Recent Technical Report List

May 1991–

New York University
Courant Institute of Mathematical Sciences
Computer Science Department

558-R246 ON-LINE SCHEDULERS FOR OVERLOADED REAL-TIME SYSTEMS

G. Koren, B. Mishra, A. Raghunathan, D. Shasha, May 1991

There are several scheduling algorithms proposed for real-time systems with sporadic tasks that exhibit optimal behavior if the system is underloaded. Examples of such scheduling algorithms include earliest-deadline-first (D), smallest-slack-time-first (SL), etc. However, these algorithms performs poorly when the system is overloaded. This is a serious problem, since even the best-modeled systems may not be able to avoid intermittent overloading situations.

In this paper, we propose an on-line scheduling algorithm for a uniprocessor in an environment of sporadic tasks that addresses this issue. We assume that the value of each task is proportional to its computation time provided it completes by its deadline; if it misses its deadline, its value is zero. An optimal scheduling algorithm, then, maximizes the sum of the values of the completed tasks. Even for a clairvoyant model, this problem is NP-hard. In this paper, we formalize the notions of when a system is overloaded and when it is underloaded. We devise a scheduling algorithm called D* , that provides the following performance guarantee. Informally speaking, D* achieves optimal performance while the system is underloaded and achieves a value that is at least 1/5 times the length of the overloaded period while the system is overloaded. To the best of our knowledge, this is the only algorithm that makes such a non-trivial guarantee for a real-time system that may face intermittent overloading. We implement D*, using simple data structures, at a cost of O(log n) time per task, where n bounds the number of tasks in the system at any instant.

We also show that no on-line algorithm can guarantee a value that is better than 1/3.224745 times the value obtained by a clairvoyant algorithm.

Our experimental data indicate that the proposed algorithm (referred to as D* algorithm) is comparable to Locke's best effort scheduler (BE) in the uniform arrival case, while guaranteeing the one-fifth bound even for the pathological cases that would cripple Locke's scheduler.

559-R247 OPTIMAL MOTION PLANNING FOR A FALLING CAT

C. Fernandes, L. Gurvits, Z. X. Li, May 1991

How does a falling cat change her orientation in midair without violating the angular momentum constraint? This has become an interesting question to both control engineers and robotists. In this paper, we address this problem together with a constructive solution. First, we show that *a falling cat problem* is equivalent to the *constructive nonlinear controllability problem*. Thus, the same principle and *algorithm* used by a falling cat can be used for space robotic applications such as reorientation of a satellite using rotors, attitude control of a space structure using internal motion and for other robotic tasks such as dextrous manipulation with multifingered robotic hands and nonholonomic motion planning for mobile robots. Then, using ideas from Ritz Approximation Theory we develop a simple algorithm for motion planning by a falling cat. Finally, we test the algorithm through simulation on two widely accepted models of a falling cat. It is interesting to note that one set of simulated trajectories is close to trajectories used by a real cat.

560 A THEORY OF NATURAL LEARNING

A. Botta, May 1991

Unsupervised learning is based on capturing regularities in data. We formalize the vague notion of regularity, using the concept of algorithmic information (Solomonoff, Chaitin, Koppel). We present a theory on how regularities are induced and accumulated. A generative model captures a regularity if it achieves compression. A basic regularity is a building block for hierarchical structures. We prove that a basic regularity may be identified as a local maximum in compressibility. Stepwise induction is a polynomial-time approach to structures whose basic components have bound complexity.

Agents exploring a universe engage in active learning. The regularities of their sensory-motor streams are similar to Piaget's schemes and constituents of an induced ontology. We illustrate these ideas on three microworlds. First are Moore automata. State representations are constructed incrementally from results of tests when in that state and from outputs perceived on the way to that state. The second world contains loosely coupled geometric objects. They are basic regularities identifiable by stepwise induction. In the third world the agent has an elaborate eye and can move objects on a tiled surface. Statistical

correlations between sets of stimuli are induced, then models are constructed to generate instances of new correlations from already known ones.

Algorithmic information theory allows a unified perspective on many areas of learning research. We define analysis as the separation of novelty in data from the already known. We present explanation based generalization as a well formalized instance of analysis, and constructive induction as an ill defined instance. We show EBG to specialize a theory through positive examples, and we prove it a language independent method, valid beyond the predicate calculus representations.

561 AMORTIZED COMPLEXITY OF DATA STRUCTURES
R. Sundar, May 1991

This thesis investigates the amortized complexity of some fundamental data structure problems and introduces interesting ideas for proving lower bounds on amortized complexity and for performing amortized analysis. The problems are as follows:

- *Dictionary Problem:* A *dictionary* is a dynamic set that supports searches of elements and changes under insertions and deletions of elements. It is open whether there exists a dictionary data structure that takes constant amortized time per operation and uses space polynomial in the dictionary size. We prove that dictionary operations require log-logarithmic amortized time under a multilevel hashing model that is based on Yao's cell probe model.

- *Splay Algorithm's Analysis: Splay* is a simple, efficient algorithm for searching binary search trees, devised by Sleator and Tarjan, that uses rotations to reorganize the tree. Tarjan conjectured that Splay takes linear time to process deque operation sequences on a binary tree and proved a special case of this conjecture called the *Scanning Theorem.* We prove tight bounds on the maximum numbers of various types of right rotations in a sequence of right rotations performed on a binary tree. One of the lower bounds refutes a conjecture of Sleator. We apply the upper bounds to obtain a nearly linear upper bound for Tarjan's conjecture. We give two new proofs of the Scanning Theorem, one of which is a potential-based proof that solves a problem of Tarjan.

- *Set Equality Problem:* The task of maintaining a dynamic collection of sets under various operations arises in many applications. We devise a fast data structure for maintaining sets under equality-tests and under creations of new sets through insertions and deletions of elements. Equality-tests require constant time and set-creations require logarithmic amortized time. This improves previous solutions.

562 DECOMPOSITION AND FICTITIOUS DOMAINS METHODS FOR ELLIPTIC BOUNDARY VALUE PROBLEMS
S. V. Nepomnyaschikh, May 1991

Boundary value problems for elliptic second order equations in three-dimensional domains with piecewise smooth boundaries are considered. Discretization of the problem is performed using a conventional version of the finite element method with piecewise linear basis functions. The main purpose of the paper is the construction of a preconditioning operator for the resulting system of grid equations. The method is based on two approaches: decomposition of the domain into subdomains and using a new version of the method of fictitious domains. The rate of convergence of the corresponding conjugate gradient method is independent of both the grid size and the number of subdomains.

563-R248 A QUOTIENT SPACE HOUGH TRANSFORM FOR SPACE-VARIANT VISUAL ATTENTION
A. S. Rojer, E. L. Schwartz, June 1991

We consider the problem of "visual attention" in the context of space-variant machine vision: Is there a general theoretical and practical formulation for the "next-look" problem to guide a space-variant sensor to a rapid choice for its next fixation point? This topic is developed in the context of Hough transform methods, by the addition of a third space to the usual feature and object spaces considered in traditional Hough methods. This third space is a "behavioural," or "motor" space, which is typically low-dimensional with respect to the feature and object spaces. For example, the motor space of particular interest for us is the two-dimensional manifold of monocular eye positions. By "collapsing" the generalized Hough transform into a low dimensional motor space, we show that it is possible to avoid a practical difficulty in applying Hough transform methods, which is the exponential dependence of the accumulator array on the dimensionality of the object space. Beginning with a simple and very general Bayesian scheme, we derive in stages the generalized Hough transform as a special case. Since "attentional" applications, by their nature, require only partial knowledge about objects, computation of all the parameters characterizing a scene object is superfluous and wasteful of computational resources. This suggests that for an attentional application, collapsing the (large) object space onto the (small) motor space provides a computationally grounded definition

of the term "visual attention". We illustrate these ideas with an example of choosing "fixation" points for a space-variant sensor in a machine vision application for real-time reading of license plates of moving vehicles.

564-R249 THE HOUGH TRANSFORM AS BAYESIAN INFERENCE
A. S. Rojer, E. L. Schwartz, June 1991

Although the Hough transform and its generalizations have been used for many years in computer vision and pattern recognition, its conceptual grounding has been largely heuristic. We demonstrate here that the generalized Hough transform may be viewed as a special case of Bayesian inference. Starting from a simple and very general Bayesian scheme, we derive in stages the generalized Hough transform(of Ballard). By doing so, we explicate a series of assumptions concerning featural independence and equal a priori occurrence of objects which are implicit in the use of the generalized Hough transform. When (if ever) these assumptions are satisfied, Hough transform methods are optimal (in a Bayesian sense) for object recognition. Consideration of the assumptions also illuminates conditions under which performance degradation of the Hough approach is expected. The goal of this paper is not to yet more variations on the Hough transform, but rather to clarify the implicit assumptions associated with this methodology, and to suggest future research in which these assumptions can verified or relaxed.

565 VIVID REPRESENTATIONS
E. Davis, June 1991

This paper criticizes the widespread idea that knowledge bases in AI systems should be complete and that representations should be "model-like." The arguments in favor of such representations are less cogent and more ambiguous than they appear at first. Levesque's suggestion that representations should be "vivid" is extremely restrictive, particularly in its uniform imposition of a closed-world assumption. Spatial representations that are adequate for reasoning about a wide range of physical phenomena must ultimately either use complex symbolic reasoning or deal with partial and imperfect approximations. Requiring that temporal representations be fully detailed simulations will often be extremely inefficient. Finally, a flexible intelligent system must necessarily deal with partial information of all kinds, and the techniques for carrying out reasoning about partial information using complete representations are very limited in their application.

566 OPTIMALITY CONDITIONS AND DUALITY THEORY FOR MINIMIZING SUMS OF THE LARGEST EIGENVALUES OF SYMMETRIC MATRICES
M. L. Overton, R. S. Womersley, June 1991

This paper gives max characterizations for the sum of the largest eigenvalues of a symmetric matrix. The elements which achieve the maximum provide a concise characterization of the generalized gradient of the eigenvalue sum in terms of a dual matrix. The dual matrix provides the information required to either verify first-order optimality conditions at a point or to generate a descent direction for the eigenvalue sum from that point, splitting a multiple eigenvalue if necessary. A model minimization algorithm is outlined, and connections with the classical literature on sums of eigenvalues are explained. Sums of the largest eigenvalues in absolute value are also addressed.

Keywords: symmetric matrix, maximum eigenvalues, spectral radius, minimax problem, max characterization, generalized gradient.

567 ON THE SUBDIFFERENTIABILITY OF FUNCTIONS OF A MATRIX SPECTRUM I: MATHEMATICAL FOUNDATIONS
J. V. Burke, M. L. Overton, June 1991

We consider analytic matrix valued mappings $A: \mathbb{C} \mapsto \mathbb{C}^{n \times n}$ and study the variational properties of the spectrum of $A(\epsilon)$. Of particular interest are $\alpha(\epsilon)$ and $\rho(\epsilon)$, respectively the spectral abscissa and the spectral radius of $A(\epsilon)$. In this paper, we introduce the mathematical techniques required for this analysis. We begin with polynomials and discuss the bifurcation of the roots of a polynomial having analytic coefficients. It is this bifurcation phenomenon that leads to the nonlipschitzian behavior of the types of functions that we wish to study. *Puiseux-Newton series and diagrams* are then introduced as a means for analyzing these bifurcations. It is shown how these techniques can be used to describe the tangent cone to certain sets of *stable* polynomials.

Matrices and polynomials are connected via characteristic polynomials. Further properties of the spectrum of a matrix $A^{(0)}$ are obtained from a block diagonalization of $A^{(0)}$, where the kth diagonal block is an n_k by n_k upper triangular matrix, with a constant diagonal whose value is an eigenvalue of $A^{(0)}$ with multiplicity n_k. By using results of Arnold on the *versal deformation* of a matrix, we show how the results concerning polynomials can be translated into into results about matrices. In the case that the block diagonal form is the Jordan form, the conditions which are obtained reduce to conditions on

generalized traces of a matrix.

568 ON THE SUBDIFFERENTIABILITY OF FUNCTIONS OF A MATRIX SPECTRUM II: SUBDIFFERENTIAL FORMULAS
J. V. Burke, M. L. Overton, June 1991

Variational properties for the spectral radius and spectral abscissa of an analytic matrix valued mapping $A: \mathbb{C}^s \mapsto \mathbb{C}^{n \times n}$ are considered. A notion of directional differentiability is introduced that allows us to exploit the perturbation results of Part I of this study. Lower bounds for the directional derivative are established which yield formulas for the directional derivative when a natural nondegeneracy condition is satisfied. These formulas are interpreted in the extreme cases where the eigenvalues attaining either the spectral radius or the spectral abscissa are nonderogatory or semisimple (nondefective). We conclude by investigating the relationship with the proximal normal subdifferential. In particular, it is shown that the proximal normal subdifferential is always nonempty.

569-R250 APPLICATIONS AND ANALYSIS OF PROBABILISTIC TECHNIQUES
P. V. Tetali, June 1991

The thesis illustrates the strength of randomness by applying some recent probabilistic techniques to solve problems in number theory, graph theory and computer science.

