

DEC 9 1999

SANDIA REPORT

SAND99-2958

Unlimited Release

Printed December 1999

LDRD Final Report: Global Optimization for Engineering Science Problems

William E. Hart

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

RECEIVED
DEC 30 1999
OSTI

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (703) 605-6000
Web site: <http://www.ntis.gov/ordering.htm>

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01



DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

SAND99-2958
Unlimited Release
Printed December 1999

LDRD Final Report: Global Optimization for Engineering Science Problems

William E. Hart[†]
Optimization/Uncertainty Estimation Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0111

Abstract

For a wide variety of scientific and engineering problems the desired solution corresponds to an optimal set of objective function parameters, where the objective function measures a solution's quality. The main goal of the LDRD "Global Optimization for Engineering Science Problems" was the development of new robust and efficient optimization algorithms that can be used to find globally optimal solutions to complex optimization problems. This SAND report summarizes the technical accomplishments of this LDRD, discusses lessons learned and describes open research issues.

[†]Dept. 9211; email wehart@sandia.gov.

Acknowledgements

The work of this LDRD involved many technical staff at Sandia-Albuquerque as well as academic collaborators. The author would like to recognize Boyd Schimel (9211) and Laura Swiler (6411) for their contributions in constrained optimization; Jonathan Eckstein (RUTGERS) and Cindy Phillips (9211) for their work on parallel branch and bound; Dean Jones (6235), Craig Lawton (6235) and Ed Kjeldgaard (6235) for their help with the Pantex production planning application; Chris Rosin (TSRI), Rik Belew (UCSD) and Garrett Morris (TSRI) for their work on the AutoDock docking application; Keith Hunter (9211) for his work with evolutionary pattern search and tabu search; John Wagner (15334), Perry Gray (15334) and Mike Trahan (15334) for their work with statistical mechanical evolutionary algorithms, and the clustering, neuroimaging, remote sensing and particle swarms applications. The details of these contributions are shown in the referenced SAND reports, conference papers and journal papers.

Preface

This document is a final report for the "Global Optimization for Engineering Science Problems" LDRD project, which was funded for fiscal years 1997 through 1999. It summarizes the main discoveries and research contributions of this LDRD activity and includes citations to the conference and journal papers that arose from the activity. Many of these papers can be retrieved online at <http://www.cs.sandia.gov/~wehart/papers.html>.

The development of the SGOPT global optimization library can be attributed in part to support from this LDRD. However, SGOPT is documented in a separate report.

Contents

Acknowledgements	iv
Preface	v
1 Introduction	1
2 Global Optimization Methods	1
3 Work Accomplished	3
3.1 Global Optimization Survey	3
3.2 Algorithmic Research	3
3.2.1 Hybrid EAs	4
3.2.2 Evolutionary Pattern Search	4
3.2.3 Parallel Branch-and-Bound	4
3.2.4 Statistical Mechanical EAs	5
3.2.5 Constrained Optimization	5
3.3 Software Development	6
3.3.1 SGOPT Optimization Abstractions	6
3.3.2 Optimization Methods	7
3.4 Applications	8
3.4.1 Canonical Test Problems	8
3.4.2 Remote Sensing	8
3.4.3 Production Planning	9
3.4.4 Clustering	9
3.4.5 Particle Swarms	9
3.4.6 Neuroimaging	10
3.4.7 Drug Docking	10
4 Presentations and Publications	11
5 Impact and Future Directions	13

1 Introduction

With decreasing budgets, more complex equipment, costly testing procedures, and national policy restrictions on nuclear testing, Sandia is relying more and more heavily on models in design, evaluation, and manufacturing of weapons systems. The ability to optimize such models will play a critical role in the maintenance and development of the nuclear stockpile. Optimization methods provide the ability to computationally examine characteristics of a model to solve problems such as: evaluate worst-case scenarios, improve initial designs, identify resource bottlenecks, and efficiently allocate resources. In all of these cases, it is highly desirable to provide globally optimal solutions, which represent the best amongst all possible candidate solutions. Even if globally optimal solutions cannot be found in a reasonable time frame, it is often valuable to apply heuristic global optimization methods that quickly “scope out” a parameter space and then refine near-optimal solutions, because even these methods provide a higher level of confidence than local solutions typically found by classical nonlinear optimization techniques.

Towards this end, a general purpose global optimization library has been developed. Written in C++, the SGOPT (Stochastic Global OPTimization) library provides a framework for defining optimization problems and optimizers. This framework is not specific to global optimization methods, but that has naturally been the focus of our efforts. The object-oriented design of these key abstractions enables SGOPT to provide a flexible and extensible optimization library for current and future problems of interest. The design of these concepts is very similar to and has been influenced by the optimization structures in OPT++ [25] and DAKOTA [7]. The principal differences are that (1) SGOPT has been designed to encompass continuous and discrete optimization parameters, and (2) the abstractions in SGOPT are specific to optimization (unlike the more general abstractions in DAKOTA).

The LDRD research activities have focused on the development, implementation and evaluation of global optimization algorithms. We have developed two new classes of evolutionary algorithms. Additionally, we have developed new hybrid evolutionary algorithms, and we have developed methods for constrained global optimization. Our software development has focused on the extension and integration of global optimization methods into SGOPT. We integrated 12 global optimization methods into SGOPT, four of which we implemented directly. Finally, we have applied these methods to several real-world applications at Sandia: Pantex production planning, remote sensing, clustering, neuroimaging, drug docking, and particle swarms.

