Recent Technical Report List
February 1989 - June 1989

New York University
Courant Institute of Mathematical Sciences
Computer Science Department

431 OPTIMAL PRECONDITIONERS OF A GIVEN SPARSITY PATTERN

A. Greenbaum, G. Rodrigue, February 1989

We consider the problem of finding the symmetric positive definite preconditioner M of a given form – e.g., having nonzero elements only in specified positions – which minimizes the ratio of the largest to smallest eigenvalue of $M^{-1}A$, for a given symmetric positive definite matrix A. We show how this problem can be expressed as one of minimizing a convex function and how an optimization code can be used to solve the problem numerically. Results are presented showing optimal preconditioners of various sparsity patterns and comparing these to preconditioners that have been proposed in the literature. Several conjectures are made, based on results from the optimization code.

432 MULTIPLE EXTENSION ALGEBRAIC NUMBER FIELDS

C-J. Ho, March 1989; also RR 190

Let α be an algebraic integer satisfying a monic irreducible polynomial $A(t) \in \mathbf{Z}[t]$. The algebraic integers contained in $\mathbf{Q}(\alpha)$ form a ring which we shall denote by $\mathcal{R}_\alpha$. It is known that if $C_\alpha = \mathbf{disc}(A)$ then $\mathcal{R}_\alpha \subseteq \frac{1}{C_\alpha}\mathbf{Z}[\alpha]$, where $\mathbf{disc}(A)$ is the *discriminant* of A. Let β be a root of a prescribed monic irreducible polynomial $B(u) \in \mathbf{Z}[\alpha][u]$. Let the algebraic number field $\mathbf{Q}(\alpha)(\beta)$ be given as the field $\mathbf{Q}(\alpha)$ extended by β. We denote by $\mathcal{R}_{\alpha\beta}$ the algebraic integers contained in $\mathbf{Q}(\alpha)(\beta)$. In this paper, we present an algorithm to compute an integer C_β such that

$$\mathcal{R}_{\alpha\beta} \subseteq \frac{1}{C_\alpha C_\beta}\mathbf{Z}[\alpha][\beta].$$

The algorithm takes $O(\delta_B^3\delta_A \log \delta_A + \delta_A^3)$ arithmetic operations where δ_A and δ_B are the degree of A and B respectively. With this result, we generalize the Lenstra's algorithm [6] for factoring polynomials over algebraic number field. Our algorithm computes the factorization of univariate polynomials over algebraic number field $\mathbf{Q}(\rho_1)(\rho_2)\cdots(\rho_z)$ where ρ_j $(j = 1, \ldots, z)$ is recursively defined as follows: ρ_1 is an algebraic number and ρ_j is algebraic of degree ≥ 2 over $\mathbf{Q}(\rho_1)(\rho_2)\cdots(\rho_{j-1})$. The number of arithmetic operation and the binary length of integers required by this algorithm are polynomial in the degree and the LOG of the coefficients of input polynomials.

433 RAY SHOOTING, IMPLICIT POINT LOCATION, AND RELATED QUERIES IN ARRANGEMENTS OF SEGMENTS

L. Guibas, M. Overmars, M. Sharir, March 1989; also RR 191

We extend and modify recent efficient techniques for half-plane range searching, to obtain algorithms for answering efficiently several types of queries on planar arrangements of segments and certain related structures. For example, we show that, given n segments in the plane, we can preprocess them in $O(n \log n)$ randomized expected time and $O(n)$ space, so that, given a query ray, emanating from an arbitrary point in an arbitrary direction, we can determine the first segment (if any) hit by the ray, in time $O(n^{\frac{2}{3}+\epsilon})$, for any $\epsilon > 0$. Other results involve "implicit" point location in an arrangement of (possibly intersecting) triangles, polygon placement queries, and certain restricted types of ray shooting and hidden surface removal for 3-dimensional polyhedral scenes.