The first part of the thesis is concerned with random construction of integer sequences with certain additive properties. A set of natural numbers is called an "asymptotic basis" of order k, if every number (sufficiently large) can be expressed as a sum of k distinct numbers from the set. We prove that for every fixed k, there exists an asymptotic basis of order k such that the number of representations of n is $Theta(logn)$. The case $k = 2$ was proved in 1956 by Paul Erdos.

The second part deals with analysis of random walks on graphs. Random walks on graphs have been known to have interesting analogies in electrical networks. A precise characterization of "effective resistance" in electrical networks is provided in this thesis in terms of random walks on the underlying graphs. The interpretation of effective resistance yields interesting new results and new proofs for some known results. The main result here is an exact formula for the "hitting time" between two vertices in terms of the effective resistances in the network, settling an open question. This is much in the spirit of the commute time result by Ashok Chandra et al.

The final part of the thesis addresses problems involving "Discrepancy" Discrepancy is a combinatorial problem that arises in various contexts such as communication complexity in multi-party protocols, quasi-randomness in graphs and hypergraphs, hypergraph coloring, and the Gale-Berlekamp switching game. In k-dimensions, the game consists of an nXn array A (of size n^k) of lights. There are kn switches, each controlling a hyperplane of n^{k-1} lights. Each switch when thrown changes each light in its control from "on" to "off" or from "off" to "on" Discrepancy is defined as the maximum difference between the number of lights on and off that one can obtain, given any initial configuration for the lights. It follows from a result of Erdos and Spencer that discrepancy = $Theta(n^{(k+1)/2})$. This result, however, is existential due to its critical use of randomness. An optimal deterministic sequential and an NC algorithm are given in this thesis. The basic idea is to use Spencer's method of conditional expectation to remove randomness. But the crucial part of the derandomization is in justifying the efficient computation of the conditional expectations involved. For the NC algorithm, we prove that 4-wise independent random setting for the switches works!

570 ADDITIVE SCHWARZ METHODS FOR ELLIPTIC FINITE ELEMENT PROBLEMS IN THREE DIMENSIONS
M. Dryja, O. B. Widlund, June 1991

Many domain decomposition algorithms and certain multigrid methods can be described and analyzed as additive Schwarz methods. When designing and analyzing domain decomposition methods, we encounter special difficulties in the case of three dimensions and if the coefficients are discontinuous and vary over a large range. In this paper, we first introduce a general framework for Schwarz methods. Three classes of applications are then considered: certain wire basket based iterative substructuring methods, Neumann-Neumann algorithms with low dimensional, global subspaces and a modified form of a multilevel algorithm introduced by Bramble, Pasciak and Xu.

571 EFFICIENT ALGORITHMS FOR CYCLIC SCHEDULING
F. Gasperoni, U. Schwiegelshohn, July 1991

This work addresses the problem of non-preemptively scheduling a cyclic set of interdependent operations, representing for instance a program loop, when p identical processors are available. For $p = \infty$ we give a simple, efficient, polynomial time algorithm producing optimum results. When $p < \infty$ the problem becomes NP-hard and a slight modification of our algorithm

generates provably close to optimum results.

572 AN OPTIMAL SCHEDULING ALGORITHM WITH A COMPETITIVE FACTOR FOR REAL-TIME SYSTEMS
G. Koren, D. Shasha, July 1991

We consider real-time systems in which the value of a task is proportional to its computation time. The system obtains the value of a given task if the task completes by its deadline. Otherwise, the system obtains no value for the task.

When such a system is underloaded (i.e. there exists a schedule for which all tasks meet their deadlines), Dertouzos showed that the earliest deadline first algorithm will achieve 100% of the possible value. We consider the case of a possibly overloaded system and present an algorithm which:
1. behaves like the earliest deadline first algorithm when the system is underloaded.
2. obtains at least 1/4 of the maximum value that an optimal clairvoyant algorithm could obtain even when the system is overloaded.

We implement our algorithm with an amortized cost of $O(\log n)$ time per task, where n bounds the number of tasks in the system at any instant.

573 ONLINE ALGORITHMS FOR FINGER SEARCHING
R. Cole, A. Raghunathan, August 1991

A well-known technique for speeding up access into search structures is to maintain *fingers* that point to various locations of the search structure. The advantage of maintaining fingers is that accessing information near a finger is normally faster than the usual method of access. In this paper, we consider the problem of choosing locations in a large search structure at which to maintain fingers. In particular, we consider a server problem in which k servers move along a line segment of length m, where m is the number of keys in the search structure. Since fingers may be arbitrarily copied, we allow a server to jump, or *fork*, to a location currently occupied by another server. We present online algorithms and analyze their competitiveness. We show that the case $k = 2$ behaves differently from the case $k \geq 3$, by showing that there is a 4-competitive algorithm for $k = 2$ that never forks its fingers. For $k \geq 3$, we show that any online algorithm that does not fork its fingers can be at most $\Omega(m^{1/2})$-competitive. By contrast, in the main result of the paper, we show that for $k = 3$, there is an online algorithm that forks and is constant competitive (independent of m, the size of the search structure). Our algorithm is simple and implementable, although our analysis is complicated. The existence of a constant competitive algorithm for $k > 3$ is left as an open question.

574 THE EXPECTED ADVANTAGE OF ASYNCHRONY
R. Cole, O. Zajicek, September 1991

This paper studies the implicit costs of synchronization and the advantage that may be gained by avoiding synchronization in asynchronous environments. An asynchronous generalization of the PRAM model called the APRAM model is used and appropriate complexity measures are defined. The advantage asynchrony provides is illustrated by 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.

575-R251 SEGMENTATION AND SURFACE-BASED MODELING OF OBJECTS IN THREE-DIMENSIONAL BIOMEDICAL IMAGES
A. Kalvin, September 1991

The rapid development of technologies for imaging the human body has led to a growing interest in the extraction and analysis of objects in 3D biomedical images for applications in fields such as clinical medicine, biomedical research, and physical anthropology.

This dissertation examines the problem of creating surface-based geometric models of biomedical objects that are suitable for analysis through visualization, mensuration, and manipulation. This is a two-stage problem. First the objects are identified by segmenting the 3D image into regions of interest, and then surface-based models of the objects are created.

We discuss the issues of segmentation and surface construction and introduce the following new methods for solving these problems.

First, we present the MLO algorithm, a general-purpose, domain-independent segmentation algorithm that has been applied successfully to identify skulls in CT images, the ventricle walls of the heart in MR images, brain ventricles in CT images, and carotid arteries in MR angiography images. It uses an iterative, cooperative procedure to segment an image by optimizing a

cost function. To achieve a fast segmentation, a coarse-to-fine strategy is employed, using a multiresolution pyramid.

The GRG algorithm is a model-driven, special-purpose algorithm for identifying thin bone in CT head images. The algorithm, developed specifically for craniofacial surgical planning, uses anatomical knowledge in the segmentation process, and can handle the abnormal anatomy of craniofacial patients. It successfully finds most of the thin bone that can not be found using previous methods.

ALLIGATOR is a surface construction algorithm that creates models using the "winged-edge" data structure of Baumgart, enabling efficient access to the topological and geometric information of the surfaces, and permitting efficient, topologically consistent modifications to the representations.

Unlike previous surface construction algorithms, ALLIGATOR is suitable not just for visualizing biomedical objects, but for measuring and manipulating them as well. Another important feature of ALLIGATOR is that it uses an adaptive face-merging process to create surface models that are significantly more concise, in terms of vertices, edges, and faces, than the models produced by other surface construction algorithms.

576-R252 ATTITUDE CONTROL OF SPACE PLATFORM/MANIPULATOR SYSTEM USING INTERNAL MOTION

C. Fernandes, L. Gurvits, Z. Li, September 1991

Attitude control of a space platform/manipulator system, using internal motion, is an example of a Nonholonomic Motion Planning (NMP) problem arising from symmetry and conservation laws. Common to NMP problems are that *an admissible configuration space path is constrained to a given nonholonomic distribution.* We formulate the dynamic equations of a system consisting of a 3-D.O.F Puma like manipulator attached to a space platform (e.g. a space station or a satellite), as a NMP problem and discuss controllability of the system. Then we describe the application of a simple algorithm for obtaining approximate optimal solutions. We conclude with a description of simulation software implementing the algorithm and simulation results for two experiments.

577-R253 FOUNDATIONS OF NONHOLONOMIC MOTION PLANNING

C. Fernandes, L. Gurvits, Z. Li, September 1991

Nonholonomic Motion Planning (NMP) problems arise not only from classic nonholonomic constraints (e.g., rolling contact), but also from symmetries and conservation laws of holonomic mechanic (or biological) systems. In real life, NMP systems exist in abundance: a (kinematic) unicycle, a front-wheel-drive cart, a system of two contacting bodies with rolling constraint, a system consisted of a dextrous robotic hand and an object to be manipulated, a spring-board diver, an astronaut, a falling-cat, a free-flying robot with a floating base, a satellite with a number of rotors, a space station and the Hubble Space Telescope. Common to these systems are that *an admissible configuration space path is constrained to a given nonholonomic distribution.* Thus, NMP deals with the problem of optimal path planning subject to a nonholonomic distribution and possibly additional holonomic constraints, e.g., collision avoidance.

In this preliminary report, we describe our recent work in NMP.

First, we study examples of NMP problems, examine the constraints or symmetries associated with these systems and show the derivation of nonholonomic distributions. We transform a NMP problem into a nonlinear control problem so that analytic techniques from nonlinear control theory can be used.

Second, we study the existence of motion using Frobenius theory and Chow's theorem. We also develop several new results for checking controllability for a class of NMP systems.

Third, we apply variational principles to study optimal trajectories of NMP. We show that an optimal solution necessarily satisfies Hamilton's equation for some well-defined Hamiltonian function. Furthermore, we prove the equivalence of two different formulations of NMP.

Fourth, we present a numerical algorithm, called the *Basis Algorithm*, for computing approximate optimal solutions of NMP. We show that the objective function is locally convex and has desirable convergence property.

Finally, we describe a simulation environment developed at our laboratory for the study of NMP, along with simulation results obtained using the Basis Algorithm.

578-R254 STABILITY PROOF OF RAIBERT'S FOUR-LEGGED HOPPER IN BOUNDING GAIT

A. Neishtadt, Z. Li, September 1991

M. Raibert and his co-workers have designed and built a four-legged hopper, or a quadruped. Using the notion of virtual legs and a 3-part control algorithm, they have achieved stable hopping motion of the system in bounding gait as well as in trotting and pacing gaits. In this paper, we provide a proof of stability of the system. The mathematical tools we use in the

proof derives from dynamical system theory, in particular perturbation techniques.
The main results of the paper are as follows:

1. *When the touch-down angle is small, there is exactly one periodic solution for each fixed energy level of the Hamiltonian system.*

2. *$j < 1$ is a necessary (and also sufficient when the touch-down angle is small) condition for Hamiltonian stability of the system, where $j = I/ml^2$ is a dimensionless parameter of the system.*

3. *For fixed levels of leg dissipation and thrusting, there exists a unique periodic solution along which energy change is balanced. The system is asymptotically stable with respect to change in energy.*

4. *With a damping term from the attitude control law added, the system becomes asymptotically stable.*

579 DIFFERENTIAL PROPERTIES OF EIGENVALUES
J. Burke, M. Overton, September 1991

We define and study a directional derivative for two functions of the spectrum of an analytic matrix valued function. These are the maximum real part and the maximum modulus of the spectrum. Results are first obtained for the roots of polynomials with analytic coefficients by way of Puiseux-Newton series. In this regard, the primary analytic tool is the so called Puiseux-Newton diagram. These results are then translated into the context of matrices. Precise results are obtained when the eigenvalues that achieve the maximum value for the function under consideration are all either nondefective or nonderogatory. In the defective derogatory cases a general lower bound for the directional derivative is given which, in particular, describes those directions in which the directional derivative attains an infinite value.

580 AN ADDITIVE SCHWARZ METHOD FOR THE P-VERSION FINITE ELEMENT METHOD
L. Pavarino, September 1991

The additive Schwarz method was originally proposed for the h-version finite element method for elliptic problems. In this paper, we apply it to the p-version, in which increased accuracy is achieved by increasing the degree of the elements while the mesh is fixed. We obtain a constant bound, independent of p, for the condition number of the iteration operator in two and three dimensions. The result holds for linear, self adjoint, second order elliptic problems and for quadrilateral elements.

581 SOME SCHWARZ METHODS FOR SYMMETRIC AND NONSYMMETRIC ELLIPTIC PROBLEMS
O. Widlund, September 1991

This paper begins with an introduction to additive and multiplicative Schwarz methods. A two-level method is then reviewed and a new result on its rate of convergence is established for the case when the overlap is small. Recent results by Xuejun Zhang, on multi-level Schwarz methods, are formulated and discussed. The paper is concluded with a discussion of recent joint results with Xiao-Chuan Cai on nonsymmetric and indefinite problems.