2 Global Optimization Methods

A wide variety of scientific and engineering problems can be posed as optimization problems in which the desired solution to the problem corresponds to an optimal set of parameters for an objective function that measures a solution’s quality. Optimization algorithms search for optimal solutions by iteratively generating candidates that are evaluated by an objective function. Conventional nonlinear optimization methods like conjugate gradient and quasi-Newton can be applied to find a locally optimal solution, which is better than all nearby

solutions. However in many cases a globally optimal solution is desired, which is at least as good as all other feasible solutions. For these types of problems, global optimization methods are used.

In order to provide a reasonable confidence that a globally optimal solution is found, global optimization algorithms generate solutions in either a deterministic, systematic fashion or with a randomized algorithm that can eventually generate any feasible solution. These methods are commonly applied to objective functions with the following features: (a) many local optima (or constraints that create local optima), (b) they are inexpensive to moderately expensive (to allow a sufficient number of function evaluations), and (c) the derivatives of the objective function, which are used extensively in conventional nonlinear optimization methods, are frequently unavailable.

The standard global optimization methods include (a) exact methods, like branch and bound, that recursively partition the search space, (b) adaptive stochastic methods, like genetic algorithms and simulated annealing, that adaptively sample the search space, (c) clustering methods that combine uniform random sampling with local minimization, and (d) Bayesian methods that use probabilistic models of the objective function to guide the search. Only exact methods are guaranteed to find global optima in a finite amount of time. The other methods provide weak asymptotic guarantees that a global optimum will be found with high probability, though they are generally quite effective in practice.

The principle challenge in developing practical global optimization techniques is to develop methods that both (1) minimize the total time required to find the globally optimal solution and (2) quickly generate near-optimal solutions. The second criterion is important because it is often difficult to guarantee that a globally optimal solution is found within a reasonable time frame for very complex problems that have many local minima. Consequently, it is important to develop methods that can find the "most optimal" solution possible within a given time frame. In addition, global optimization algorithms are often used in an exploratory manner to learn more about the landscape of the objective function.

Our research has included both an evaluation of existing global optimization techniques and the development of new global optimization methods. When developing new methods, the following algorithmic issues must be considered to develop practical methods that find near optimal solutions quickly while provide a global search overall.

Global vs Local Optimization Many of the most effective global optimization algorithms combine global sampling (to cover the space of feasible solutions) with local optimization (to refine solutions to local optima).

Stopping Rules The design of stopping rules is particularly important for randomized global optimization algorithms, because it has proven difficult to terminate these algorithms after a global optimum is sampled.

Constraints While unconstrained problems have been the traditional focus of global optimization research, constrained global optimization problems are becoming increasingly important because constraints are a common component of many applications.

Use of Function Evaluations The effective use of function evaluations becomes an important factor as the cost of each function evaluation increases. For even moderately

expensive function evaluations, it can be more efficient to spend additional time carefully selecting new solutions for evaluation.

Parallelism Exploiting parallelism is a powerful capability that can enable the solution of much larger problems than are available for serial computers. Exploiting parallelism to perform function evaluations simultaneously is relatively easy to achieve. However, there are significant challenges for the design of global optimization algorithms that parallelize the global search itself.

3 Work Accomplished

Our research efforts can be roughly divided into a survey of related work, algorithmic research, software development and applications. In almost all cases the algorithms that we designed were implemented and evaluated with standard test problems and real-world applications.

3.1 Global Optimization Survey

We prepared a survey of global optimization methods that was published on the web at <http://www.cs.sandia.gov/opt/survey/>. The focus of the survey was on general techniques that are applicable to a wide variety of combinatorial and continuous optimization problems that are likely to arise in applications at Sandia National Laboratories. The survey complements existing web surveys like the global (and local) optimization web page by providing more information about the methods. The description of each optimization method includes (a) an overview of the method, (b) a list of application domains in which it has been successfully applied, (c) notes about the software that is available, (d) references for further information and (e) links to other web resources.

The methods that we surveyed included: branch and bound methods like mixed-integer linear programming, clustering methods, evolutionary algorithms, simulated annealing, statistical methods, tabu search, and various hybrid methods. Software is publicly available for most of the methods that we surveyed, and the links to other web resources provide access to a wide range of optimization-related web pages.

3.2 Algorithmic Research

The principle focus of our algorithmic research is on evolutionary algorithms (EAs). This reflects the fact that these are the methods that the team members know best, as well as the fact that EAs did very well in most of our comparative evaluations. EAs are characterized by (a) the fact that they search with a set of candidate solutions (called a *population*) and (b) they use a *mutation* operator to generate small steps and a *crossover* search method to generate new candidate solutions from two (or more) candidate solutions. For descriptions of different EAs, we refer the reader to Fogel [9, 10], Bäck and Schwefel [2], Bäck, Hoffmeister and Schwefel [1], Goldberg [11] and Davis [4].

3.2.1 Hybrid EAs

We worked with Richard Belew (UCSD) to develop new techniques for hybridizing EAs on continuous domains [20, 26]. Our model for hybridization is to refine a subset of the population of candidate solutions with a local search that is run for a fixed duration. Our research has focused on using direct search methods for local search. We performed comparison tests between different direct search methods to confirm that methods with better theoretical convergence guarantees led to better hybrid EAs [20]. Furthermore, we evaluated strategies for adaptively initializing the step length in the local search to improve the efficiency and robustness of the optimization. Preliminary experiments with standard test problems suggest that these strategies can significantly improve the robustness of these methods, though our work with a drug docking application did not confirm these results [20].