434 COUNTING AND REPORTING INTERSECTIONS IN ARRANGEMENTS OF LINE SEGMENTS

L. Guibas, M. Overmars, M. Sharir, March 1989; also RR 192

We present efficient algorithms for counting and reporting all intersections in an arrangement of n line segments in the plane. Specifically, we have a randomized algorithm for finding all k intersections in such an arrangement in expected time $O(n^{\frac{4}{3}+\delta} + k)$, for any $\delta > 0$, and linear working storage. A variant of the algorithm counts the number of intersections in $O(n^{\frac{4}{3}+\delta})$ randomized expected time, for any $\delta > 0$ and linear space. Our techniques are based on recursive decomposition of the problem into subproblems of smaller size, using plane partition methods that involve random sampling of the given segments, akin to the techniques of [HW] and [Cl].

435 COMPILATION TECHNIQUES FOR VLIW ARCHITECTURES

F. Gasperoni, March 1989

Experiments have shown that reasonable amounts of fine grain parallelism are available globally in serial programs. New compaction algorithms, as well as a new architecture, have been put forward to exploit such parallelism.

Due to its global nature, ordinary local compaction techniques are inadequate to fully extract the fine grain parallelism available in serial programs. Three global compaction algorithms are presented:

1. *Trace scheduling* compacts the most probable execution paths of a program, using some local compaction algorithm. Compensation code is inserted to preserve semantic equivalence with the original program.

2. *Percolation scheduling* is based on a series of elementary semantic-preserving program transformations which are repeatedly applied using higher level guidance rules.

3. *Region scheduling* evenly redistributes the parallelism available throughout a program, taking into account machine capabilities. Parallelism redistribution is again achieved via repetitive application of elementary program transformations.

Global compaction algorithms alone are insufficient to speedup loops effectively, since after compaction loop's iterations are still executed serially. A new compilation technique, *software pipelining* (the software version of hardware pipelining), overcomes this problem by overlapping successive iterations as much as possible. Several software pipelining algorithms are presented, some achieving optimal or quasi-optimal overlapping in the absence of resource constraints.

The idea behind software pipelining algorithms is to (explicitly or implicitly) determine the (fixed or variable) time interval which must be respected before starting the execution of a new iteration. Most of the algorithms create a rigid pipeline, that is iterations are scheduled when the compiler is certain that they will execute without interruptions. More flexible pipelining algorithms are also presented.

The above compilation techniques have been developed for a specific type of architecture termed very long instruction word (VLIW) architecture. VLIWs vaguely resemble horizontal microengines, but their structure is cleaner and their instruction set more orthogonal.

Three sorts of VLIW machines, based on three different computational models are presented. All three computational paradigms allow multiple operations and multiple conditional branches per machine instruction. What distinguishes the models is the way multiple conditional branches are handled.

The first model was specifically designed to accommodate trace scheduling: semantically multiple conditional jumps are organized like in a Lisp `cond` statement.

The second model, used by percolation scheduling, allows multiple conditional jumps to have the structure of an arbitrary decision tree.

The last model extends percolation scheduling by allowing conditional register assignments to be performed within a decision tree. This model is used to generate flexible software pipelines.

436 DATA COMMUNICATION IN ROBOT CONTROL SYSTEMS

D. Clark, March 1989; also RR 193

Robots and robot controllers are becoming more sophisticated. Consequently, the demands on the controller's operating system are increasing. The lower levels of robot control systems (indeed, most real-time control systems) are characterized by *servo loops*. This thesis examines servo loops and how they affect data communications within robot control systems. In the two systems described in this thesis the special characteristics of servo loops are exploited to enhance the data communications.

HIC is an operating system for hierarchies of servo loops. It uses rate monotonic scheduling for the periodic servo loop processes. HIC events (or processes) which are used to implement servo loops are not allowed to block. They will only surrender the processor upon completion or when preempted by a higher priority process. A non-blocking communication structure, Periodic Data Buffers (PDB's) was developed for inter-process communication. HIC has been implemented and is used successfully in a controller for the Utah/MIT hand.