582 MULTILEVEL ADDITIVE SCHWARZ METHODS
X. Zhang, September 1991

We consider the solution of the algebraic system of equations which result from the discretization of elliptic equations. A class of multilevel algorithms are studied using the additive Schwarz framework. We establish that, in the general case, the condition number of the iterative operator grows at most linearly with the number of levels. The bound is independent of the mesh sizes and the number of levels under the H^2-regularity assumption of the equation. This is an improvement on Dryja and Widlund's result on a multilevel additive Schwarz algorithm, as well as Bramble, Pasciak and Xu's result on the BPX algorithm.

583 DOMAIN DECOMPOSITION ALGORITHMS FOR THE BIHARMONIC DIRICHLET PROBLEM
X. Zhang, September 1991

We consider additive Schwarz methods for the biharmonic Dirichlet problem and show that the algorithms have optimal

convergence properties for some conforming finite elements. Some multilevel methods are also discussed.

584 STUDIES IN DOMAIN DECOMPOSITION: MULTILEVEL METHODS AND THE BIHARMONIC DIRICHLET PROBLEM

X. Zhang, September 1991

A class of multilevel methods for second order problems is considered in the additive Schwarz framework. It is established that, in the general case, the condition number of the iterative operator grows at most linearly with the number of levels. The bound is independent of the mesh sizes and the number of levels under a regularity assumption. This is an improvement of a result by Dryja and Widlund on a multilevel additive Schwarz algorithm, and the theory given by Bramble, Pasciak and Xu for the BPX algorithm.

Additive Schwarz and iterative substructuring algorithms for the biharmonic equation are also considered. These are domain decomposition methods which have previously been developed extensively for second order elliptic problems by Bramble, Pasciak and Schatz, Dryja and Widlund and others.

Optimal convergence properties are established for additive Schwarz algorithms for the biharmonic equation discretized by certain conforming finite elements. The number of iterations for the iterative substructuring methods grows only as the logarithm of the number of degrees of freedom associated with a typical subregion. It is also demonstrateed that it is possible to simplify the basic algorithms. This leads to a decrease of the cost but not of the rate of convergence of the iterative methods. In the analysis, new tools are developed to deal with Hermitian elements. Certain new inequalities for discrete norms for finite element spaces are also used.

585 EFFICIENT LOOP-LEVEL PARALLELISM IN ADA

M. Hind, September 1991

Parallelism in scientific applications can most often be found at the loop level. Although Ada supports parallelism via the *task* construct, its coarseness renders it unsuitable for this *light-weight* parallelism. In this work, we propose Ada constructs to achieve efficient loop-level parallelism in ANSI-Ada. This is accomplished in two steps. First, we present an idiom that allows the specification of light-weight tasks. Second, we give an efficient implementation of this idiom that is considerably more efficient than a standard Ada task.

In addition, we present an idiom that makes the fetch_and_add synchronization primitive available at the Ada level. Our implementation of this idiom is more efficient in both time and space than previous results. In addition to providing universal synchronization, using fetch_and_add simplifies program analysis (e.g. proving the absence of race conditions in the implementation of a parallel algorithm). Since all these idioms are written in standard Ada, they maintain the portability that is central to the mandated uses of the language.

586 ON SHAPE OPTIMIZING THE RATIO OF THE FIRST TWO EIGENVALUES OF THE LAPLACIAN

J. Haeberly, October 1991

We investigate numerically a 1956 conjecture of Payne, Polya, and Weinberger. The conjecture asserts that the ratio of the first two eigenvalues of the Laplacian on a bounded domain Ω of the plane with Dirichlet boundary conditions reaches its minimum value precisely when Ω is a disk. A crucial feature of this problem is the loss of smoothness of the objective function at the solution. The following results form the core of our numerical treatment. First, we construct finite dimensional families of deformations of a disk equipped with a uniform triangulation. This permits the formulation of a discrete model of the problem via finite element techniques. Second, we build on the work of M. Overton to derive optimality conditions in terms of Clarke's generalized gradients for nonsmooth functions. These ideas are then combined into an algorithm and implemented in Fortran.

587 NEW THEORETICAL AND COMPUTATIONAL RESULTS FOR REGULAR LANGUAGES

C. Chang, R. Paige, October 1991

We show how to turn a regular expression into an $O(s)$ space representation of McNaughton and Yamada's NFA, where s is the number of NFA states. The standard adjacency list representation of McNaughton and Yamada's NFA takes up $s + s^2$ space in the worst case. The adjacency list representation of the NFA produced by Thompson takes up between $2r$ and $5r$ space, where $r \geq s$ in general, and can be arbitrarily larger than s. Given any set T of NFA states, our representation can be used to compute the set N of states one transition away from the states in T in optimal time $O(|T| + |N|)$. McNaughton and Yamada's NFA requires $\Theta(|T| \times |N|)$ in the worst case. Using Thompson's NFA, the equivalent calculation requires $\Theta(r)$ time in the worst case.

An implementation of our NFA representation confirms that it takes up an order of magnitude less space than McNaughton and Yamada's machine. An implementation to produce a DFA from our NFA representation by subset construction shows linear and quadratic speedups over subset construction starting from both Thompson's and McNaughton and Yamada's NFA's. It also shows that the DFA produced from our NFA is as much as one order of magnitude smaller than DFA's constructed from the two other NFA's.

588-R255 A MINIATURIZED ACTIVE VISION SYSTEM
B. B. Bederson, R. S. Wallace, E. L. Schwartz, November 1991

We have developed a prototype miniaturized active vision system whose sensor architecture is based on a logarithmic structured space-variant pixel geometry. This system integrates a CCD sensor, mini pan-tilt actuator, controller, general purpose processor and display. Due to the ability of space-variant sensors to cover large work-spaces yet provide high acuity with an extremely small number of pixels, space-variant active vision system architectures provide the potential for radical reductions in system size and cost. In this paper, we describe a prototype space-variant active vision system which performs such tasks as connected components analysis and motion tracking. The potential application domains for systems of this type include vision systems for mobile robots and robot manipulators, traffic monitoring systems, security and surveillance, and consumer video communications.

589-R256 SPACE-VARIANT IMAGE PROCESSING
R. S. Wallace, P. Ong, B. B. Bederson, E. L. Schwartz, November 1991

Space-variant sensors have nonuniform sampling across the image plane. Partially motivated by the observation that human vision is strongly space-variant, yielding an image compression for the human system that is estimated to be as much as four orders of magnitude, a number of research groups have been experimenting with space-variant sensors. Such systems cover wide solid angles yet maintain high acuity in their central regions. Implementation of space-variant systems pose at least two outstanding problems. First, such a system must be active, in order to utilize its high acuity region; second, there are fascinating image processing problems introduced by the non-uniform pixel size, shape and connectivity. Familiar image processing operations such as connected components, convolution, template matching, and even image translation, take on new and different forms when defined on space-variant arrays. The present paper provides a general method for space-variant image processing, based on a connectivity graph which represents the neighbor-relations in an arbitrarily structured sensor. We illustrate this approach with the following applications:

- Connected components is reduced to its graph theoretic counterpart. We illustrate this on a logmap sensor, which possesses a difficult topology due to the branch cut associated with the complex logarithm function.
- We show how to write local image operators in the connectivity graph that are independent of the sensor geometry, and solve a sensor calibration problem with these operations.
- We relate the connectivity graph to pyramids over irregular tessalations, and implement a local binarization operator in a 2-level pyramid.
- Finally, we expand the connectivity graph into a structure we call a translation graph, representing the effects of translation in space-variant image sensors. Using the translation graph, we define an efficient algorithm for translation in the logmap image and solve the template matching problem for space-variant images.

Because of the very small number of pixels typical of logarithmic structured space-variant arrays, the connectivity graph approach to image processing is suitable for real-time implementation, and provides a generic solution to a wide range of image processing applications with space-variant sensors.

591-R257 A WAVELET BASED SPACE-TIME ADAPTIVE NUMERICAL METHOD FOR PARTIAL DIFFERENTIAL EQUATIONS
E. Bacry, S. Mallat, November 1991

We describe a time and space adaptive numerical method for solving partial differential equations, based on wavelet orthonormal bases. The multiresolution structure of wavelet orthonormal bases provide a simple approach to adapt computation refinements to the local regularity of the solution High resolution computations are performed only in regions where singularities or sharp transitions occur. For many evolution equations, to maintain the stability and precision of the numerical scheme, it is necessary to adapt the time steps to the spatial resolution. We describe an algorithm that modifies the time discretization at each resolution, depending upon the properties of the partial differential equation. The stability of this space-time adaptive scheme is studied for the heat equation and the linear advection equation. We explain how this algorithm

applies to the one-dimensional Burgers equation with periodic boundary conditions. Numerical results on the precision and complexity of the algorithm are given.

592-R258 CHARACTERIZATION OF SIGNALS FROM MULTISCALE EDGES
S. Mallat, S. Zhong, December 1991

A multiscale Canny edge detection is equivalent to finding the local maxima of a wavelet transform. We study the properties of multiscale edges through the wavelet formalism. For pattern recognition, one often needs to discriminate different types of edges. We show that the evolution of wavelet local maxima across scales characterize the local shape of irregular structures. Numerical descriptors of edge types are derived. The completeness of a multiscale edge representation is also studied. We describe an algorithm that reconstructs a close approximation of one and two-dimensional signals from their multiscale edges. For images, the reconstruction errors are below our visual sensitivity. As an application, we implement a compact image coding algorithm that selects important edges and compresses the image data by factors over 30.

593-R259 CORTICAL HYPERCOLUMNS AND THE TOPOLOGY OF RANDOM ORIENTATION MAPS
E. Schwartz, A. Rojer, December 1991

If a set of points is assigned a random choice of orientations, and then is locally smoothed, patterns appear consisting of orientation "vortices" containing central regions of weak amplitude. This pattern is reminiscent of the "puff-extra-puff" pattern of orientation tuning in monkey visual cortex, and of the "pin-wheels" that have been recently observed in cat area 18. We show in this paper that the formation of such patterns is a consequence of the topology (non-retraction of $S^1 \rightarrow \mathbb{R}^2$) and stochastic properties of random orientation distributions. These observations provide insight into experimental, theoretical, and descriptive aspects of functional architecture in visual cortex. Specifically:

1.) Observation of "vortex" structure in any visually excitable neural tissue must be evaluated in the light of the null hypothesis that these patterns are readily formed from random orientation distributions. This is particularly problematic for observational techniques which bin together local neuronal orientation response (e.g. multiple unit recording, 2-deoxyglucose, optical dye/blood flow). We show that similar patterns could in principle be artifactually observed in the retina, LGN, or any other visually responsive tissue.

2.) Since observed hypercolumn patterns in cat and monkey cortex are similar to noise patterns in orientation maps, it is possible that the developmental specification of the brain for creating hypercolumns need be nothing more than a local correlation of orientation tuning in an otherwise random, or weakly correlated map.

3.)Since similar applications of (band-pass) filtered noise can be shown to accurately reproduce the ocular dominance column patterns of cats and monkeys, it appears that the columnar structure of visual cortex in higher vertebrates may be described most economically as the result of a Wiener process: a linear filter applied to an array of random variables.

594 AN OPTIMAL ON-LINE SCHEDULING ALGORITHM FOR OVERLOADED REAL-TIME SYSTEMS
G. Koren, D. Shasha, January 1992

Every task in a real-time system has a deadline by which time it should complete and a value that it obtains only if it completes by its deadline. The problem is to design an on-line scheduling algorithm (i.e., the scheduler has no knowledge of a task until it is released) that maximizes the obtained value.

When such a system is underloaded (i.e. there exists a schedule for which all tasks meet their deadlines), Dertouzos showed that the earliest deadline first algorithm will achieve 100% of the possible value. Locke showed that earliest deadline first performs very badly when the system is overloaded and proposed heuristics to deal with overload. Baruah et al. showed that the best that an on-line algorithm can guarantee is $\frac{1}{(1+\sqrt{k})^2}$ of the value that an offline (clairvoyant) algorithm can obtain where k is the ratio between the highest value density and the lowest value density task in the system (where the value density of a task is its value divided by its computation time).

This paper presents three results:

- An on-line algorithm that achieves the above best possible guarantee when computation time is known exactly when the task is released and a task obtains no value if it completes after its deadline (the model studied by Baruah et al.). The algorithm also obtains 100% of the possible value when the system is underloaded.
- The same guarantee in certain situations where tasks still obtain some value if they complete after their deadlines.
- A weaker guarantee of $\frac{1}{(1+\sqrt{k})^2+(\epsilon \cdot k)(1+\sqrt{k})+1}$ of the value of a clairvoyant algorithm when the exact computation time

of a task is not given upon its release. That is, for every task T, an upper bound on its possible computation time denoted by c_{max}, is given. The actual computation time of T denoted by c satisfies: $(1-\epsilon)\cdot c_{max} \leq c \leq c_{max}$

595 MULTIPLICATIVE SCHWARZ ALGORITHMS FOR SOME NONSYMMETRIC AND INDEFINITE PROBLEMS
X.-C. Cai, O. Widlund, February 1992

The classical Schwarz alternating method has recently been generalized in several directions. This effort has resulted in a number of new powerful domain decomposition methods for elliptic problems, in new insight into multigrid methods and in the development of a very useful framework for the analysis of a variety of iterative methods. Most of this work has focused on positive definite, symmetric problems. In this paper a general framework is developed for multiplicative Schwarz algorithms for nonsymmetric and indefinite problems. Several applications are then discussed including two- and multi-level Schwarz methods and iterative substructuring algorithms. Some new results on additive Schwarz methods are also presented.

