

SANDIA REPORT

SAND97-8275 • UC-406

Unlimited Release

Printed August 1997

Parallel Optimization Methods for Agile Manufacturing

J. C. Meza, C. D. Moen, T. D. Plantenga, P. A. Spence, C. H. Tong, B. A. Hendrickson,
R. W. Leland, G. M. Reese

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; distribution is unlimited.



Sandia National Laboratories

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, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents 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
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

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

NTIS price codes
Printed copy: A08
Microfiche copy: A01

Parallel Optimization Methods for Agile Manufacturing

J. C. Meza¹, C. D. Moen, T. D. Plantenga, P. A. Spence, C. H. Tong
Sandia National Laboratories
Livermore, CA 94551-0969

B. A. Hendrickson, R. W. Leland, G. M. Reese
Sandia National Laboratories
Albuquerque, NM 87185

ABSTRACT

The rapid and optimal design of new goods is essential for meeting national objectives in advanced manufacturing. Currently almost all manufacturing procedures involve the determination of some optimal design parameters. This process is iterative in nature and because it is usually done manually it can be expensive and time consuming. This report describes the results of an LDRD, the goal of which was to develop optimization algorithms and software tools that will enable automated design thereby allowing for agile manufacturing. Although the design processes vary across industries, many of the mathematical characteristics of the problems are the same, including large-scale, noisy, and non-differentiable functions with nonlinear constraints. This report describes the development of a common set of optimization tools using object-oriented programming techniques that can be applied to these types of problems. We give examples of several applications that are representative of design problems including an inverse scattering problem, a vibration isolation problem, a system identification problem for the correlation of finite element models with test data and the control of a chemical vapor deposition reactor furnace. Because the function evaluations are computationally expensive, we emphasize algorithms that can be adapted to parallel computers.

¹ Email: meza@ca.sandia.gov

1. Introduction

The automated design problem involves several issues in nonlinear optimization. Unlike standard optimization problems, this class of problems involves functions that are usually described by the solution of a complex simulation. This report describes the results of an LDRD, the major goal of which was the development of a set of portable optimization tools for design optimization problems. In this work we consider four applications that are representative of optimization problems in manufacturing: 1) an inverse scattering problem, 2) a parameter identification problem, 3) a vibration isolation problem and 4) the optimal control of a chemical vapor deposition reactor furnace. In the case of the inverse scattering problem, an integrated circuit is illuminated by a laser that creates an interference pattern in the diffracted light. Given this pattern, geometric properties describing the IC can be inferred which can then be used to determine the quality of the IC. The second application involves the identification of design parameters for a finite element model. The goal of this problem is to determine the design parameters of a finite element model that accurately represents the modal response consisting of both frequencies and mode shape of some given test results. The third application of interest involves developing methods for vibration isolation which entails determining the optimal location and design of passive, semi-active, and active isolation components to minimize machine vibration. In the fourth example, the goal is to determine a set of powers that are applied to a set of heaters that control the temperature inside a chemical vapor deposition reactor furnace. In each of these four cases, the objective functions are expensive to compute, the accuracy of the function values varies, and derivatives are not usually available or the function may be non-differentiable. In addition, there may be various nonlinear constraints on the design parameters including discrete or integer constraints.

Our approach was based on developing a nonlinear optimization framework using object-oriented programming techniques. Our first goal was the development of a C++ class library for nonlinear equations and optimization. Using this approach, we developed a general framework in which new algorithms can be easily developed and tested on various application problems. This approach also has the advantage of generating re-usable code for future applications.

The next major step of this project was the development of efficient optimization algorithms for problems where the function values correspond to the result of a simulation such as the ones mentioned above. The solution from any simulation code always has some inherent discretization or roundoff error due to the numerical techniques used. This goal involved the development of robust optimization techniques for problems with noisy functions. One characteristic of these types of problems that proved advantageous was the ability to control the level of noise, at the price of increasing the cost of the function evaluator.