GANGLIA is a proposed real-time communication network. It is intended to allow the processors in a robot controller to be distributed within the robot. Thus the processors can be close to the sensors and actuators they control. Much of the traffic on such a network would be periodic. GANGLIA uses a central controller which allocates access to the network. For the

periodic traffic a fixed schedule, produced off-line, is used. For the aperiodic traffic round-robin polling is used. Unlike most protocols, messages do not contain the address of the destination node. Instead, the messages are labeled with the "name" of its contents. Each node examines each message and decides whether or not it is interested in the message. A special communication controller in each node (the Communication Memory Management Unit) examines and selects the messages. The result of this protocol is a network-wide common memory. In this thesis, the GANGLIA protocol is described in detail and some preliminary analysis of its effectiveness in some real robot systems is given.

437 OUR EXAGMINATION OF WORK IN PROGRESS: TELE AND DEXTEROUS MANIPULATION
B. Mishra, March 1989; also RR 193

We focus on the research issues concerning models, algorithms and applications in tele and dexterous manipulation areas and survey past, present and planned future research work at the Courant Institute of Mathematical Sciences, New York University.

438 SOME DOMAIN DECOMPOSITION ALGORITHMS FOR ELLIPTIC PROBLEMS
M. Dryja, O. Widlund, April 1989; also UCN 155

We discuss domain decomposition methods with which the often very large linear systems of algebraic equations, arising when elliptic problems are discretized by finite differences or finite elements, can be solved with the aid of exact or approximate solvers for the same equations restricted to subregions. The interaction between the subregions, to enforce appropriate continuity requirements, is handled by an iterative method, often a preconditioned conjugate gradient method. Much of the work is local and can be carried out in parallel. We first explore how ideas from structural engineering computations naturally lead to certain matrix splittings. In preparation for the detailed design and analysis of related domain decomposition methods, we then consider the Schwarz alternating algorithm, discovered already in 1869. That algorithm can conveniently be expressed in terms of certain projections. We develop these ideas further and discuss an interesting additive variant of the Schwarz method. This also leads to the development of a general framework, which already has proven quite useful in the study of a variety of domain decomposition methods and certain related algorithms. We demonstrate this by developing several algorithms and by showing how their rates of convergence can be estimated. One of them is a Schwarz-type method, for which the subregions overlap, while the others are so called iterative substructuring methods, where the subregions do not overlap. Compared to previous studies of iterative substructuring methods, our proof is simpler and in one case it can be completed without using a finite element extension theorem. Such a theorem has, to our knowledge, always been used in the previous analysis in all but the very simplest cases.

439 DECISION PROCEDURES FOR ELEMENTARY SUBLANGUAGES OF SET THEORY. XVI. MULTILEVEL SYLLOGISTIC EXTENDED BY SINGLETON, RANK COMPARISON AND UNARY INTERSECTION
D. Cantone, V. Cutello, April 1989

In this paper we show that the satisfiability problem for a sublanguage of set theory involving the unary intersection operator, the singleton operator and the rank comparison predicate is decidable. By unary intersection of a set we mean the intersection of all its members. We also recall that the rank of a set s is the minimum ordinal α such that $s \subseteq \mathcal{V}_\alpha$, where $\mathcal{V}_\alpha$ is the α-th level in the Von Neumann hierarchy of sets defined by

$$\mathcal{V}_0 = \emptyset \text{ and}$$

$$\mathcal{V}_\alpha = \bigcup_{\beta<\alpha} pow(\mathcal{V}_\beta).$$

Many extension of the basic theory MLS (i.e., the class of unquantified formulae in the language $\in, =, \cup, \setminus$) involving various set theoretic operators and/or the rank comparison predicate have already been shown to be decidable

In this paper we introduce the syntactic notion of I-graph and use it to derive decidable conditions for the satisfiability problem of our class of formulae. It turns out that when such conditions are satisfied, a model can be effectively constructed by propagating elements through the I-graph.

440 A DECISION PROCEDURE FOR SET-THEORETIC FORMULAE INVOLVING RANK AND CARDINALITY COMPARISON
D. Cantone, V. Cutello, April 1989

In this paper we consider the satisfiability problem for the class MLSSRC of unquantified set-theoretic formulae in the language $\cup$, $\setminus$, $=$, $\{\cdot\}$, together with the predicates rank and cardinality comparison. We recall that given a set, its rank is a measure of the nesting level of elements within it, whereas its cardinality is the number of its elements.