596-R260 SHORTEST PATHS FOR LINE SEGMENTS
C. Icking, G. Rote, E. Welzl, C. Yap, February 1992

We study the problem of shortest paths for a line segment in the plane. As a measure of the distance traversed by a path, we take the average curve length of the orbits of prescribed points on the line segment. This problem is nontrivial even in free space (i.e., in the absence of obstacles). We characterize all shortest paths of the line segment moving in free space under the measure d_2, the average orbit length of the two endpoints.

The problem of d_2 optimal motion has been solved by Gurevich and also by Dubovitskij, who calls it Ulam's problem. Unlike previous solutions, our basic tool is Cauchy's surface-area formula. This new approach is relatively elementary, and yields new insights.

597 SEMANTIC ANALYSES FOR STORAGE MANAGEMENT OPTIMIZATIONS IN FUNCTIONAL LANGUAGE IMPLEMENTATIONS
Y. G. Park, February 1992

One of the major overheads in implementing functional languages is the storage management overhead due to dynamic allocation and automatic reclamation of indefinite-extent storage. This dissertation investigates the problems of statically inferring lifetime information about dynamically-allocated objects in *higher-order polymorphic* functional languages, both *strict* and *non-strict*, and of applying that information to reduce the storage management overhead.

We have developed a set of compile-time semantic analyses for a higher-order, monomorphic, strict functional language based on denotational semantics and abstract interpretation. They are 1) *escape analysis*, which provides information about the relative lifetimes of objects such as arguments and local objects defined within a function with respect to an activation of the function call, 2) *refined escape analysis* which, as a refinement of escape analysis, provides information about the lifetimes of components of aggregate structures, and 3) *reference escape analysis* which provides information about the relative lifetimes of references created within a function with respect to an activation of the function.

We also have developed a compile-time semantic analysis called *order-of-demand analysis* for higher-order, monomorphic, non-strict functional languages, which provides information about the order in which the values of bound variables are demanded and thus allows one to compute a range of information including strictness, evaluation-order, and evaluation-status information.

Using the notion of *polymorphic invariance*, we describe a method for analyzing a polymorphic language by using the analyses for a monomorphic language. We then extend those analyses for a strict language to a non-strict language using non-strict program transformation and evaluation-status information.

Based on statically inferred escape information, we propose a combination of storage management optimization techniques including stack allocation, explicit reclamation, in-place reuse, reference counting elimination, block allocation/reclamation, and improving generational garbage collection.

598 SMOOTH TIME-PERIODIC FEEDBACK SOLUTIONS FOR NONHOLONOMIC MOTION PLANNING
L. Gurvits and Z. X. Li February 1992

In this paper, we present an algorithm for computing time-periodic feedback solutions for nonholonomic motion planning with collision-avoidance. For a first-order Lie bracket system, we begin by computing a holonomic collision-free path using the potential field method. Then, we compute a nonholonomic path approximating the collision-free path within a predetermined

bound. For this we first solve for extended inputs of an extended system using Lie bracket completion vectors. We then use averaging techniques to calculate the asymptotic trajectory of the nonholonomic system under application of a family of highly-oscillatory inputs. Comparing the limiting trajectories with the extended system we obtain a system of nonlinear equations from which the desired admissible control inputs can be solved. For higher-order Lie bracket systems we use multi-scale averaging and apply recursively the algorithm for first-order Lie bracket systems. Based on averaging techniques we also provide error bounds between a nonholonomic system and its averaged system.

599 INFINITE LOOPS IN FINITE TIME: SOME OBSERVATIONS
E. Davis, February 1992

Formulating physical theories using real-valued space and time often raises the difficult issue of clustered variation: states that change their value infinitely often in a finite interval, or infinite sequences of events that occur in a finite interval. Some physical theories permit nonsensical models unless clustered variation is prohibited; others force clustered variation to occur. This paper surveys a number of technical issues involved with clustered variation, without attempting to find a uniform treatment of the issue. We present examples of theories where clustered variation must be allowed and of theories where it must be prohibited. We discuss the meaning of plans containing infinite series of events. We present constraints on real-valued fluents and spatial fluents that guarantee discrete variation of important associated states. We show that requiring discrete variation may be inconsistent with allowing infinitely many objects in the universe, even if each object is individually well-behaved. We show how the need for requiring discrete variation may depend subtly on the particular ontology chosen for the domain.

600-R262 A SOLUTION TO KRONECKER'S PROBLEM
G. Gallo, B. Mishra, March 1992

"Kronecker believed God made the natural numbers and all the rest was man's work. We only know of this opinion by hearsay evidence, however, and his paper 'Ueber den Zahlbegriff' indicates to me that he thought God made a bit more: 'Buchstabenrechnung', or calculation with letters. In modern terms, Kronecker seems to envisage a cosmic computer which computes not just with natural numbers, but with polynomials with natural number coefficients (in any number of indeterminates)." (see the article by Harold M. Edwards, entitled "Kronecker's Views on the Foundations of Mathematics.")

The preceding notion of constructive mathematics, as postulated by Kronecker, can be expanded further and leads to some very interesting algorithmic problems. While it is quite apparent that Kronecker regarded these algorithmic questions as at the heart of his formulation of mathematics, it is unclear whether Kronecker had been able to resolve these questions in a satisfactory manner.

The present notes make a fresh attempt at resolving the algorithmic questions raised by Kronecker, while remaining faithful to the notion of constructivity espoused by Kronecker.

601-R264 A MINIATURE PAN-TILT ACTUATOR: THE SPHERICAL POINTING MOTOR
B. Bederson, R. Wallace, E. Schwartz, April 1992

A pan-tilt mechanism is a computer-controlled actuator designed to point an object such as a camera sensor. For applications in active vision, a pan-tilt mechanism should be accurate, fast, small, inexpensive and have low power requirements. We have designed and constructed a new type of actuator meeting these requirements , which incorporates both pan and tilt into a single, two degree of freedom device. The Spherical Pointing Motor (SPM) consists of three orthogonal motor windings in a permanent magnetic field, configured to move a small camera mounted on a gimbal. It is an absolute positioning device and is run open loop. The SPM is capable of panning and tilting a load of 15 grams, for example a CCD image sensor, at rotational velocities of several hundred degrees per second with an accuracy of 0.15 degrees. We have also built a miniature camera consisting of a single CCD sensor chip and miniature lens assembly that fits on the rotor of this motor. In this paper, we discuss the theory of the SPM which includes its basic electromagnetic principles and derive the relationship between applied currents and resultant motor position. We present an automatic calibration procedure and discuss open and closed-loop control strategies. Finally, we present the physical characteristics and results of our prototype.

602 AN $O(m\log n)$-TIME ALGORITHM FOR THE MAXIMAL PLANAR SUBGRAPH PROBLEM
J. Cai, X. Han, R. Tarjan, April 1992

Based on a new version of Hopcroft and Tarjan's planarity testing algorithm, we develop an $O(m\log n)$-time algorithm to

find a maximal planar subgraph.

603 COUNTING EMBEDDINGS OF PLANAR GRAPHS USING DFS TREES
J. Cai, April 1992

Previously counting embeddings of planar graphs used P-Q trees and was restricted to biconnected graphs. Although the P-Q tree approach is conceptually simple, its implementation is complicated. In this paper we solve this problem using DFS trees, which are easy to implement. We also give formulas that count the number of embeddings of general planar graphs (not necessarily connected or biconnected) in $O(n)$ arithmetic steps, where n is the number of vertices of the input graph. Finally, our algorithm can be extended to generate all embeddings of a planar graph in linear time with respect to the output.

604 MORE EFFICIENT BOTTOM-UP MULTI-PATTERN MATCHING IN TREES
J. Cai, R. Paige, R. Tarjan, April 1992

Pattern matching in trees is fundamental to a variety of programming language systems. However, progress has been slow in satisfying a pressing need for general purpose pattern matching algorithms that are efficient in both time and space. We offer asymptotic improvements in both time and space to Chase's bottom-up algorithm for pattern preprocessing. A preliminary implementation of our algorithm runs ten times faster than Chase's implementation on the hardest problem instances. Our preprocessing algorithm has the advantage of being on-line with respect to pattern additions and deletions. It also adapts to favorable input instances, and on Hoffmann and O'Donnell's class of Simple Patterns, it performs better than their special purpose algorithm tailored to this class. We show how to modify our algorithm using a new decomposition method to obtain a space/time tradeoff. Finally, we trade a log factor in time for a linear space bottom-up pattern matching algorithm that handles a wide subclass of Hoffmann and O'Donnell's Simple Patterns.

605-R265 COMPLEXITY ISSUES IN COMPUTATIONAL ALGEBRA
G. Gallo, May 1992

The ideal membership problem for the ring of multivariate polynomials is a central problem in Computational Algebra. Relatively tight computational complexity bounds for this problem are known in the case of polynomials with coefficients in a field. After reviewing these results we give an algorithm, together with an upper bound on its complexity, for the solution of the membership problem in the case of polynomials with integer coefficients. This result is obtained adapting Buchberger's algorithm to the integer case. As an application, we also provide a more general upper bound on the length of strictly ascending chain of ideals in the ring Z[x1,...,xn].

Many applications of Computational Algebra, however, do not require the complete solution of the membership problems and alternative approaches are available. In this thesis we survey the method of the characteristic sets, originally introduced by Ritt in the forties and now widely applied, particularly in Elementary Geometry Theorem Proving. We present optimal algorithms for computing a characteristic set with simple-exponential sequential and polynomial parallel time complexities.

We finally present an attempt to generalize some of the constructive methods of Commutative Algebra to the Algebra of differential polynomials. Rings of differential polynomials do not resemble their purely algebraic counterparts: we prove that there exist non-recursive differential ideals in Zx and hence that, in general, no effective method can be devised to solve the membership problem in this case. However special classes of ideals can be algorithmically approached and toward this goal, we generalize the concept of H-basis, first introduced by Macaulay for algebraic ideals, to differential rings.

606 DOMAIN DECOMPOSITION ALGORITHMS WITH SMALL OVERLAP
M. Dryja, O. Widlund, May 1992

Numerical experiments have shown that two-level Schwarz methods often perform very well even if the overlap between neighboring subregions is quite small. This is true to an even greater extent for a related algorithm, due to Barry Smith, where a Schwarz algorithm is applied to the reduced linear system of equations that remains after that the variables interior to the subregions have been eliminated. In this paper, a supporting theory is developed.

607 GMRES/CR AND ARNOLDI/LANCZOS AS MATRIX APPROXIMATION PROBLEMS
A. Greenbaum, L. Trefethen, May 1992

The GMRES and Arnoldi algorithms, which reduce to the CR and Lanczos algorithms in the symmetric case, both minimize $\|p(A)b\|$ over polynomials p of degree n. The difference is that p is normalized at $z = 0$ for GMRES and at $z = \inf$ for Arnoldi. Analogous "ideal GMRES" and "ideal Arnoldi" problems are obtained if one removes b from the discussion and minimizes $\|p(A)\|$ instead. Investigation of these true and ideal approximation problems gives insight into how fast GMRES

converges and how the Arnoldi iteration locates eigenvalues.

608 MATRICES THAT GENERATE THE SAME KRYLOV RESIDUAL SPACES
A. Greenbaum, Z. Strakos, May 1992

Given an n by n nonsingular matrix A and an n-vector v, we consider the spaces of the form $AK_k(A,v)$, $k = 1, ..., n$, where $K_k(A,v)$ is the k^{th} Krylov space, equal to $span\{v, Av, ..., A^{k-1}v\}$. We characterize the set of matrices B that, with the given vector v, generate the same spaces; i.e., those matrices B for which $BK_k(B,v) = AK_k(A,v)$, for all $k = 1, ..., n$. It is shown that any such sequence of spaces can be generated by a unitary matrix. If zero is outside the field of values of A, then there is a Hermitian positive definite matrix that generates the same spaces, and, moreover, if A is close to Hermitian then there is a nearby Hermitian matrix that generates the same spaces. It is also shown that any such sequence of spaces can be generated by a matrix having any desired eigenvalues.

Implications about the convergence rate of the GMRES method are discussed. A new proof is given that if zero is outside the field of values of A, then convergence of the GMRES algorithm is strictly monotonic. It is shown that if A is close to Hermitian, then the eigenvalues of A essentially determine the behavior of the GMRES iteration but that, in general, eigenvalue information alone is never sufficient to ensure rapid convergence of the GMRES algorithm.