3.2.2 Evolutionary Pattern Search

Evolutionary pattern search algorithms (EPSAs) are a class of self-adapting evolutionary algorithms that modify the step size of the mutation operator in response to the success of previous optimization steps. Previously, we proved a stationary point convergence theory for EPSAs that shows that

$$P\left(\liminf_{k \rightarrow \infty} \|g(x_k^*)\| = 0\right) = 1,$$

where $g(x)$ is the gradient of the objective function at x_k^* , where x_k^* is the best point in the population at iteration k [14, 15, 18]. Our initial analysis applied to EPSAs that were not allowed to increase the mutation step size [15]. We subsequently generalized this analysis to prove a convergence theory for EPSAs that are allowed to increase the step size, as well as proving convergence for bound-constrained optimization [14, 18].

In addition to this analysis, we implemented and evaluated EPSAs [16, 17, 19]. The framework for proving convergence of EPSAs is quite general, so in part this empirical evaluation was directed at identifying appropriate algorithmic options for these methods. However, these evaluations also included comparison with standard methods on a drug docking application, which demonstrated that these methods often converge more quickly to better solutions.

3.2.3 Parallel Branch-and-Bound

Parallel branch-and-bound provides a key capability for exactly solving large-scale combinatorial problems. With Jonathan Eckstein (RUTCOR), we developed a simple thread management tool that can be used to flexibly decompose parallel branching algorithms into threads that are scheduled at different priorities. Task decomposition with threads offers a flexible programming methodology, which is essential for complex, asynchronous algorithms like parallel branch-and-bound. We used this tool to implement a simple parallel branch-and-bound algorithm as well as a parallel genetic algorithm [5]. Because the threads were relatively independent, we were also able to quickly hybridize these algorithms to use the parallel genetic algorithm to find feasible solutions that could be used to prune the search tree explored by branch-and-bound. The development of this parallel branch-and-bound

engine has been continued under the "Parallel Combinatorial Optimization for Scheduling Problems" LDRD.

3.2.4 Statistical Mechanical EAs

We developed a methodology for adapting the step size of EAs on continuous problems using a statistical mechanical model of the global search.

This model uses population statistics to adaptively bias the mutation steps.

One of the major limitations of EAs is that for fixed step size, the tendency of a particular scheme is for the population to initially converge on a solution and then stay with a fixed (or at least very slowly decreasing) dispersion about the solution with a spread characterized by the step size. The common solution to this is to implement some scheme by which the step size is decreased as the EA solution is approached. These methods tend to be *ad hoc*. As an alternate approach we developed a methodology by which the step size of the EA is tied directly to the dispersion of the population in the search space. This technique can be thought of as a hybridization of a simulated annealing method applied to an EA. The scheme by which the step size is adjusted from iteration to iteration is similar to the annealing method of Metropolis algorithm [24] with the additional twist that 'temperatures' are computed from the EA population and hence the system is not constrained to cool. This method was incorporated into our existing EA codes and compared against the code with annealing turned of on a standard suite of test problems [11]. The comparisons were made using both single and multiple temperature versions of the the step size annealing scheme; these preliminary results are encouraging in that when hybrid scheme improved performance of the EA it did so by as much as an order of magnitude. Although improvement was not achieved in all case, we are evaluating algorithmic parameters to tune the performance of this method.

3.2.5 Constrained Optimization

Our research with constrained global optimization methods has focused on EAs, although many of the techniques that we have developed should apply equally well to other heuristic global optimization methods. We surveyed the literature on constrained EAs and were surprised to learn how little has been done in this area. Although several methods have been proposed for solving nonlinearly constrained problems, we noted that virtually nothing has been done to solve bound-constrained and linearly constrained problems (which are arguably easier to solve).

We have begun research on bound-constrained and linearly-constrained EAs. We extended our analysis of EPSAs to include bound constraints [18]. Further, we have begun to evaluate the design of standard self-adaptive methods on bound-constrained problems. The initial effort has focused on defining the theoretical properties of fine tuning mutation operators for bound constrained problems, and determining how these propties can be exploited in practical algorithms to best enable a rapid search capability. The methods developed for bound constrained problems should extend readily to linearly constrained problems.

3.3 Software Development

Global optimization methods were developed within the SGOPT global optimization library. The following sections describe significant software development activities supported by this LDRD.

3.3.1 SGOPT Optimization Abstractions

We extended and implemented several optimization abstractions in SGOPT to facilitate and extend its ability to be applied to new problem domains.

Generic Optimizers First, we adapted the optimization hierarchy to include optimizers for a generic search space. The majority of optimizers initially included in SGOPT were designed to perform optimization over \mathbf{R}^n . However, some optimizers like EAs can be readily adapted to perform optimization over a wide variety of search domains. Two approaches were taken in SGOPT to accommodate the application of optimizers to a new search domain. First, a facility was implemented that allows a user to provide a generic class that defines the basic domain-specific operations that are used by an optimizer. This class also encompasses the definition of a generic point in the search domain that the user wants to use for optimization. Using this “generic-point” class to perform optimization is straightforward since the user does not need to develop a domain-specific optimizer. However, the user’s ability to customize the optimizer to a particular search domain is limited by the set of available domain-specific operations that are allowed for a generic-point class.

