-time supervisor that is free of costly synchronization points. The run-time supervisor SMARTS (Shared-memory Ada Run Time Supervisor) depends on the hardware synchronization primitive *fetch&* Φ, and supports the tasking features of Ada in an efficient highly parallel manner.

We further reduce the overhead of Ada tasking, by micro-tasking, i.e., the explicit scheduling of a family of Ada tasks on a specified number of processors. Thus, Ada tasks are implemented as *lightweight* processes managed by SMARTS, rather than full-blown operating system processes.

Finally, SMARTS implements Ada shared variables efficiently by means of *relay sets*. Relay sets not only provide a means for effectively storing shared variables in a hierarchy of memories, but also for identifying shared variables. Furthermore, relay sets facilitate optimizations wherein rendezvous are executed in the environment of the calling task.

496-R226 DYNAMICS AND OPTIMAL CONTROL OF A LEGGED ROBOT IN FLIGHT PHASE
Z. Li, R. Montgomery, March 1990

It is well known that the angular momentum of a legged robot in the flight phase is conserved. This paper discusses the control of the body orientation in flight using the internal motion of the leg. First, we use the (nonholonomic) angular momentum constraint to recast the problem into a nonholonomic motion planning problem. Then we apply Chow's Theorem to verify that the system is controllable and the concept of holonomy is introduced for constructing an optimal path that does the job. Finally, we apply linearization control in the internal motion space to realize the planned path. This study, hopefully, will provide an additional means of control to dynamically balance a legged robot that ran.

497-R227 AVERAGE TIME FOR BOOLEAN MATRIX MULTIPLICATION
J. Hong, X. Tan, March 1990

This paper presents a new algorithm for Boolean matrix multiplication. By this algorithm, we reduce the average time complexity of multiplying two Boolean matrices of size *ntimesn* to $O(nsup2log\ n)$, with a very small constant, based on an RAM model. In our simulation, for $N = 16$ we run 10,000 random examples. Interestingly enough, the maximum time cost among all these 10,000 different random examples is only fifty percent higher than the average cost.

498 PERFORMANCE OF SHARED MEMORY IN A PARALLEL COMPUTER
K. Donovan, March 1990

Suppose that in a single memory cycle, n independent random accesses are made to m separate memory modules, with each access equally likely to go to any of the memories. Let L_{avg} then represent the expected value of the maximum number of references to a single memory module. Here we show a new method for analyzing this problem. It allows one to efficiently compute narrow upper and lower bounds for L_{avg} as a function of m and n. We also determine the asymptotic behavior of as m and n grow to infinity at a constant ratio $\lambda = n/m$. For any $\lambda > 0$, this paper proves that $L_{avg} = (1 + o(1)) \log m / \log \log m$ as m and $n \to \infty$. An equivalent result was previously obtained by Gonnet in connection with a hashing problem. Using different methods, Gonnet found the asymptotic value of $\Gamma^{-1}(m)$ (plus lower order terms).

499-U168 TIME SERIES TRANSFORMATIONS IN CLOCKED QUEUEING NETWORKS
O. E. Percus, J. K. Percus, March 1990

We consider a network in which message packets arriving on a clock cycle are directed along paths which may have segments in common. When they do, the messages involved are queued and transmitted, one per cycle. The statistics of queue length are investigated by an expansion of the time series in a channel about the regime of independent arrivals. Results for the second stage of such a network are in good agreement with computer simulations.

500-R228 FAT TRIANGLES DETERMINE LINEARLY MANY HOLES
J. Matoušek, J. Pach, Micha Sharir, S. Sifrony, Emo Welzl, April 1990

We show that for every fixed $\delta > 0$ the following holds: If F is a union of n triangles, all of whose angles are at least δ, then the complement of F has $O(n)$ connected components, and the the boundary of F consists of $O(n \log n)$ straight segments. This latter complexity becomes linear if all triangles are of roughly the same size or if they are all infinite wedges.

501-R229 AN ALGORITHM FOR POINT CLUSTERING AND GRID GENERATION
M. Berger, I. Rigoutsos, April 1990

We describe a new point clustering algorithm, and its application to automatic grid generation, a technique used to solve partial differential equations. Algorithms from the computer vision and pattern recognition literature are used to partition points into a set of enclosing rectangles. We show examples from two dimensional calculations, but the algorithm generalizes readily to three dimensions.

502 ANALYZING PARTITION TESTING STRATEGIES
E. Weyuker, B. Jeng, April 1990

In this paper partition testing is assessed analytically. An investigation of what conditions affect the efficacy of partition testing is performed, and comparisons of the fault detection capabilities of partition testing and random testing are made. The effects of subdomain modifications on partition testing's ability to detect faults is also studied.

Index Terms: partition testing, random testing, software testing

503 TREE LOCKING ON CHANGING TREES
V. Lanin, D. Shasha, April 1990

The tree locking protocol is a deadlock-free method of concurrency control defined and verified by Silberschatz and Kedem for data organized in a directed tree. Can the tree protocol work for applications that change the tree? We define a set of three operations capable of changing any tree to any other tree and show that the tree protocol continues to ensure serializability and deadlock-freedom in the presence of these operations.

86-U100 AN OVERVIEW OF THE NYU ULTRACOMPUTER PROJECT
A. Gottlieb, Revised October 1987

The NYU Ultracomputer is a shared memory MIMD parallel computer design to contain thousands of processors connected by an Omega network to a like number of memory modules. A new coordination primitive fetch-and-add is defined and the network is enhanced to combine simultaneous requests, including fetch-and-adds, directed at the same memory location. The present report, an edited amalgam of papers I have coauthored with other members of the Ultracomputer project, discusses our work on architecture, programming models, system software, hardware and VLSI design, performance analysis, and simulation studies.