609 MULTISET DISCRIMINATION – A METHOD FOR IMPLEMENTING PROGRAMMING LANGUAGE SYSTEMS WITHOUT HASHING
J. Cai, R. Paige, June 1992

It is generally assumed that hashing is essential to many algorithms related to efficient compilation; e.g., symbol table formation and maintenance, grammar manipulation, basic block optimization, and global optimization. This paper questions this assumption, and initiates development of an efficient alternative compiler methodology without hashing or sorting. Underlying this methodology are several generic algorithmic tools, among which special importance is given to *Multiset Discrimination*, which partitions a multiset into blocks of duplicate elements. We show how multiset discrimination, together with other tools, can be tailored to rid compilation of hashing without loss in asymptotic performance. Because of the simplicity of these tools, our results may be of practical as well as theoretical interest. The various applications presented culminate with a new algorithm to solve iterated strength reduction folded with useless code elimination that runs in worst case asymptotic time and auxiliary space $\Theta(|L| + |L^*|)$, where $|L|$ and $|L^*|$ represent the lengths of the initial and optimized programs respectively. The previous best solution due to Cocke and Kennedy takes $\Omega(|L|^3|L^*|)$ hash operations in the worst case.

610 HESFCN — A FORTRAN PACKAGE OF HESSIAN SUBROUTINES FOR TESTING NONLINEAR OPTIMIZATION SOFTWARE
V. Averbukh, S. Figueroa, T. Schlick, June 1992

We report the development of Hessian FORTRAN routines for testing unconstrained nonlinear optimization. An existing package, "Algorithm 566" (J. Moré, B. G. Garbow, and K. Hillstrom, *ACM Trans. Math. Softw.* **7**, 14–41 and 136–140, 1981) provides function and gradient subroutines of 18 test functions for multivariate minimization. Our supplementary Hessian segments will enable users to test optimization software that requires second derivative information. Eigenvalue analysis throughout the minimization will also be possible in the goal of better understanding minimization progress by different algorithms and the relation of progress to eigenvalue distribution and condition number.

611-R266 MINIATURE SPACE-VARIANT ACTIVE VISION SYSTEM: CORTEX-I
B. Bederson, July 1992

We have developed a prototype miniaturized active vision system whose sensor architecture is based on a logarithmically structured space-variant pixel geometry. A space-variant image's resolution changes across the image. Typically, the central part of the image has a very high resolution, and the resolution falls off gradually in the periphery. Our system integrates a miniature CCD-based camera, pan-tilt actuator, controller, general purpose processors and display. Due to the ability of space-variant sensors to cover large work-spaces, yet provide high acuity with an extremely small number of pixels, space-variant active vision system architectures provide the potential for radical reductions in system size and cost. We have realized this by creating an entire system that takes up less than a third of a cubic foot. In this thesis, we describe a prototype space-variant active vision system ($\mathcal{V}$) which performs such tasks as tracking moving objects and license plate reading, and functions as a video telephone.

We report on the design and construction of the camera (which is $8 \times 8 \times 8$mm), its readout, and a fast mapping algorithm to convert the uniform image to a space-variant image. We introduce a new miniature pan-tilt actuator, the Spherical Pointing

Motor (SPM), which is $4 \times 5 \times 6$cm. The basic idea behind the SPM is to orient a permanent magnet to the magnetic field induced by three orthogonal coils by applying the appropriate ratio of currents to the coils. Finally, we present results of integrating the system with several applications. Potential application domains for systems of this type include vision systems for mobile robots and robot manipulators, traffic monitoring systems, security and surveillance, telerobotics, and consumer video communications.

The long-range goal of this project is to demonstrate that major new applications of robotics will become feasible when small low-cost machine vision systems can be mass-produced. This notion of "commodity robotics" is expected to parallel the impact of the personal computer, in the sense of opening up new application niches for what has until now been expensive and therefore limited technology.

612-R267 THE SPIDERWEB ALGORITHM FOR SURFACE CONSTRUCTION FROM MEDICAL VOLUME DATA: GEOMETRIC PROPERTIES OF ITS SURFACE

D. Karron, J. Cox, B. Mishra, July 1992

The "SpiderWeb" algorithm is a set of rules to construct a triangular mesh isosurface (simplicial complex) from density data. The algorithm is robust in the presence of noisy data in the sense that the surface grown will be an oriented two manifold surface. Further, the algorithm constructs the triangulation using only local triangle-spinning rules, thus needing no additional global information. The triangle spinning rules use local geometric symmetries that translate into face-hit rules. These rules are not artifacts of the algorithm or any data assumptions, but rather inherent in the nature of density data, sampled at a given resolution or spatial frequency. SpiderWeb does not require ad hoc rules to deal with inconsistencies caused by so-called nongeometric noise. Thus the algorithm conforms to and strongly exploits previously unrecognized geometric properties of density data, while requiring minimal assumptions on the distribution of such data.

613 COMPUTING READY SIMULATIONS EFFICIENTLY

B. Bloom, R. Paige, September 1992

Ready simulation is the finest fully abstract notion of process equivalence in the CCS setting. We give an $O(mn + n^2)$-time algorithm for deciding when n-state, m-transition processes are ready similar; a substantial improvement on the $\Theta(mn^6)$ algorithm presented in Bard Bloom's Thesis.

614 SOME SCHWARZ ALGORITHMS FOR THE P-VERSION FINITE ELEMENT METHOD

L. Pavarino, September 1992

Domain decomposition methods based on the Schwarz framework were originally proposed for the h-version finite element method for elliptic problems. In this paper, we consider instead the p-version, in which increased accuracy is achieved by increasing the degree of the elements while the mesh is fixed. We consider linear, scalar, self adjoint, second order elliptic problems and quadrilateral elements in the finite element discretization. For a class of overlapping additive Schwarz methods, we prove a constant bound, independent of the degree p and the number of elements N, for the condition number of the iteration operator. This optimal result holds in two and three dimensions for additive and multiplicative schemes, as well as variants on the interface. We then study local refinement for the same class of overlapping methods in two dimensions. A constant bound holds under certain hypotheses on the refinement region, while in general an almost optimal bound with logarithmic growth in p is obtained.

615 SOME RECENT RESULTS ON SCHWARZ TYPE DOMAIN DECOMPOSITION ALGORITHMS

M. Dryja, O. Widlund September 1992

Numerical experiments have shown that two-level Schwarz methods, for the solution of discrete elliptic problems, often perform very well even if the overlap between neighboring subregions is quite small. This is true to an even greater extent for a related algorithm, due to Barry Smith, where a Schwarz algorithm is applied to the reduced linear system of equations that remains after that the variables interior to the subregions have been eliminated. A supporting theory is outlined.

616 DOMAIN DECOMPOSITION ALGORITHMS FOR THE P-VERSION FINITE ELEMENT METHOD FOR ELLIPTIC PROBLEMS

L. Pavarino, September 1992

Domain decomposition algorithms based on the Schwarz framework were originally proposed for the h-version finite element method for elliptic problems. In this thesis, we study some Schwarz algorithms for the p-version finite element method, in

which increased accuracy is achieved by increasing the degree p of the elements while the mesh is fixed. These iterative algorithms, often of conjugate gradient type, are both parallel and scalable, and therefore very well suited for massively parallel computing.

We consider linear, scalar, self adjoint, second order elliptic problems and quadrilateral elements in the finite element discretization. For a class of overlapping methods, we prove a constant bound, independent of the degree p, the mesh size H and the number of elements N, for the condition number of the iteration operator. This optimal result holds in two and three dimensions for additive and multiplicative schemes, as well as variants on the interface.

We consider then local refinement for the same class of overlapping methods in two dimensions. Optimal bounds are obtained under certain hypothesis on the choice of refinement points, while in general almost optimal bounds with logarithmic growth in p are obtained. In the analysis of these local refinement methods, we prove some results of independent interest, such as a polynomial discrete Sobolev inequality and a bounded decomposition of discrete harmonic polynomials.

Iterative substructuring methods in two dimensions are also considered. We use the additive Schwarz framework to prove almost optimal bounds as in the h-version finite element method.

Results of numerical experiments, confirming the theoretical results, are conducted in two dimensions for model problems.

617 AN EFFICIENT ALGORITHM TO DECIDE PLANAR IMMOBILITY

M. Teichmann, B. Mishra, October 1992

(Out of Print)

In robotics, it is necessary to immobilize an object by placing "fingers" (clamps or fixtures) whose contacts with the object prevent its motion. Depending on the areas of applications, the problem has been variously labelled as grasping, workholding, part-mating.

Also, there have been a wide class of definitions of grasping: for instance, force/torque closure grasps, form closure grasps, stable grasps, zero total freedom and recently, a new simple geometric formulation, called immobility.

In this paper, we propose a characterization of the sets of immobilizing points, which leads to an algorithm to test whether a given set immobilizes a polygon. This algorithm runs in time linear in the number of points in the set.

618 ON THE NUMBER OF CELLS DEFINED BY A SET OF POLYNOMIALS

R. Pollack, M. Roy, October 1992

We show that the sum of the number of connected components of all non- empty realizations of sign conditions of s polynomials in n variables, each of degree at most d, is bounded by $(O(sd/n))^n$. The previously known bound, based on a result of J. Heintz, was $(sd)^{O(n)}$ (see for example [7]).

619 MATCHING PURSUITS WITH TIME-FREQUENCY DICTIONARIES

S. Mallat, Z. Zhang, October 1992, Revised August 1993

We introduce an algorithm, called matching pursuit, that decomposes any signal into a linear expansion of waveforms that are selected from a redundant dictionary of functions. These waveforms are chosen in order to best match the signal structures. Matching pursuits are general procedures to compute adaptive signal representations. With a dictionary of Gabor functions, a matching pursuit defines an adaptive time-frequency transform. We derive a signal energy distribution in the time-frequency plane, which does not include interference terms, unlike Wigner and Cohen class distributions [?]. A matching pursuit isolates the signal structures that are coherent with respect to a given dictionary. An application to pattern extraction from noisy signals is described. We compare a matching pursuit decomposition with a signal expansion over an optimized wavepacket orthonormal basis, selected with the algorithm of Coifman and Wickerhauser.

620 BACKWARD ANALYSIS FOR HIGHER-ORDER FUNCTIONS USING INVERSE IMAGES

T-R. Chuang, B. Goldberg, November 1992

We propose a method for performing backward analysis on higher–order functional programming languages based on computing inverse images of functions over abstract domains. This method can be viewed as abstract interpretation done backward. Given an abstract semantics which supports forward analysis, we can transform it into an abstract semantics which performs backward analysis. We show that if the original abstract semantics is correct and computable, then the transformed version of the abstract semantics is also correct and computable.

More specifically, given a forward abstract semantics of a higher–order functional language which is expressed in terms of Scott–closed powerdomains, we derive an backward abstraction semantics which is expressed in terms of Scott–open powerdomains. The derivation is shown to be correct and the relationships between forward analysis and backward analysis

is established.

We apply this method to the classic strictness analysis in functional languages and obtain promising results. We show that the time complexity of inverse image based backward analysis is not much worse than the forward analysis from which it is derived. We then compare this work with previous works of Wadler and Hughes (1987), Hughes (1988), and Burn (1990), showing that some special restrictions and constructions in previous works have natural interpretation in the Scott–closed/Scott–open powerdomain framework. A brief outline of applying the inverse image method to other backward semantics analysis is also given.

621 STATISTICAL APPROACH TO AFFINE INVARIANT MATCHING WITH LINE FEATURES
F. Tsai, November 1992

One of the most important goals of computer vision research is *object recognition*. Currently, most object recognition algorithms assume reliable quality of image segmentation, which in practice is often not the case. This report examines the combination of the *Hough Transform* with a variation of *Geometric Hashing* as a technique for model-based object recognition in *seriously degraded* single intensity images.

There is recently much focus on the performance analysis of geometric hashing. However, to our knowledge, all of them are focusing on applying the paradigm to point features and show that the technique is sensitive to noise. There is as yet no exploration of line features. In this report, we use lines as the primitive features to compute the *geometric invariant*s for fast indexing into the geometric hash table containing the pre-processed model information. In addition, we analytically determine the effect of perturbations of line parameters on the computed *invariant* for the case where models are allowed to undergo affine transformation.

We have implemented the system with a series of experiments on polygonal objects, which are modeled by lines. It shows that the technique is noise resistant and suitable in an environment containing many occlusions.

622 POLYMORPHIC TYPE INFERENCE AND ABSTRACT DATA TYPES
K. Läufer, December 1992

Many statically-typed programming languages provide an abstract data type construct, such as the package in Ada, the cluster in CLU, and the module in Modula2. However, in most of these languages, instances of abstract data types are not first-class values. Thus they cannot be assigned to a variable, passed as a function parameter, or returned as a function result.

The higher-order functional language ML has a strong and static type system with parametric polymorphism. In addition, ML provides type reconstruction and consequently does not require type declarations for identifiers. Although the ML module system supports abstract data types, their instances cannot be used as first-class values for type-theoretic reasons.

In this dissertation, we describe a family of extensions of ML. While retaining ML's static type discipline, type reconstruction, and most of its syntax, we add significant expressive power to the language by incorporating first-class abstract types as an extension of ML's free algebraic datatypes. In particular, we are now able to express

- multiple implementations of a given abstract type,
- heterogeneous aggregates of different implementations of the same abstract type, and
- dynamic dispatching of operations with respect to the implementation type.