Second, whenever possible an object oriented design has been adopted which allows a user to easily define a new C++ class to perform optimization over a novel search domain. Specifically, the optimizers have been designed to enable extensive code reuse. Developing an optimizer for a new search domain simply requires the definition of virtual methods of the base optimizer class that encapsulate the domain-specific operations. Instantiating a new optimizer in this manner is considerably more complex than using the generic-point facility, but this option does allow the user to make algorithmic modifications that tailor the optimizers in SGOPT for a particular search domain.

Smooth and Constrained Problems The architecture for defining an optimization problem in SGOPT was also extended to enable: (1) optimization problems with derivatives, (2) optimization problems with constraints (including constraint derivatives), and (3) both continuous and combinatorial search domains. With these extensions, the current SGOPT framework encompasses the class of problems which can be specified by OPT++ and the DAKOTA toolkit.

As part of this work, a new application interface for SGOPT was developed which is similar to the application interfaces of OPT++ and DAKOTA. Further, the specification of an optimization problem was designed to offer a simple manner in which to specify asynchronous evaluations along with a synchronization point. This enables the independent, parallel execution of function evaluations (and constraint and gradient calculations) in a generic fashion, so long as an optimizer indicates that these evaluations are independent. A

master-slave implementation was developed to illustrate this capability, which can currently be exploited by the evolutionary algorithm and Monte Carlo class hierarchies in SGOPT.

3.3.2 Optimization Methods

One of the main goals of our work was to extend SGOPT to include a wider range of global optimization methods. We have succeeded in developing and integrating 12 new optimization methods into SGOPT. These methods include

- two clustering methods,
- pattern search methods,
- Bayesian optimization,
- evolutionary pattern search,
- iterated Bayesian line search,
- grid search (a super-uniform, deterministic version of random sampling),
- the Nelder-Mead simplex method,
- a quasi-Newton method, and
- tabu search,
- GRASP (a biased multistart search method),
- COBYLA2 (a polyhedral search method that can handle constraints).

Most of these techniques are suited for solving problems over \mathbf{R}^n . This bias reflects the fact that our main applications right now involve problems with this character. Additionally, we have integrated both local and global optimization methods to facilitate the development of hybrid methods using these new local optimizers.

Many of the global optimization methods integrated into SGOPT were taken from publicly available software. However, we developed the GRASP, tabu search, evolutionary pattern search and a clustering method ourselves. Our implementations of GRASP and tabu search include a unique, abstract search engine that provides a generic definition of these search methods. The implementation of evolutionary pattern search meets the structure of the convergence analysis for these methods, so this algorithm is asymptotically convergent. Finally, our implementation of the cluster method provides a more scalable implementation than publically available implementations.

3.4 Applications

In addition to our core algorithmic and software development, we have surveyed and developed global optimization applications at Sandia. Our survey provided a sense of the general problem characteristics, which helped focus our research efforts. The global optimization applications that we identified include applications in engineering design, reliability analysis, remote sensing, sensor placement, transportation, manufacturing, parallel scheduling, robotics, site security and stockpile stewardship. All of the basic algorithmic issues discussed in Section 2 apply to methods for these applications. An additional factor that we noted is the ability to perform optimization with some degree of confidence. Many of the optimization methods we surveyed are often used in a heuristic fashion that provides little confidence in the final solution (though all of the methods have proven successful in various application domains). The following sections describe applications that we have used to evaluate global optimization methods.

3.4.1 Canonical Test Problems

Comparing the relative performance of global optimization methods is notoriously difficult. In particular, a standard set of difficult, interesting test functions have not been developed for global optimization. The standard test sets for optimization include the following. The standard global optimization test functions from Dixon and Szego are generally considered too easy (in part because they are low dimensional problems). The DeJong test functions commonly used with genetic algorithms also are perceived as relatively easy, at least for evolutionary algorithms. The MINPACK-2 test problem collection was developed for testing large-scale optimization methods. Although this collection was not necessarily developed for testing global optimization methods, it does contain a lot of multimodal test functions. The netlib uncon test set is similar, but it contains mostly problems that are unimodal. There are many additional domain-specific test sets, especially for combinatorial problems (e.g. TSPLIB).

Because a standard set of difficult, interesting test functions have not been developed for global optimization, we evaluated our algorithms with a set of difficult test problems that we selected from among these existing test suites.

3.4.2 Remote Sensing

Sandia is technically well positioned to lead the new era in intelligent remote sensing for national security in the areas of counter proliferation, non-proliferation and counter terrorism. A key enabling element in these programs is the intelligent software that can detect and discriminate biologicals, chemicals, and explosives, as well as the design and control of the sensing systems themselves. In this application area, we developed the capability to recognize the 3-D spectral signatures of threat species embedded in the spectra of natural backgrounds using UV fluorescence data. This detection problem is fundamentally different from classical target recognition because the spectral signatures are broad, featureless and overlapping. We applied EAs to optimize neural network parameters and architectures [28, 29]. Specifically, the objective criteria for optimization was the minimization of prediction errors for simulated spectral data.

3.4.3 Production Planning

Sandia National Laboratories has developed and implemented the Pantex Process Model [23, 22, 21], a computerized model to support the planning and scheduling activities at Pantex, a US Department of Energy production plant in Amarillo, Texas. The Evaluation Planning Module (EPM) is one of the core algorithms in the Pantex Process Model. This module provides production planning tools that are used to project how facilities and technicians will be utilized over a given planning horizon (typically a year).