It is shown that the satisfiability problem for MLSSRC can be reduced to the problem of solving linear systems of integer equations and of verifying a set of combinatorial conditions on objects depending solely on the input formula.

441 COMPUTING EXTERNAL-FARTHEST NEIGHBORS FOR A SIMPLE POLYGON

P.K. Agarwal, A. Aggarwal, B. Aronov, S. R. Kosaraju, B. Schieber, S. Suri, April 1989; also RR 195

Let $\mathcal{P}$ be (the boundary of) a simple polygon with n vertices. For a vertex p of $\mathcal{P}$, let $\phi(p)$ be the set of points on $\mathcal{P}$ that are farthest from p, where the distance between two points is the length of the (Euclidean) shortest path that connects them without intersecting the interior of $\mathcal{P}$. In this paper, we present an $O(n \log n)$ algorithm to compute a member of $\phi(p)$ for every vertex p of $\mathcal{P}$. As a corollary, the external diameter of $\mathcal{P}$ can also be computed in the same time.

442 ON THE OPTIMALITY OF AN ADDITIVE ITERATIVE REFINEMENT METHOD

M. Dryja, O. Widlund, April 1989; also UCN 156

We consider the solution of the linear systems of algebraic equations which arise from elliptic finite element problems defined on composite meshes. Such problems can systematically be constructed by introducing a basic finite element approximation on the entire region and then selecting subregions, and subregions of subregions etc., where the finite element model is further refined in order to gain higher accuracy. We consider conjugate gradient algorithms where, in each iteration, problems representing finite element models on the original region and the subregions prior to further refinement are solved. We can therefore use solvers for problems with uniform or relatively uniform mesh sizes, while the composite mesh can be strongly graded.

The central theoretical issue concerning these so-called iterative refinement algorithms is the design and analysis of algorithms for which the rate of convergence is independent of the number of subproblems, i.e. the number of refinement levels, as well as the mesh sizes. Algorithms for which such bounds can be obtained are called optimal. In this paper, we extend our previous results and show that an additional important algorithm is optimal.

As in our previous studies, we use a basic mathematical framework introduced in a recent study of a variant of Schwarz' alternating algorithm. Some tools from older work on iterative substructuring methods are also used.

443 POLYMORPHIC TYPE INFERENCE AND SEMI-UNIFICATION

F. Henglein, April 1989

In the last ten years declaration-free programming languages with a *polymorphic* typing discipline (ML, B) have been developed to approximate the flexibility and conciseness of dynamically typed languages (LISP, SETL) while retaining the safety and execution efficiency of conventional statically typed languages (Algol68, Pascal). These polymorphic languages can be type checked at compile time, yet allow functions whose arguments range over a variety of types.

We investigate several polymorphic type systems, the most powerful of which, termed Milner-Mycroft Calculus, extends the so-called **let**-polymorphism found in, e.g., ML with a polymorphic typing rule for recursive definitions. We show that semi-unification, the problem of solving inequalities over first-order terms, characterizes type checking in the Milner-Mycroft Calculus to polynomial time, even in the restricted case where nested definitions are disallowed. This permits us to extend some infeasibility results for related combinatorial problems to type inference and to correct several claims and statements in the literature.

We prove the existence of unique most general solutions of term inequalities, called most general semi-unifiers, and present an algorithm for computing them that terminates for all known inputs due To a novel "extended occurs check". We conjecture this algorithm to be uniformly terminating even though, at present, general semi-unification is not known to be decidable. We prove termination of our algorithm for a restricted case of semi-unification that is of independent interest.

Finally, we offer an explanation for the apparent practicality of polymorphic type inference in the face of theoretical intractability results.

444 PROBABILISTIC METHODS: ALGORITHMIC ASPECTS

P. Tetali, February 1989

The basic probabilistic method was first introduced by Paul Erdos in proving lower bounds for Ramsey numbers. The method since then has been a useful tool in extremal graph theory, combinatorics, and more recently in theoretical computer science as well. We review here some of the (basic and non-basic) probabilistic methods. The emphasis of the report is on algorithmic aspects of these methods. Of special interest is the Lovasz local lemma for which no constructive analogs are known.