Following Mitchell and Plotkin, we formalize abstract types in terms of existentially quantified types. We prove that our type system is semantically sound with respect to a standard denotational semantics.

We then present an extension of Haskell, a non-strict functional language that uses type classes to capture systematic overloading. This language results from incorporating existentially quantified types into Haskell and gives us first-class abstract types with type classes as their interfaces. We can now express heterogeneous structures over type classes. The language is statically typed and offers comparable flexibility to object-oriented languages. Its semantics is defined through a type-preserving translation to a modified version of our ML extension.

We have implemented a prototype of an interpreter for our language, including the type reconstruction algorithm, in Standard ML.

623 RESIDUAL RELATIONSHIPS WITHIN THREE PAIRS OF ITERATIVE ALGORITHMS FOR SOLVING $Ax = b$
J. Cullum, A. Greenbaum, February 1993

In [Brown, SIAM Jour. Sci. Stat. Comput. 20, pp. 58-78], a relationship was established between the norms of the residuals produced by the Arnoldi algorithm and those produced by the GMRES algorithm. This, combined with the relationship between successive GMRES residual norms obtained in [Saad and Schultz, SIAM Jour. Sci. Stat. Comput. 7, pp. 856-869], shows that when the GMRES algorithm is converging rapidly the norms of the Arnoldi residuals are approximately equal to those of the GMRES residuals, while when the GMRES algorithm is converging very slowly the norms of the Arnoldi residuals are significantly larger than those of the GMRES residuals. We state this relationship in a slightly different form, showing the precise connection between the residual norm reduction at a particular GMRES step and the ratio of the Arnoldi to GMRES residual norms at that step. Assuming exact arithmetic, we show that a similar relationship exists between the residuals of the biconjugate gradient algorithm and the "quasi-residuals" of the QMR algorithm, where the "quasi-residuals" are the vectors whose norms are actually minimized by the QMR algorithm. It is demonstrated numerically that the norms of the quasi-residuals are of the same order of magnitude as the norms of the actual residuals.

We then consider this relationship for the corresponding real symmetric procedures, Lanczos and MINRES. We prove that for a certain class of real symmetric problems the relationship holds to a close approximation, even in finite precision arithmetic. In particular, we show that for well-conditioned problems the nonsymmetric bidiagonalization procedures considered in [Cullum, IBM Res. Rept. RC 18084] fall into this class and hence satisfy the residual relationship to a close approximation in finite precision arithmetic.

624 A HYBRID ALGORITHM FOR OPTIMIZING EIGENVALUES OF SYMMETRIC DEFINITE PENCILS

J. Haeberly, M. Overton, February 1993

We present an algorithm for the optimization of the maximum eigenvalue of a symmetric definite pencil depending affinely on a vector of parameters. The algorithm uses a hybrid approach, combining a scheme based on the method of centers, developed by Boyd and El Ghaoui, with a new quadratically convergent local scheme. A convenient expression for the generalized gradient of the maximum eigenvalue of the pencil is also given, expressed in terms of a dual matrix. The algorithm computes the dual matrix which establishes the optimality of the computed solution.

625 USING LINE INVARIANTS FOR OBJECT RECOGNITION BY GEOMETRIC HASHING

F. Tsai, February 1993

Geometric Hashing is a model-based object recognition technique for detecting objects which can be partially overlapping or partly occluded. It precompiles, from local geometric features, redundant transformation-invariant information of the models in a hash table and uses the invariants computed from a scene for fast indexing into the hash table to hypothesize possible matches between object instances and object models during recognition.

In its simplest form, the geometric hashing method assumes relatively noise-free data and is applied to objects with points as local features. However, extracting of the locations of point features is inherently error-prone and the analysis of geometric hashing on point sets shows considerable noise sensitivity. Line features can generally be extracted with greater accuracy.

We investigate the use of line features for geometric hashing applied to 2-D (or flat 3-D) object recognition and derive, from a combination of line features, invariants for lines under various geometric transformations.

626 SCHWARZ METHODS OF NEUMANN-NEUMANN TYPE FOR THREE-DIMENSIONAL ELLIPTIC FINITE ELEMENT PROBLEMS

M. Dryja, O. Widlund, March 1993

Several domain decomposition methods of Neumann-Neumann type are considered for solving the large linear systems of algebraic equations that arise from discretizations of elliptic problems by finite elements. We will only consider problems in three dimensions. Several new variants of the basic algorithm are introduced in a Schwarz method framework that provides tools which have already proven very useful in the design and analysis of other domain decomposition and multi-level methods.

The Neumann-Neumann algorithms have several advantages over other domain decomposition methods. The subregions, which define the subproblems, only share the boundary degrees of freedom with their neighbors. The subregions can also be of quite arbitrary shape and many of the major components of the preconditioner can be constructed from subprograms available in standard finite element program libraries. However, in its original form, the algorithm lacks a mechanism for global transportation of information and its performance therefore suffers when the number of subregions increases. In the new variants of the algorithms, considered in this paper, the preconditioners include global components, of low rank, to overcome this difficulty. Bounds are established for the condition number of the iteration operator, which are independent of the number of subregions, and depend only polylogarithmically on the number of degrees of freedom of individual local subproblems. Results are also given for problems with arbitrarily large jumps in the coefficients across the interfaces separating

the subregions.

627 SECOND DERIVATIVES FOR OPTIMIZING EIGENVALUES OF SYMMETRIC MATRICES
M. Overton, R. Womersley, March 1993

Let A denote an $n \times n$ real symmetric matrix-valued function depending on a vector of real parameters, $x \in \Re^m$. Assume that A is a twice continuously differentiable function of x, with the second derivative satisfying a Lipschitz condition. Consider the following optimization problem: minimize the largest eigenvalue of $A(x)$. Let x^* denote a minimum. Typically, the maximum eigenvalue of $A(x^*)$ is multiple, so the objective function is not differentiable at x^*, and straightforward application of Newton's method is not possible. Nonetheless, the formulation of a Newton-based method with local quadratic convergence is possible. This formulation is derived by parameterizing the eigenvectors using a matrix exponential representation, introducing an appropriate Lagrangian function. Consideration of the Hessian of this Lagrangian function leads to the second derivative matrix used by Newton's method. The derivation of the method and the proof of quadratic convergence take into account the following key fact: the number of variables used in the eigenvector parameterization drops, in the limit, by $t(t-1)/2$, where t is the multiplicity of the maximum eigenvalue of $A(x^*)$. In the special case $t = 1$, the maximum eigenvalue is a smooth function and the method reduces to the standard Newton method. The same ideas are applicable to a wide variety of other eigenvalue and singular value optimization problems.

628 THE COMPLEXITY OF RESOLVENT RESOLVED
G. Gallo, B. Mishra, March 1993

The concept of a resolvent of a prime ideal was originally introduced by J.F. Ritt along with the notion of a characteristic set. The motivation for studying resolvents comes from its connections with the birational isomorphisms that describe structures of irreducible algebraic varieties by means of an equivalent hypersurface and a one-to-one rational map. As a result, these ideas have a wide range of applications in such areas as solid modeling, computer aided design and manufacturing. An algorithm to compute the resolvent by means of characteristic sets was first proposed by Ritt. This and some related algorithms have resurfaced as interest in resolvent structures have grown, spurred by its applicability.

Unfortunately, the algebraic complexity of the resolvent and the computational complexity of the associated algorithms have never been satisfactorily explored. In this paper, we give single exponential upper and lower bounds for the degrees of the resolvent and its associated parametrizing polynomials. We also show that the resolvent can be computed deterministically in single exponential sequential and polynomial parallel time complexity. All previous algorithms for resolvent had relied on a random choice of certain extraneous parameters.

629 TWO-LEVEL SCHWARZ METHODS FOR NONCONFORMING FINITE ELEMENTS AND DISCONTINUOUS COEFFICIENTS
M. Sarkis, March 1993

A two-level domain decomposition methods are developed for a simple nonconforming approximation of second order elliptic problems. A bound is established for the condition number of these iterative methods, which grows only logarithmically with the number of degrees of freedom in each subregion. This bound holds for two and three dimensions and is independent of jumps in the value of the coefficients.

630 THE CELL PROGRAMMING LANGUAGE
P. Agarwal, March 1993

We describe the Cell Programming Language (CPL), which we have designed to write programs that mimic the life of a biological cell. The aim is to study the complex interactions between cells which lead to diverse shapes of cellular aggregates.

Each cell is treated as a two-dimensional homogeneous polygon with a specific area. A cell goes through a series of states in its lifetime. In each state, one can specify the cell's growth rate; information about cell division and cell differentiation; the chemical constituents of the cell and their interactions; and cell motion. This behavior may be conditional based on the cell's own status (chemical concentrations and size) or on its neighborhood (the type of cells surrounding it, the contact lengths with each of them, their areas, the directions to these cells, and their chemical concentrations). The language is explored by modeling cellular sorting *in vitro*, and aggregation in the *Dictyostelium discoidea* (cellular slime mold).

631 TOWARDS SECOND-ORDER METHODS FOR STRUCTURED NONSMOOTH OPTIMIZATION
M. Overton, X. Ye, April 1993

Structured nonsmooth optimization objectives often arise in a composite form $f = h \circ a$, where h is convex (but not necessarily polyhedral) and a is smooth. We consider the case where the structure of the nonsmooth convex function h is known. Specifically, we assume that, for any given point in the domain of h, a parameterization of a manifold Ω, on which h reduces locally to a smooth function, is given. We discuss two affine spaces: the tangent space to the manifold Ω at a point, and the affine hull of the subdifferential of h at the same point, and explain that these are typically orthogonal complements. We indicate how the construction of locally second-order methods is possible, even when h is nonpolyhedral, provided the appropriate Lagrangian, modeling the structure, is used. We illustrate our ideas with two important convex functions: the ordinary max function, and the max eigenvalue function for symmetric matrices, and we solicit other interesting examples with genuinely different structure from the community.

632 HIGHLY EFFICIENT DICTIONARY MATCHING IN PARALLEL
S. Muthukrishnan, K. Palem, April 1993

We present highly efficient parallel algorithms for several well-studied dictionary matching problems. Our algorithms are faster and more efficient in terms of their *parallel work*, compared to previously known results.

- For *static dictionary matching*, we present an algorithm that preprocesses the dictionary and matches the text in $O(\log m)$ parallel time and $O(M + n \log m)$ work, given any dictionary of size M whose longest pattern is m characters long, and a text of size n. We have further improved this algorithm to solve static dictionary matching with only $O((M + n)\sqrt{\log m})$ work, if the characters are drawn from an alphabet of constant size. A distinguishing feature of these algorithms and the one stated below for matching in higher dimensions, is that in contrast with previous work, the running times, and work overheads when applicable, are dependent only on the length of the longest pattern m.

- We present a parallel algorithm for *d-dimensional dictionary matching* that runs in $O(\log m)$ time and matches the text in $O(M + n \log m)$ work for any fixed d.

- We present a new and more efficient parallel algorithm for *dynamic dictionary matching. Insertions* into and *deletions* from the dictionary, as well as *matching* the text can be done with *optimal speedup* in $O(\lambda \log M)$ work and $O(\log M)$ time. Here, λ denotes the length of the string to be inserted, deleted or matched into a dictionary of size M.

All of the above algorithms are designed by applying the *shrink-and-spawn* technique that we introduce in this paper. We also show that this technique leads to parallel algorithms that only do optimal (linear) work, for *multi-dimensional pattern matching* and related problems in [KLP89], [Rabin93]. Our algorithms are deterministic, as those in [KLP89], but however, are much simpler and preserve the efficiency as well as the speed of those presented there.

633 SPACE VARIANT IMAGE PROCESSING
R. S. Wallace, P.-W. Ong, B. B. Bederson, E. L. Schwartz, April 1993

This paper describes a graph-based approach to image processing, intended for use with images obtained from sensors having space variant sampling grids. The connectivity graph (CG) is presented as a fundamental framework for posing image operations in any kind of space variant sensor. Partially motivated by the observation that human vision is strongly space variant, a number of research groups have been experimenting with space variant sensors. Such systems cover wide solid angles yet maintain high acuity in their central regions. Implementation of space variant systems pose at least two outstanding problems. First, such a system must be active, in order to utilize its high acuity region; second, there are significant image processing problems introduced by the non-uniform pixel size, shape and connectivity. Familiar image processing operations such as connected components, convolution, template matching, and even image translation, take on new and different forms when defined on space variant images. The present paper provides a general method for space variant image processing, based on a connectivity graph which represents the neighbor-relations in an arbitrarily structured sensor. We illustrate this approach with the following applications: (1) Connected components is reduced to its graph theoretic counterpart. We illustrate this on a logmap sensor, which possesses a difficult topology due to the branch cut associated with the complex logarithm function. (2) We show how to write local image operators in the connectivity graph that are independent of the sensor geometry. (3) We relate the connectivity graph to pyramids over irregular tessalations, and implement a local binarization operator in a 2-level pyramid. (4) Finally, we expand the connectivity graph into a structure we call a transformation graph, which represents the effects of geometric transformations in space variant image sensors. Using the transformation graph, we define an efficient algorithm for matching in the logmap images and solve the template matching problem for space variant images. Because of the very small number of pixels typical of logarithmic structured space variant arrays, the connectivity graph approach to image processing is suitable for real-time implementation, and provides a generic solution to a wide range of

image processing applications with space variant sensors.