We have developed new formulations for this application that can be solved with standard heuristic global optimization methods. In a preliminary study, an EA was able to find better solutions than the standard v -variable formulation currently implemented in the Pantex Process Model in over 75% of the random trials. On average, the solutions were 4.6% better than the v -variable solution. We have also implemented a simple Tabu Search code for this application. This code currently can only generate solutions that are twice as bad as the v -variable formulation, but we expect this to improve when more sophisticated methods for handling precedence constraints between scheduled tasks are developed.

3.4.4 Clustering

A major difficulty with clustering methods [8] (e.g. K-means or minimal spanning tree) is that the algorithms are biased by the initial conditions. For example in using K-means, one assumes there are clusters present and requires the algorithm to break the data into a specified number of clusters. This approach requires a great deal of human intervention and interpretation of the results. EA methods allow one to relax the constraint of specifying the number of clusters assumed. Our implementation of EA cluster analysis entails EA optimization of the number of clusters generated by a K-means algorithm. One specifies some metric for the EA to minimize with respect to the resultant clusters (e.g the total variance of cluster members about the mean of their respective cluster) and then allows the EA to solve for the number of clusters present. We have implemented a C version of K-means for this project, and we have begun developing an MP version of this code using MPI.

3.4.5 Particle Swarms

We developed two prototypes for optimizing the behavior of swarms of robobugs which attack moving objects: fighter planes and ICBMs. These prototypes utilized particle-in-cell (PIC) simulations to model the swarm behavior, and a genetic algorithm was used to optimize the parameters for a particle's behavior [33]. PIC simulation code provides a useful model of the swarm behavior of a heterogeneous class of particles because PIC codes have proven stable and accurate models of particle dynamics. These methods have well-understood theoretical underpinnings and they can be computed efficiently both serially and in parallel. The models that we examined used standard force equations and equations of motion to allow localized, collective behavior and collision avoidance to be handled. Force laws for friction, drag, and inertia, a pursuer's swarming and target seeking forces, and a target's swarm avoidance force were used.

EAs can also be used to determine optimal design features for robotic control that are consistent with physics-based models. We examined a two-dimensional lattice model that

modeled a swarm of robots following a chemical gradient [31, 3, 32]. The EA was used to optimize physical parameters for different regions of the lattice to enforce the desired behavior of the robotic particles. For example, obstacle avoidance was enforced by the EA by evolving parameters that added repulsive interactions between a wall and the robots while limiting the shear forces of the system.

3.4.6 Neuroimaging

The capabilities of several neural network (NN) architectures were evaluated as candidate tools for use in magneto encephalography (MEG) data analysis. Magneto encephalography is a powerful noninvasive technique that yields specific and important information about the brain. Normal brain activity induces both electric and magnetic fields in the brain. Electric fields are studied using electroencephalography (EEG), and magnetic fields are studied using MEG. In general, clinical applications have focused on identification of specific, usually malfunctioning, parts of the brain. In the future these methods will likely be applied as screening tools for exposure to Weapons of Mass Destruction, and for improving human performance on complex tasks.

The volume of MEG data is very large, and high levels of noise in MEG data make the signal very difficult to interpret. Currently, the analysis of MEG data is a time-intensive nonautomated process. We used an EA to optimize different NN architectures that were applied to simulated MEG data. Two classes of NNs were used: traditional feedforward NNs and recurrent NNs. Recurrent neural networks (RNNs) generalize feedforward NNs by allowing connections between two nodes to flow in both directions and by allowing a connection from a node to itself. By adding a recursion to the NN architecture, we establish a *smarter* NN that can recognize and/or reproduce time sequences. Our initial results with this technique were quite promising.

3.4.7 Drug Docking

Computational methods for molecular docking are valuable tools for structure-based drug discovery. Conformational search methods typically model the ligand in some detail, and they often allow conformational flexibility in either the ligand or receptor site, or both. These methods employ a simulation or optimization method to search through the space of ligand-receptor configurations. We have worked with Art Olson (TSRI), Rik Belew (UCSD) and their collaborators to apply new global optimization strategies to drug docking problems in AutoDock [12, 27]. AutoDock uses a physically detailed model that allows for a fixed receptor site and flexible ligand. It also employs a rapid grid-based method for energy evaluation and precalculates ligand-protein pairwise interaction energies so that they may be used as a look-up table during the conformational search.

Docking problems in AutoDock are typically not very large (10-30 dimensions), and the docking potential is not smooth because of way that precalculated energies are interpolated. We have successfully applied self-adaptive EAs, EPSAs, and hybrid EAs to this application [17, 20, 26, 30]. The hybrid EAs are particularly effective at finding optimal solutions, and our hybrid EAs have found the best solutions for a testbed of six standard docking problems [20].

4 Presentations and Publications

The following summarizes presentations and publications that we prepared which are related to this work. Complete citations for publications are provided in the reference section.