445 ON k-SETS IN ARRANGEMENTS OF CURVES AND SURFACES

M. Sharir, May 1989; also RR 196

We extend the notion of k-sets and $(\leq k)$-sets, to arrangements of curves and surfaces. In the case of curves in the plane, we assume that each curve is simple and separates the plane. The k-set is the set of all intersection points of the curves which are covered by exactly k interiors of (or halfplanes bounded by) other curves; the $(\leq k)$-set is the union of all j-sets, for $j \leq k$. Adapting the probabilistic analysis technique of Clarkson, we obtain bounds that relate the maximum size of the $(\leq k)$-set to the maximum size of a 0-set of a sample of the curves. Using known bounds on the size of such 0-sets, we obtain asymptotically tight bounds for the maximum size of the $(\leq k)$-set in the following special cases: (i) If each pair of curves intersect at most twice, the maximum size is $\Theta(nk)$. (ii) If the curves are unbounded arcs and each pair of them intersect at most three times, then the maximum size is $\Theta(nk\alpha(n/k))$. (iii) If the curves are x-monotone arcs and each pair of them intersect in at most some fixed number s of points, then the maximum size of the $(\leq k)$-set is $\Theta(k^2\lambda_s(n/k))$, where $\lambda_s(m)$ is the maximum length of (m,s)-Davenport-Schinzel sequences. We also obtain generalizations of these results to certain classes of surfaces in three and higher dimensions. Finally we present various applications of these results to arrangements of segments and curves, high-order Voronoi diagrams, partial stabbing of disjoint convex sets in the plane, and more.

446 PARTITIONING ARRANGEMENTS OF LINES: I. AN EFFICIENT DETERMINISTIC ALGORITHM

P.K. Agarwal, May 1989; also RR 197

In this paper we consider the following problem: Given a set $\mathcal{L}$ of n lines in the plane, partition the plane into $O(r^2)$ triangles so that no triangle meets more than $O(\frac{n}{r})$ lines of $\mathcal{L}$. We present a deterministic algorithm for this problem with $O(nr\log n\log^{\omega} r)$ running time, where ω is a constant < 3.33.

447 PARTITIONING ARRANGEMENTS OF LINES: II. APPLICATIONS

P.K. Agarwal, May 1989; also RR 198

In this paper we present efficient deterministic algorithms for various problems involving lines or segments in the plane, using the partitioning algorithm described in a companion paper [Ag]. These applications include: (i) An $O(m^{2/3}n^{2/3}\cdot\log^{2/3} n\cdot\log^{\omega/3}\frac{m}{\sqrt{n}}+(m+n)\log n)$ algorithm to compute all incidences between m points and n lines, where ω is a constant < 3.33; (ii) An $O(m^{2/3}n^{2/3}\log^{5/3} n\log^{\omega/3}\frac{m}{\sqrt{n}}+(m+n)\log n)$ algorithm to compute m faces in an arrangement of n lines; (iii) An $O(n^{4/3}\log^{(\omega+2)/3} n)$ algorithm to count the number of intersections in a set of n segments; (iv) An $O(n^{4/3}\log^{(\omega+2)/3} n)$ algorithm to count "red-blue" intersections between two sets of segments, and (v) An $O(n^{3/2}\log^{\omega+1} n)$ algorithm to compute spanning trees with low stabbing number for a set of n points. We also present an algorithm that, given set of n points in the plane preprocesses it, in time $O(n\sqrt{m}\log^{\omega+1/2} n)$, into a data structure of size $O(m)$ for $n\log n\leq m\leq n^2$, so that the number of points of S lying inside a query triangle can be computed in $O(\frac{n}{\sqrt{m}}\log^{3/2} n)$ time.