634 MINIATURE DIRECT DRIVE ROTARY ACTUATORS
R. S. Wallace, April 1993

This paper reports the development of direct drive DC motor actuators for miniature robots. The motors are based on Nd-Fe-B rare earth permanent magnets and controlled by low cost microcontrollers. The motors have low friction, small size, high speed, low construction cost, no gear backlash, operate safely without limit switches, have limited self-braking, and generate moderate torque. Significantly, one motor can generate enough torque to lift a second motor of about the same size against the force of gravity, at a distance approximately equal to the size of the motor, without resorting to the use of a counterweight. We demonstrated the feasibility of using these actuators to make a two-link robot finger or leg.

635 A LANGUAGE FOR SEMANTIC ANALYSIS
J. Cai, May 1993

Semantic analysis is important for compilers. In the APTS program transformation system, semantics is specified by rules in the language RSL. The semantic rules are interpreted by APTS to generate the semantic information of the program, which is then used by the rewriting engine for program translation. This approach has proved to be convenient and powerful in our construction of a SETL-to-C compiler. In this paper, we discuss the features, applications, implementation strategy, and performance of RSL.

636 VOICE-BANDWIDTH VISUAL COMMUNICATION THROUGH LOGMAPS: THE TELECORTEX
R. Wallace, B. Bederson, E. Schwartz, May 1993

We present a robotic video telephone application of the Cortex-1 miniaturized space-variant active vision system. The embedded processor architecture of Cortex-1 enables it to implement a variety of functions not found in conventional video telephones, for example the camera tracks moving users with its pantilt mechanism. We also report an analog channel coding scheme to transmit logmap video images through band-limited analog channels such as the public switched telephone network (PSTN). The transmitter divides the voice frequency band into 768 channels, and modulates two values in quadrature on each channel. Some channels are reserved for special calibration signals enabling the receiver to recover both the phase and magnitude of the transmitted signal. The remaining channels carry pixel intensities. We synthesize the signal in the frequency domain and run the FFT algorithm to implement a fast conversion to a real, time-domain signal. A phase-lock loop keeps the receiver frame-synchronized with the transmitter. We constructed an experimental video telephone that sends 1376 pixel logmap images at 3.9 frames per second through the PSTN. Using the analog channel coding scheme, we achieve an effective data transfer rate in excess of 40000 bits per second.

637 KNOWLEDGE PRECONDITIONS FOR PLANS
E. Davis, May 1993

For an agent to be able to rely on a plan, he must know both that he is physically capable of carrying out the physical actions involved, and that he knows enough to carry out the plan. In this talk, we advance and discuss new definitions of "knowing enough to carry out a plan", for the case of a single agent carrying out a sequence of primitive actions one at a time. We consider both determinate and indeterminate plans.

Definition: A plan P (determinate or indeterminate) is *executable* for agent A at time T if and only if,

a. P terminates when executed starting in T; and

b. After any beginning of the execution of P starting in T,

b.i A will know whether P has successfully finished;

b.ii A will know of every action whether or not it is a next step of P; and

b.iii All the next steps of P are feasible.

Definition: An indeterminate plan P is *epistemically feasible as a task* for agent A at time T if there exists a plan P' such that A knows in T that

a. P' is executable for A at T; and

b. Any execution of P' starting in T will constitute an execution of P starting in T.

We show how these definition can be expressed in a formal logic, using a situation calculus model of time and a possible worlds model of knowledge. The definitions strictly subsume previous theories for the single-agent case without concurrent actions.

We illustrate the power of the definition by showing that it supports results of the following kinds:

- Positive verification: Showing that a plan is feasible.
- Negative verification: Showing that a plan is infeasible.
- Monotonicity: The more an agent knows, the more plans are executable.
- Reduction for omniscient agent: For an omniscient agent, a plan is executable if and only if it is necessarily physically feasible, and a plan is epistemically feasible as task if and only if it is possibly physically feasible.
- Simple recursive rules that are sufficient conditions for the feasibility of a plan described as a sequence or a conditional combination of subplans.

638 SCHWARZ ANALYSIS OF ITERATIVE SUBSTRUCTURING ALGORITHMS FOR ELLIPTIC PROBLEMS IN THREE DIMENSIONS

M. Dryja, B. Smith, O. Widlund, May 1993

Domain decomposition methods provide powerful preconditioners for the iterative solution of the large systems of algebraic equations that arise in finite element or finite difference approximations of partial differential equations. The preconditioners are constructed from exact or approximate solvers for the same partial differential equation restricted to a set of subregions into which the given region has been divided. In addition, the preconditioner is often augmented by a coarse, second level approximation, that provides additional, global exchange of information, and which can enhance the rate of convergence considerably. The iterative substructuring methods, based on decompositions of the region into non-overlapping subregions, form one of the main families of such

Many domain decomposition algorithms can conveniently be described and analyzed as Schwarz methods. These algorithms are fully defined in terms of a set of subspaces and auxiliary bilinear forms. A general theoretical framework has previously been developed and, in this paper, these techniques are used in an analysis of iterative substructuring methods for elliptic problems in three dimensions. A special emphasis is placed on the difficult problem of designing good coarse models and obtaining robust methods for which the rate of convergence is insensitive to large variations in the coefficients of the differential equation.

Domain decomposition algorithms can conveniently be built from modules, which represent local and global components of the preconditioner. In this paper, a number of such possibilities is explored and it is demonstrated how a great variety of fast algorithms can be designed and analyzed.

639 COMPETITIVE ALGORITHMS AND LOWER BOUNDS FOR ON-LINE SCHEDULING OF MULTIPROCESSOR REAL-TIME SYSTEMS

G. Koren, D. Shasha, June 1993

We study competitive on-line scheduling in multi-processor real-time environments. In our model, every task has a deadline and a value that it obtains only if it completes by its deadline. A task can be assigned to any processor, all of which are equally powerful. The problem is to design an on-line scheduling algorithm (i.e., the scheduler has no knowledge of a task until it is released) with worst case guarantees as to the total value obtained by the system.

We study systems with two or more processors and with uniform or non-uniform value density. We present an inherent limit on the best competitive guarantee that any on-line parallel real-time scheduler can give. Then we present a competitive algorithm that achieves a worst case guarantee which is only within a factor of 2 from the best possible guarantee in many cases. These are the most general results yet known for parallel overloaded real-time scheduling.

640 A PROBABILISTIC APPROACH TO GEOMETRIC HASHING USING LINE FEATURES

F. Tsai, June 1993

Most current object recognition algorithms assume reliable image segmentation, which in practice is often not available. We examine the combination of the Hough Transform with a variation of Geometric Hashing as a technique for model-based object recognition in seriously degraded single intensity images. Prior work on the performance analysis of geometric hashing has focused on point features and shown its noise sensitivity. This paper uses line features to compute recognition invariants

in a potentially more robust way. We investigate the statistical behavior of these line features analytically. Various viewing transformations, which 2-D (or flat 3-D) objects undergo during image formation, are considered. For the case of affine transformations, which are often suitable substitutes for more general perspective transformations, we show experimentally that the technique is noise resistant and can be used in highly occluded environments.

641 CHARACTERIZATION OF SELF-SIMILAR MULTIFRACTALS WITH WAVELET MAXIMA
W.-L. Hwang, S. Mallat, August 1993

Self-similar multifractals have a wavelet transform whose maxima define self-similar curves in the scale-space plane. We introduce an algorithm that recovers the affine self-similarity parameters through a voting procedure in the corresponding parameter space. The voting approach is robust with respect to renormalization noises and can recover the value of parameters having random fluctuations. We describe numerical applications to Cantor measures, dyadique multifractals and to the study of Diffusion Limited Aggregates.

642 CONTROL OF A DEXTEROUS ROBOT HAND: THEORY, IMPLEMENTATION, AND EXPERIMENT
N. Silver, August 1993

Advanced robotic systems, such as multi-fingered hands are becoming more complex, and, as yet, many of the basic questions involved remain unanswered. What control law should we use? What constitutes a good control law? How should we describe motions? What constitutes a broad, yet efficient description of motions for a grasped object?

In addition to the complexity, robotic systems frequently undergo upgrades, and it is therefore necessary to design the system in an organized manner. This includes such things as hierarchical software, system independent descriptions of motion, system independent control laws.

In this thesis, we address these issues for a less complex system. We have focused our attention on describing motions for objects being grasped by a multi-fingered hand. We present a formulation for motion primitives, which allow object manipulation and require limited parameter specification. We have attempted to find a control law which will perform under adverse conditions. We have built the system and tested it on the NYU Four Finger Manipulator which is a two dimensional hand. Even for this simplified problem there remains a large degree of complexity and there are as yet no definitive solutions to these problems.

643 ED I: NYU EDUCATIONAL ROBOT DESIGN AND EVALUATION
B. Mishra, M. Antoniotti, August 1993

The primary goal of the NYU educational robot project 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.

The main work-horse of the NYU educational project was chosen to be a multi-functional ED I robot system, consisting of a 4 DOF DD arm and several auxiliary devices. The system was designed to be simple, inexpensive, flexible and safe.

In this report, we describe the history, design, structure and evaluation of this robot system. We also describe several robotics and related course sequence that can use the ED I system effectively. We also provide some example experiments that have been run on ED I successfully.

This report has benefited from the labor, contribution, discussions, advice and criticisms of several people on the ED I project team and the credit for the final product goes to the entire team.

644 NEW TECHNIQUES FOR THE ANALYSIS AND IMPLEMENTATION OF FUNCTIONAL PROGRAMS
T.-R. Chuang, August 1993

Functional programming languages provide programmers with clean semantics and side–effect free computation, which make easier the tasks of designing programs and reasoning about them. Efficient implementations of purely functional programs, however, can pose certain challenges. Our purpose in this dissertation is to develop new techniques for the efficient analysis and implementation of functional programs.

Our first goal is to investigate a syntactic approach, contrary to the usual semantic approaches, of finding the least fixed points of higher–order functions over finite domains. The second objective is to develop implementation techniques for aggregate data structures for functional programs such that accesses to aggregates are both efficient and side–effects free.

Finding the least fixed point of a monotonic function over a finite domain is an essential task when analyzing a functional

program in the framework of abstract interpretation. Previous methods for least fixed point finding have primarily used semantic approaches, which often traverse large portions of the semantic domain and may be very inefficient even for simple programs. We propose a syntactic method based on an augmented simply typed lambda calculus. It is shown that, for finite domains, the syntactic method is both sound and complete with respect to the semantics. Moreover, we demonstrate that the proposed syntactic method can be quite effective in cases where the usual semantic method is very inefficient.

Efficient implementations of aggregate data structures for functional programs has been an active research topic. The problem arises because once an aggregate is updated, both the old version and newly updated copy must be preserved to maintain the side–effect free semantics of functional languages. We modify the shallow binding scheme of Baker to implement functional arrays for efficient incremental updates and voluminous reads. The scheme, however, uses side–effects and cannot be implemented in purely functional languages themselves. We then investigate the possibility of implementing efficient aggregates without using side–effects, and show that real–time deques can be implemented in a purely functional way. We describe several interesting applications of this technique.

645 NORMS OF FUNCTIONS OF MATRICES

A. Greenbaum, August 1993

Let A be an n by n matrix and let Λ be a diagonal matrix of eigenvalues of A. An expression is given for the smallest number γ such that $\|p(A)\|_F \leq \gamma \cdot \|p(\Lambda)\|_F$, for all polynomials p, where $\|\cdot\|_F$ denotes the Frobenius norm. It is shown that if A has distinct eigenvalues, then the well-known bound – $\gamma = \min_V \kappa_F(V)$, where κ_F denotes the Frobenius norm condition number and the minimum is taken over all eigenvector matrices V of A – is within a factor of n of the optimal constant γ. Classes of matrices A and polynomials p for which the sharp upper bound is also a reasonable estimate of $\|p(A)\|_F$ are described.

646 BIDIRECTIONAL EDGES PROBLEM: PART I A SIMPLE ALGORITHM

B. Mishra, September 1993

The "bidirectional edges problem" is to find an edge-labelling of an undirected network, $G = (V, E)$, with a source and a sink, such that an edge $[u, v]$ in E is labelled $< u, v >$ or $< v, u >$ (or both) depending on the existence of a (simple) path from the source to sink that visits the vertices u and v, in the order u, v or v, u, respectively. We provide several algorithms for this problem in the current paper and the accompanying sequel. In this paper, we show the relation between this problem and the classical two-vertex-disjoint-paths problem and then devise a simple algorithm with a time complexity of $O(|E||V|^2)$. In the sequel, we improve the time complexity to $O(|E||V|)$. The main technique exploits a clever partition of the graph into a set of paths and bridges which are then analyzed recursively.