- A survey of global optimization methods. P. Gray, W.E. Hart, L. Painton, C. Phillips, M. Trahan, J. Wagner. <http://www.cs.sandia.gov/SandiaOpt/survey>. Feb, 1997.
- A survey of optimization research at Sandia National Labs. W.E. Hart. *INFORMS*. May, 1997.
- Resource management in a parallel mixed integer programming package. *Intel Supercomputer Users Group*. J. Eckstein, W.E. Hart, C. Phillips. June, 1997.
- A generalized stationary point convergence theory for evolutionary algorithms. W. Hart. *ICGA*. July, 1997.
- A comparison of global and local search methods in drug docking. C. Rosin, S. Halliday, W. Hart, and R. Belew. *ICGA*. July, 1997.
- SGOPT: A C++ library of global optimization methods. *IMSL*. Aug, 1997.
- Status Report of the Joint Sandia-ERDEC Multispectral UV Fluorescence Measurement Analysis Using Computationally Intelligent Algorithms. J. Wagner. *Proc. MASINT Biological Defense Sciences and Technology Symposium 1997*.
- On the dynamics of hybrid genetic algorithms that use local search. *UCSD*. Oct, 1997.
- The Use of Intelligent Algorithms in Multispectral UV Analysis, M. W. Trahan, J. S. Wagner, I. R. Shokair, G. C. Tisone, P. C. Gray, *CALIOPE Program Fourth Interim Technical Review Proc.* 1997.
- On the convergence of evolutionary pattern search algorithms. *AMS Fall Western Section Mtg.*. Nov, 1997.
- Methods for parallel optimization. *ORNL*. Feb, 1998.
- On the application of evolutionary pattern search algorithms. W. Hart. *Evolutionary Programming VII*. March, 1998.
- Several Presentations on Using GA's and Neural nets to analyze Mass Spec data. *ORNL*. May, 1998.
- On neural net analysis of mass spectral data, *LANL* May, 1998.
- On neural net analysis of mass spectral data, *ERDEC Symposium* June, 1998.
- Real time analysis of mass spectral data, *ASMS meeting* June, 1998.
- Computationally Intelligent Algorithms for Design, Control, Analysis and Optimization. J. Wagner. *Proc. AIPA 1998 Symposium*.

- Design and implementation of multilevel parallel optimization on the Intel teraflops. M. Eldred and W. Hart. *Proc Symp Multidisciplinary Analysis and Optimization*. Sept, 1998.
- AutoDock 3.0: Automated Docking using a Lamarkian Genetic Algorithm and an Empirical Binding Free Energy Function. C. Morris et al. *J. Comp. Chem.* 1998.
- Swarms of UAVs and Fighter Aircraft. M. Trahan, J. Wagner, K. Stantz, P. Gray, R. Robinett. *Proc Second Intl Conf on Nonlinear Problems in Aviation and Aerospace*. 1998.
- Adaptive Remote-Sensing Techniques Implementing Swarms of Mobile Agents. K. Stantz, R. Asher, S. Cameron, G. Loubriel, R. Robinett, M. Trahan, J. Wagner. *Proc Fourth Joint Workshop on Standoff Detection for Chemical and Biological Defense*. 1998.
- Adaptive Remote-Sensing Techniques Implementing Swarms of Mobile Agents. S. Cameron, G. Loubriel, R. Robinett, K. Stantz, M. Trahan, J. Wagner. *SPIE Intl Symp Aerospace/Defense Sensing, Simulation, and Controls*. 1999.
- Dynamical Behavior of Multi-Robot Systems Using Lattice Gas Automata. K. Stantz, S. Cameron, R. Robinett, M. Trahan, J. Wagner. *SPIE Intl Symp Aerospace/Defense Sensing, Simulation, and Controls*. 1999.
- Architecting department-scale simulation. D. Greenberg, W. Hart, and C. Phillips. *Algorithms for Parallel Processing*. 1999.
- Improved evolutionary hybrids for flexible ligand docking in AutoDock. W. Hart, C. Rosin, R. Belew, G. Morris. *Intl Conf on Optimization in Comp Chem and Mol Bio*. May, 1999.
- Comparing evolutionary programs and evolutionary pattern search algorithms: A drug docking application. W. Hart. *GECCO*. July, 1999.
- A performance analysis of evolutionary pattern search with generalized mutation steps. W. Hart, K. Hunter. *CEC*. July, 1999.
- Genetic Algorithms as Decision Aids to First Responders. J. S. Wagner. *Weapons of Mass Destruction (WMD) Emergency Responders Symp, Pacific Command Orientation, and Pacific Disaster Center Orientation for Western National Guard States*. Sept, 1999.
- Real-Time Analysis of Individual Airborne Microparticles using Laser Ablation Mass Spectroscopy and Genetically Trained Neural Networks. E. Parker, S. Rosenthal, M. Trahan, J. Wagner, W. Whitten, R. Gieray, P. Reilly, A. Lazar, M. Ramsey. *J Process Analytical Chemistry*. (to appear).

- Detection and Classification of Individual Airborne Microparticles using Laser Ablation Mass Spectroscopy and Multivariate Analysis. E. Parker, S. Rosenthal, M. Trahan, J. Wagner, W. Whitten, R. Gieray, P. Reilly, A. Lazar, M. Ramsey. *Field Analytical Chemistry and Technology*. (to appear).

Additionally, the following is unpublished work supported by this LDRD.