448 RAY SHOOTING AND OTHER APPLICATIONS OF SPANNING TREES WITH LOW STABBING NUMBER

P.K. Agarwal, May 1989; also RR 199

In this paper we consider the following problem: Given a set $\mathcal{G}$ of n (possibly intersecting) line segments in the plane, preprocess it so that, given a query ray ρ emanating from a point p, one can quickly compute the intersection point $\Phi(\mathcal{G},\rho)$ of ρ with a segment of $\mathcal{G}$ that lies nearest to p. We present an algorithm that preprocesses $\mathcal{G}$, in time $O(n^{3/2}\log^{\omega} n)$, into a data structure of size $O(n\alpha(n)\log^4 n)$, so that for a query ray ρ, $\Phi(\mathcal{G},\rho)$ can be computed in time $O(\sqrt{n\alpha(n)}\log^2 n)$, where ω is a constant < 4.33 and $\alpha(n)$ is a functional inverse of Ackermann's function. If the given segments are non-intersecting, the storage goes down to $O(n\log^3 n)$ and the query time becomes $O(\sqrt{n}\log^2 n)$. The main tool that we use is spanning trees (on the set of segment endpoints) with low stabbing number, i.e. with the property that no line intersects more than $O(\sqrt{n})$ edges of the tree. Using such trees we also obtain faster algorithms for several other problems, including implicit point location, polygon containment and implicit hidden surface removal.

449 SET-THEORETIC REDUCTIONS OF HILBERT'S TENTH PROBLEM

D. Cantone, V. Cutello, A. Policriti, May 1989

The problem to set "upper bounds" to the decidable in set theory has been recently investigated in Parlamento-Policriti ("Decision procedures for elementary sublanguages of set theory. IX. Unsolvability of the decision problem for a restricted subclass of the Δ_0-formulas in set theory", *Comm. Pure App. Math.*, XLI:221–251, 1988) where it is shown that Gödel's

first incompleteness theorem can be proved in a surprisingly weak theory of sets with formulae of very low complexity. As a consequence, it is shown that the satisfiability problem for propositional combinations of restricted prenex formulae with two quantifiers alternations (the so-called $(\forall\exists\forall)_0$-formulae) is unsolvable.

In this paper, we show that just one quantifier alternation suffices to obtain undecidability. This class of formulae will be denoted by $(\forall\exists)_0$.

In particular, we show how to effectively associate a suitable $(\forall\exists)_0$-formula φ to any given polynomial Diophantine equation D in such a way that φ is satisfiable by a set model if and only if D has integer solutions. Then, the undecidability of the $(\forall\exists)_0$-formulae will follow directly from the unsolvability of Hilbert's tenth problem, i.e. the problem to decide for a polynomial Diophantine equation whether or not it has integer solutions.

As by-products, the undecidability of other classes of quantified and unquantified formulae will be also derived.

If, as we conjecture, the class of $(\forall)_0$-formulae (i.e. propositional combinations of restricted purely universal formulae) is decidable, then our result allows to locate precisely the boundary between the decidable and the undecidable in set theory.

450 ON THE DECIDABILITY OF PRESBURGER ARITHMETIC WITH HOMOGENEOUS EXPONENTIATION

D. Cantone, May 1989

One of the basic problems of mathematical logic and theoretical computer science is to find as a tight boundary as possible between the decidable and the undecidable. Many advances in this direction have recently been made for some fragments of set theory (see for example Cantone-Ferro-Omodeo, *Computable set theory*, Oxf. Univ. Press).

In this paper we show that Presburger arithmetic (which is well known to be decidable) can be extended with "homogeneous" exponentiation (in short: h.e.) without disrupting decidability.

It is known that the most interesting and expressive theories are intractable, when decidable. This is, for instance, the case for every (logical) theory which can express the propositional calculus. In fact, from the result of Fisher and Rabin, it follows that Presburger arithmetic has a nondeterministic time lower bound of $2^{2^{cn}}$. Obviously, the same result also holds when h.e. is admitted.

Nevertheless, even intractable decision procedures are amenable of practical utilizations for small fragments of their domain of application. For instance, despite of the super-exponential lower bound cited above, unquantified Presburger arithmetic is $\mathcal{NP}$-complete. This fact has been generalized by various authors who gave nondeterministic polynomial time decision procedures for unquantified Presburger formulae extended respectively with uninterpreted predicate and function symbols and with array constructs. Likewise, it turns out that unquantified Presburger arithmetic with h.e. is $\mathcal{NP}$-complete.