Bidirectional edges problem arises naturally in the context of the simulation of an MOS transistor network, in which a transistor may operate as a unilateral or a bilateral device, depending on the voltages at its source and drain nodes. For efficient simulation, it is required to detect the set of transistors that may operate as bilateral devices. Also, sometimes it is intended to propagate information in one direction only, and propagation in the wrong direction (resulting in a "sneak path" can cause functional error. Our algorithms can be used to detect all the sneak paths.

647 BIDIRECTIONAL EDGES PROBLEM: PART II AN EFFICIENT ALGORITHM

B. Mishra, September 1993

The "bidirectional edges problem" is to find an edge-labelling of an undirected network, G = (V,E), with a source and a sink, such that an edge [u,v] in E is labelled ¡u, v¿ or ¡v, u¿ (or both) depending on the existence of a (simple) path from the source to sink that visits the vertices u and v, in the order u, v or v, u, respectively. In this paper, building upon the machinery developed in a prequel, we devise an efficient algorithm for this problem with a time complexity of O(—E— —V—). The main technique exploits a clever partition of the graph into a set of paths and bridges which are then analyzed recursively.

Bidirectional edges problem arises naturally in the context of the simulation of an MOS transistor network, in which a transistor may operate as a unilateral or a bilateral device, depending on the voltages at its source and drain nodes. Here, it is required to detect the set of transistors that may operate as bilateral devices.

648 ITERATIVE SUBSTRUCTURING METHODS FOR SPECTRAL ELEMENTS IN THREE DIMENSIONS

L Pavarino, Olof B. Widlund, September 1993

Iterative substructuring methods form an important family of domain decomposition algorithms for elliptic finite element problems. The p-version finite element method based on continuous, piecewise Q_p functions is considered for second order elliptic problems in three dimensions; this special method can also be viewed as a spectral element method. An iterative

method is designed for which number of the relevant operator grows only in proportion to $(1 + \log p)^2$.

649 ITERATIVE SOLUTION OF ELLIPTIC FINITE ELEMENT PROBLEMS ON PARTIALLY REFINED MESHES AND THE EFFECT OF USING INEXACT SOLVERS
H. Cheng, October 1993

We will consider some solution methods for the linear systems of algebraic equations which arise from second-order elliptic finite element problems. We first consider iterative refinement methods. Historically these methods were called FAC and AFAC methods. Optimal bounds of the condition number for the AFAC iterative operator are established by proving a strengthened Cauchy-Schwarz inequality using an interpolation theorem for Hilbert scales. This work completes earlier work by Dryja and Widlund. We also apply an extension theorem for finite element functions to get a weaker bound under a more general assumption. The optimality of the FAC method when spectrally equivalent inexact solvers are used is also proved by using similar techniques.

We next consider multilevel additive Schwarz methods with partial refinement. These algorithms are generalizations of the multilevel additive Schwarz methods developed by Dryja and Widlund and many others. We use quasi-interpolants to prove that this class of methods has an optimal condition number, which is independent of the selection of refinement subregions, under quite general assumptions.

The third algorithm, which we develop, is a domain decomposition parallel preconditioning method (DD-PPM). The preconditioner is closely related to the construction of an operator providing a uniformly bounded extension from the boundary of the nonoverlapping subdomains to the interior of these subdomains. This method is well suited for the computation of finite element solutions on multiprocessor computer architectures. We also obtain a new condition number estimate for general DD-PPM which is sharper than those previously known by using the framework of additive Schwarz methods.

Finally, we develop an iterative substructuring method which is a variation of Smith's vertex space method by using the similar idea of the previous algorithm. Some numerical experiments illustrate the difference between this algorithm and that of Smith.

650 A SURVEY OF COMPUTATIONAL DIFFERENTIAL ALGEBRA
B. Mishra, October 1993

In this note, we explore the computational aspects of several problems in differential algebra with concrete applications in dynamics and motion-planning problems in robotics, automatic synthesis of control schemes for nonlinear systems and simulation of physical systems with fixed degrees of freedom.

Our primary aim is to study, compute and structurally describe the solution of a system of differential equations with coefficients in a field (say, the field of complex numbers,). There seem to have been various approaches in this direction: e.g. ideal theoretic approach of Ritt, Galois theoretic approach of Kolchin and Singer and group theoretic technique of Lie. It is interesting to study their interrelationship and effectivity of various problems they suggest.

In general, these problems are known to be uncomputable; thus, we need to understand under what situations these problems become feasible. As related computer science questions, we also need to study the complexity of these problems, underlying data-structures, effects of the representation (e.g. sparsity).

Of related interest are some closely-related problems such as symbolic integration problem, solving difference equations, integro-differential equations and differential equations with algebraic constraints.

651 Miniature Direct Drive Rotary Actuators II: Eye, Finger and Leg
R. Wallace, November 1993

We have developed miniature direct drive DC motor actuators for robotics. These actuators have low friction, small size, high speed, low construction cost, no gear backlash, operate safely without the use of limit switches and generate moderate torque at a high torque to weight ratio. Our initial experiments indicated the feasibility of constructing a variety of new high speed low cost actuators, for applications in camera pointing, robot hands, and robot legs. In this work we study some prototype devices in each of these categories.

652 Feedback Control of Miniature Direct Drive Devices
D. Max, R. Wallace

We discuss dynamics and control of miniature direct drive actuators, specifically for the two axis spherical pointing motor and for the one axis direct drive finger joint. These actuators can move both accurately and at high speed, however the capabilities of these devices cannot be fully exploited using open loop control techniques. We derive an ideal PID feedback control scheme. Our initial experiments indicate that PID feedback control for the SPM and finger joint is highly feasible

and a significant improvement over open loop methods.

653 NYU Educational Robotics Project:A Pedagogic Overview
B. Mishra, M. Antoniotti, F. Hansen and R. Wallace, November 1993

The primary goal of the NYU educational robotics project (NYU-ED) 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.

The earlier approaches to robotics laboratory education have been to use either industrial robot arms or commercially available low-power arms. In each case, there have been considerable problems in the lack of an ideal interface, in not providing enough flexibility to add other devices, or in the absence of adequate safety. Also, the underlying dynamical model is usually so complicated that performing any control experiment has been beyond the scope of all but advanced graduate students.

In this report, we describe our approach to deal with these problems in constructing a modern robotics laboratory. The main work-horse of the NYU educational project was chosen to be a multi-functional ED I robot system, consisting of a 4 DOF DD arm and several auxiliary devices. The system was designed to be simple, inexpensive, flexible and safe.

We also describe our experience with some advanced laboratory experiments using miniature direct drive robots that have proven to be ideal for teaching as they are reconfigurable, safe and easy to program.

654 Multilevel Schwarz Methods with Partial Refinement
H. Cheng

We consider multilevel additive Schwarz methods with partial refinement. These algorithms are generalizations of the multilevel additive Schwarz methods developed by Dryja and Widlund and many others. We will give two different proofs by using quasi-interpolants under two different assumptions on selected refinement subregions to show that this class of methods has an optimal condition number. The first proof is based purely on the localization property of quasi-interpolants. However, the second proof use some results on iterative refinement methods. As a by-product, the multiplicative versions which corresponds to the FAC algorithms with inexact solvers consisting of one Gauss-Seidel or damped Jacobi iteration have optimal rates of convergence. Finally, some numerical results are presented for these methods.

655 Pscheme: Extending Continuations to Express Control and Synchronization in a Parallel LISP
C. Yao, B. Goldberg, March, 1994

In this paper, we describe Pscheme, a parallel dialect of Scheme. The primary construct for specifying parallelism, synchronization, and communication is a natural extension of first-class continuations which we call a port. We describe the behavior of ports, along with the other parallel constructs of Pscheme. Because the user has precise control over the parallel computation, the Pscheme constructs can be used to build higher-level parallel programming abstractions, such as futures, semaphores, and Ada-style rendezvous. We provide the Pscheme code for these abstractions and discuss the current implementation of Pscheme on a shared-memory multiprocessor.

657 Adaptive Time-Frequency Approximations with Matching Pursuits
G. Davis, S. Mallat, Z. Zhang

Computing the optimal expansion of a signal in a redundant dictionary of waveforms is an NP-complete problem. We introduce a greedy algorithm called a matching pursuit which computes a sub-optimal expansion. The dictionary waveforms which best match a signal's structures are chosen iteratively. An orthogonalized version of the matching pursuit is also developed. Matching pursuits are general procedures for computing adaptive signal representations. With a dictionary of Gabor functions, a matching pursuit defines an adaptive time-frequency transform. We derive a signal energy distribution in the time-frequency plane which does not contain interference terms, unlike the Wigner and Cohen class distributions. A matching pursuit is a chaotic map whose asymptotic properties are studied. We describe an algorithm which isolates the coherent structures of a signal and show an application to pattern extraction from noisy signals.

658 Stability Theory for Dissipatively Perturbed Hamiltonian Systems
J. Maddocks, M. Overton

It is shown that for an appropriate class of dissipatively perturbed Hamiltonian systems, the number of unstable modes of the dynamics linearized at a nondegenerate equilibrium is determined solely by the index of the equilibrium regarded as a critical point of the Hamiltonian. In addition, the movement of the associated eigenvalues in the limit of vanishing dissipation

is analyzed.

659 **A New Primal-Dual Interior-Dual Interior-Point Method for Semidefinite Programming**
F. Alizadeh, J.P. Haeberly, M. Overton

Semidefinite programming (SDP) is a convex optimization problem in the space of symmetric matrices. Primal-dual interior-point methods for SDP are discussed. These generate primal and dual matrices X and Z which commute only in the limit. A new method is proposed which iterates in the space of commuting matrices.

660 **Optimizing Eigenvalues of Symmetric Definite Pencils**
J.P. Haeberly, M. Overton

We consider the following quasiconvex optimization problem: minimize the largest eigenvalue of a symmetric definite matrix pencil depending on parameters. A new form of optimality conditions is given, emphasizing a complementarity condition on primal and dual matrices. Newton's method is then applied to these conditions to give a new quadratically convergent interior-point method which works well in practice. The algorithm is closely related to primal-dual interior-point methods for semidefinite programming.

661 **A Polylogarithmic Bound for an Iterative Substructuring Method for Special Elements in Three Dimensions**
L. Pavarino, O. Widlund

Iterative substructuring methods form an important family of domain decomposition algorithms for elliptic finite element problems. A p-version finite element method based on continuous, piecewise Q_p functions is considered for second order elliptic problems in three dimensions; this special method can also be viewed as a conforming spectral element method. An iterative method is designed for which the condition number of the relevant operator grows only in proportion to $(1+\log p)^2$. This bound is independent of jumps in the coefficient of the elliptic problem across the interfaces between the subregions. Numerical results are also reported which support the theory.

662 **Multilevel Schwarz Methods for Elliptic Problems with Discontinous Coefficients in Three Dimensions**
M. Dryja, M. Sarkis, O. Widlund

Multilevel Schwarz methods are developed for a conforming finite element approximation of second order elliptic problems. We focus on problems in three dimensions with possibly large jumps in the coefficients across the interface separating the subregions. We establish a condition number estimate for the iterative operator, which is independent of the coefficients, and grows at most as the square of the number of levels. We also characterize a class of distributions of the coefficients, called quasi-monotone, for which the weighted L^2-projection is stable and for which we can use the standard piecewise linear functions as a coarse space. In this case, we obtain optimal methods, i.e. bounds which are independent of the number of levels and subregions. We also design and analyze multilevel methods with new coarse spaces given by simple explicit formulas. We consider nonuniform meshes and conclude by an analysis of multilevel iterative substructuring methods.

663 **Iterative Substructuring Methods for Spectral Elements: Problems in Three Dimensions Based on Numerical Quadrature**
L. Pavarino, O. Widlund

Iterative substructuring methods form an important family of domain decomposition algorithms for elliptic finite element problems. Two preconditioners for p-version finite element methods based on continuous, piecewise Q_p functions are considered for second order elliptic problems in three dimensions; these special methods can also be viewed as spectral element methods. The first iterative method is designed for the Galerkin formulation of the problem. The second applies to linear systems for a discrete model derived by using Gauss-Lobatto-Legendre quadrature. For both methods, it is established that the condition number of the relevant operator grows only in proportion to $(1+\log p)^2$. These bounds are independent of the number of elements, into which the given region has been divided, their diameters, as well as the jumps in the coefficients of the elliptic equation between elements. Results of numerical computations are also given, which provide upper bounds on the condition numbers as functions of p and which confirms the correctness of our theory.

664 **Conceptual and Pragmatic Tools for Design and Control of Manufacturing Systems**
M. Antoniotti

In this survey we explore two conceptual tools that have been used in recent (and not so recent) years to build "models of and to "control" various Manufacturing Systems (MS) and processes. Our aim is to investigate some "computational frameworks" that could be interpreted and used to (partially) perform the tasks of

- Simulation of an (arbitrary) MS; - Preventive Verification of some of its properties; - Evaluation of some of its characteristic parameters; - Control of its behavior.

These tasks require well defined computational models, that should be easily equipped with an appropriate "semantics", capable of expressing several properties of MS's. Petri Nets and RW-Systems satisfy these requirements.