- Optimizing Noisy Functions with an Adaptive GA. M. Anderson and W. Hart. Unpublished Research, 1998.
- Hybrid Simulated Annealing Evolutionary Algorithms. P. Gray, M. Trahan, J. Wagner. (in preparation)
- Dense Mepectral UV Fluorescence Detection of a Dilute Constituent in an Optically Dense Matrix. P. Gray, C. Wehlburg, J. Wagner G. Tisone, H. Chan. *Applied Optics*. (submitted)
- SGOPT: A C++ library of (stochastic) global optimization algorithms. W.E. Hart. (in preparation)
- Multi-Spectral UV fluorescence detection and discrimination of biological spectra. P.C. Gray et. al *Army ERDEC report*. (submitted)
- An analysis of evolutionary algorithms for bound-constrained optimization. B. Schimel, W. Hart. (in preparation)

5 Impact and Future Directions

Research on global optimization has the potential for high impact on a wide variety of existing Sandia programs. Global optimization algorithms have been successfully applied in a wide range of application domains, including a variety of mission-related areas such as nonproliferation, energy, weapons manufacturing, transportation and environment. Our research has had a direct impact on some of these applications, and our involvement in other Sandia programs will certainly lead to the ongoing application of these new optimization methods.

The object-oriented software design of SGOPT facilitates the application of this technology to new problems. Furthermore, this library has been integrated into the DAKOTA toolkit, which makes these new algorithms immediately available to Sandians that perform optimization using a variety of complex simulation tools. Plans are also underway to integrate SGOPT in the PPM as well as tools for vehicle routing for TSD.

Despite the accomplishments of this project, there remain a number of important research directins for our work:

- Integrate nonlinear optimizers into SGOPT: more powerful optimizers from libraries like SGOPT would provide more powerful global optimization hybrids for smooth applications.

- Integrate nonsmooth optimizers into SGOPT: nonsmooth optimizers are not commonly available, but they can be applied to applications that are locally smooth but which contain well-known points of nondifferentiability. For example, the Pantex production planning application is piece-wise linear, with well-known nondifferentiable points. These methods could lead to more powerful global optimization hybrids.
- Constrained global optimization: Our work with constrained EAs represents only a first step towards a complete understanding of how to design constrained global optimization methods. For example, we have not begun to consider whether the methods we have designed for EAs can be applied to other heuristic global optimization methods.
- Nonlinear Branch-and-Bound: For continuous domains, nonlinear branch-and-bound can be used to exactly solve applications for which an algebraic problem definition is available. These methods have proven particularly effective for molecular biology applications like drug docking and protein folding.
- Flexible Optimization Frameworks: One of the software challenges that we have faced in the development of SGOPT is the adaptation of the abstract optimization methods to problem-specific parameter representations. We believe that wrapping SGOPT with a high-level language like Python or Perl could substantially reduce the development time for new applications. In particular, high level languages do not enforce syntactic type checking, which facilitates the use of an abstract, compiled kernel for a general purpose optimization method.

References

- [1] T. Bäck, F. Hoffmeister, and H.-P. Schwefel. A survey of evolution strategies. In R. K. Belew and L. B. Booker, editors, *Proc. of the Fourth Intl. Conf. on Genetic Algorithms*, pages 2–9, San Mateo, CA, 1991. Morgan-Kaufmann.
- [2] T. Bäck and H.-P. Schwefel. An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, 1(1):1–23, 1993.
- [3] S. M. Cameron, G. M. Loubriel, R. D. Robinett III, K. M. Stantz, M. W. Trahan, and J. S. Wagner. Adaptive remote-sensing techniques implementing swarms of mobile agents. In *Proc SPIE 13th Annual Intl Symp Aerospace/Defense Sensing, Simulation, and Controls*, volume 3713, 1999.
- [4] L. Davis, editor. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, 1991.
- [5] J. Eckstein, W. E. Hart, and C. A. Phillips. Resource management in a parallel mixed integer programming package. In *Proc Intel Supercomputer Users Group, 1997*. <http://www.cs.sandia.gov/ISUG97/program.html>.
- [6] M. Eldred and W. E. Hart. Design and implementation of multilevel parallel optimization on the intel teraflops. In *Proc 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, volume 1, pages 44–54, 1998.

- [7] M. S. Eldred, W. E. Hart, W. J. Bohnhoff, V. J. Romero, S. A. Hutchinson, and A. G. Salinger. Utilizing object-oriented design to build advanced optimization strategies with generic implementation. In *Proc 6th Symp on Multidisciplinary Analysis and Optimization*, pages 1568–1582, 1996.
- [8] B. Everitt. *Cluster Analysis*. Oxford University Press, 1993.
- [9] D. B. Fogel. An introduction to simulated evolutionary optimization. *IEEE Trans Neural Networks*, 5(1):3–14, 1994.
- [10] D. B. Fogel. *Evolutionary Computation*. IEEE Press, Piscataway, NJ, 1995.
- [11] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Co., Inc., 1989.
- [12] D. S. Goodsell and A. J. Olson. Automated docking of substrates to proteins by simulated annealing. *Proteins: Structure, Function and Genetics*, 8:195–202, 1990.
- [13] D. S. Greenberg, W. E. Hart, and C. A. Phillips. Architecting department-scale simulation. In *Algorithms for Parallel Processing*, volume 105 of *IMA Volumes in Mathematics and Its Applications*, pages 321–344, 1999.
- [14] W. E. Hart. A generalized stationary point convergence theory for evolutionary algorithms. In T. Baeck, editor, *Proc 7th Intl Conf on Genetic Algorithms*, pages 127–134, San Francisco, CA, 1997. Morgan Kaufmann.
- [15] W. E. Hart. A stationary point convergence theory for evolutionary algorithms. In R. K. Belew and M. D. Vose, editors, *Foundations of Genetic Algorithms 4*, pages 325–342, San Francisco, CA, 1997. Morgan Kaufmann Publishers, Inc.
- [16] W. E. Hart. On the application of evolutionary pattern search algorithms. In *Proc Evolutionary Programming VII*, pages 303–312, New York, 1998. Springer-Berlin.
- [17] W. E. Hart. Comparing evolutionary programs and evolutionary pattern search algorithms: A drug docking application. In *Proc. Genetic and Evolutionary Computation Conf*, pages 855–862, 1999.
- [18] W. E. Hart. A convergence analysis of unconstrained and bound constrained evolutionary pattern search. 1999. (submitted).
- [19] W. E. Hart and K. Hunter. A performance analysis of evolutionary pattern search with generalized mutation steps. In *Proc Conf Evolutionary Computation*, pages 672–679, 1999.
- [20] W. E. Hart, C. R. Rosin, R. K. Belew, and G. M. Morris. Improved evolutionary hybrids for flexible ligand docking in autodock. In *Proc Intl Conf on Optimization in Computational Chemistry and Molecular Biology*, 1999. (to appear).