451 FRAGMENTATION OF VARIABLE LENGTH PACKETS

D. Clark, B. Mishra, May 1989; also RR 200

This note analyzes the loss of bandwidth due to fragmentation caused by the variable length packets in the asynchronous frame of a communication cycle in the protocol used by the GANGLIA network [Clark 89]. We provide a simple back-of-the-envelop computation, which is later justified by a more rigorous analysis.

452 MOSAIC SYSTEM DESCRIPTION, SPECIFICATION AND PLANNING

I. Greenfeld, F. Hansen, J. Fehlinger, E. Pavlakos, June 1989; also RR 201

The MOSAIC system is specified and described. MOSAIC is an open-system controller intended for use on the next generation of machine tools. Its configuration is based on the open-system computer technology, thus providing a suitable solution for the multi-device multi-sensor environment of future machine-tools. The implementation of the system at the Robotics and Manufacturing Laboratory of New York University is described. The current configuration is based on a Sun workstation, an extended VMEbus, and a multi-processor environment. The system is demonstrated on a 3-axis milling machine, equipped with subsidiary devices like a manipulator and a vision system. The document contains information on the current configuration, future planned enhancements, and various preliminary ideas. Details of the hardware and software are included. Our short-term and long-term plans for research and implemetation are also presented.

453 RANDOM B-TREES WITH INSERTS AND DELETES

T. Johnson, D. Shasha, June 1989

The utilization of B-tree nodes determines the number of levels in the B-tree and hence its performance. Yao showed that the utilization of B-tree nodes under pure inserts was 69%. We derive analytically and verify by simulation the utilization of B-tree nodes constructed from a mixture of insert and delete operations. Assuming that nodes only merge when they are empty we show that the utilization is 39% when the number of inserts is the same as the number of deletes. However, if there are just 5% more inserts than deletes, then the utilization is over 62%. We also calculate the probability of splitting

and merging. We derive a simple rule-of-thumb that accurately calculates the probability of splitting. We also model B-trees that merge half-empty nodes. The utilization of merge-at-half B-trees is slightly larger than the utilization of merge-at-empty B-trees (within 10% if if there are at least 5% more inserts than deletes), but the restructuring rate is much higher. We present two models for computing B-tree utilization, the more accurate of which remembers items inserted and then deleted in a node. After analyzing the leaf level of a random B-tree, we analyze the upper levels, and show that they have the structure of the leaf level of a pure-insert tree, independent of the percentage of operations that are deletes (if the percentage of deletes is less than 50%).

454 A PROPERTY OF GRAPHS OF POLYTOPES

N. Prabhu, June 1989

We prove that the subgraph obtained by removing the vertices of a k-face from the graph of a d-polytope ($0 \leq k < d$) is $(d-k-1)$-connected. Further we show that this lower bound is tight for $k < d-1$. We also show that for $k = d-1$ the known lower bound is tight.

455 INDUCTIVE PROOF OF MACAULAY'S THEOREM

T. Dube, June 1989; also RR 202

Macaulay's theorem states that for every ideal I in a multi-variate polynomial ring $\mathcal{R} = K[x_1, \ldots, x_n]$ there is a lexicographic monomial ideal L_I such that L_I has the same Hilbert function as I. This theorem allows one to characterize exactly the form of Hilbert polynomials.

This report uses a recursive characterization of a Borel-fixed monomial ideal to provide a new inductive proof of Macaulay's famous theorem.

456 ON THE EXISTENCE OF GENERIC COORDINATES

T. Dube, June 1989; also RR 203

A well-known technique in computational mathematics is to change to a coordinate system which exhibits special attractive properties. A somewhat unusual case is the change to generic coordinates in proving a degree bound for Groebner bases. Here, the generic coordinates exhibit the worst case behavior, and thus allow the construction of upper bounds.

A sketch of the existence of such coordinates was provided in Bayer's thesis. This report expands upon that sketch to provide a fully detailed proof that such a coordinate system always exists.