- [21] D. A. Jones, C. R. Lawton, G. F. List, and M. A. Turnquist. The Pantex model: Formulations of the evaluation planning module. Technical Report SAND99-2095, Sandia National Laboratories, 1998.
- [22] E. A. Kjeldgaard, D. A. Jones, G. F. List, and M. A. Turnquist. Planning and scheduling for agile manufacturers: The Pantex process model. Technical Report SAND98-0030, Sandia National Laboratories, 1998.
- [23] E. A. Kjeldgaard, D. A. Jones, G. F. List, M. A. Turnquist, J. W. Angelo, R. D. Hopson, and J. Hudson. Swords into plowshares: Nuclear weapon dismantlement, evaluation, and maintenance at Pantex. *INTERFACES*, 1999. (to appear).
- [24] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087-1092, 1953.
- [25] J. C. Meza. OPT++: An object-oriented class library for nonlinear optimization. Technical Report SAND94-8225, Sandia National Laboratories, 1994.
- [26] G. M. Morris, D. S. Goodsell, R. S. Halliday, R. Huey, W. E. Hart, R. K. Belew, and A. J. Olson. Automated docking using a Lamarckian genetic algorithm and an empirical binding free energy function. *J Comp Chem*, 19(14):1639-1662, 1998.
- [27] G. M. Morris, D. S. Goodsell, R. Huey, and A. J. Olson. Distributed automated docking of flexible ligands to proteins: Parallel applications of Autodock 2.4. *J. Comp.-Aid. Mol. Des.*, 10:293-304, 1996.
- [28] E. P. Parger, S. E. Rosenthal, M. W. Trahan, J. S. Wagner, W. B. Whitten, R. A. Gieray, P. T. Reilly, A. C. Lazar, and J. M. Ramsey. Real-time analysis of individual airborne microparticles using laser ablation mass spectroscopy and genetically trained neural networks. *J Process Analytical Chemistry*. (to appear).
- [29] E. P. Parker, S. E. Rosenthal, M. W. Trahan, J. S. Wagner, W. B. Whitten, R. A. Gieray, P. T. Reilly, A. C. Lazar, and J. M. Ramsey. Detection and classification of individual airborne microparticles using laser ablation mass spectroscopy and multivariate analysis. *Field Analytical Chemistry and Technology*. (to appear).
- [30] C. D. Rosin, S. Halliday, W. E. Hart, and R. K. Belew. A comparison of global and local search methods in drug docking. In T. Baeck, editor, *Proc 7th Intl Conf on Genetic Algorithms*, pages 221-228, San Francisco, CA, 1997. Morgan Kaufmann.
- [31] K. M. Stantz, R. B. Asher, S. M. Cameron, G. M. Loubriel, R. D. Robinett, M. W. Trahan, and J. S. Wagner. Adaptive remote-sensing techniques implementing swarms of mobile agents. In *Proc Fourth Joint Workshop on Standoff Detection for Chemical and Biological Defense*, 1998.
- [32] K. M. Stantz, R. D. R. Cameron, S. M. Robinett, M. W. Trahan, and J. S. Wagner. Dynamical behavior of multi-robot systems using lattice gas automata. In *Proceedings of SPIE 13th Annual Intl Symp on Aerospace/Defense Sensing, Simulation, and Controls*, volume 3693, 1999.

- [33] M. W. Trahan, J. S. Wagner, K. M. Stantz, P. C. Gray, and R. D. Robinett. Swarms of uavs and fighter aircraft. In *Proc Second Intl Conf Nonlinear Problems in Aviation and Aerospace*, volume 2, pages 745–752, 1998.

UNLIMITED RELEASE
INITIAL DISTRIBUTION:

1 MS 0321 W. CAMP, 9200
1 1111 S. S. Dosanjh, 9221
2 1110 W. E. Hart, 9211
1 1110 D. E. Womble, 9222
1 1188 J. S. Wagner, 15334
1 1188 P. C. Gray, 15334
1 1188 M. W. Trahan, 15334
1 1188 M. W. Trahan, 15334
1 0746 L. P. Swiler, 6411
1 1110 C. A. Philips, 9211
1 1110 K. O. Hunter, 9211
1 0847 B. D. Schimel, 9211

1 MS 9018 Central Technical files, 8940-2
2 0899 Technical Library, 4916
1 0612 Review and Approval Desk, 4912
For DOE/OSTI