A third issue involved the availability of derivatives. Since the objective function is implicitly defined as the result of some simulation, neither first nor second derivatives are generally available. We therefore investigated optimization methods that do not require derivative information or that can efficiently compute the necessary information. The goal of this part of the research was to explore the use of both stochastic methods, such as genetic algorithms and simulated annealing and deterministic methods such as direct search methods for design optimization. We also investigated the use of automatic differentiators such as ADIFOR that can automatically generate codes that evaluate derivatives.

The research described in this report touches on various critical technologies. The major contributions of this research are in the areas of high performance computing and the modeling and simulation of complex systems. The development of portable optimization tools provide a new capability to our existing efforts in parallel algorithms.

Another major contribution is in the area of advanced manufacturing. The DOE and Sandia have ongoing work in several technology areas that are essential to design optimization, including structural system identification, structural control, smart materials, and computational mechanics code development. Sandia has performed research in system identification and structural dynamics for many years, including an integrated analytical and experimental effort to develop production system ID software, and has been developing control, actuation and sensing technology for active control of structures for over seven years. Applications of this technology are currently being made to Sandia projects in weapons, submarine technology, space structures and sensors, wind energy, robotics for waste cleanup, rocket and missile structures, and advanced materials and sensors for smart weapon systems. The work on the inverse scattering problem has applications in semiconductor manufacturing where the ability to detect processing errors and to correct them can have a major impact in the yield. This work therefore has important implications for the semiconductor industry and consortiums such as Sematech.

Related areas that would also benefit from this research include Sandia's effort in environment science and technology. This research is closely related to the solution of the seismic inverse problem that has the potential of increasing the accuracy and the efficiency with which geophysicists can develop reservoir characterizations. This has implications for several technological areas that require the accurate determination of geophysical properties. For example, the oil industry spends hundreds of millions of dollars yearly in determining the location of new oil reserves that depends on accurate reservoir descriptions. These geophysical models are also important in the DOE's efforts in waste management and several Sandia programs, including soil remediation, contaminant transport, and waste disposal.

2. Background

Much of the past work in optimization has dealt with small-scale problems where the objective function can easily be manipulated. In recent years, more attention has been focused on large-scale optimization problems (for example the MINPACK2 project), although much work remains to be done. In particular, researchers still make many of the same underlying theoretical assumptions as for small-scale problems. Although these assumptions make the analysis of algorithms easier, they normally do not hold for typical design engineering problems. This has led to a situation where design engineers have had to rely on ad-hoc optimization methods that rely heavily on carefully chosen algorithm parameters. As such, efficiency and reliability are often sacrificed to meet certain design goals.

One of the goals of this LDRD was to develop a set of efficient and robust algorithms that engineers can apply to automated design. This work entailed research in several areas of optimization and parallel algorithms. In optimization the crucial new areas involve the development of methods for large-scale, noisy, optimization problems. Since parameter identification problems are usually computationally expensive we also need to address methods for parallel optimization.

Although some work has been done in parallel optimization, research is still in the early stages. Most of this work has concentrated on the parallelization of the function evaluator or of the linear solution methods for the Newton equations. This work is different in that we concentrate on developing new algorithms that can be parallelized at a higher level such as the parallel direct search methods. Finally, because design problems can have vastly different structures, we needed to develop a general framework that will encompass many applications and has a wide applicability. In contrast to other projects such as the MINPACK2 project (which is written entirely in FORTRAN) our codes use object-oriented class libraries for greater flexibility and re-usability.

3. Accomplishments

This LDRD had a three year plan with eight milestones:

Date	Milestone
3/94	Implementation of an object-oriented class library for nonlinear optimization.
6/94	Development of a test problem data base including two problems derived from the inverse optical scattering problem and the system identification problem.
9/94	Implementation of a set of parallel algorithms for noisy optimization.
3/95	Application and comparison of new methods to the test problem data base.
9/95	Implementation of derivative-free optimization methods.
3/96	Comparison of methods using derivative information versus derivative-free methods on the test problems.
8/96	Implementation of an optimization method for non-differentiable functions.
9/96	Final SAND report.

In the following sections, we report on each one of these milestones.

3.1 FY94 Accomplishments

Milestone 1: Our first major goal was to develop a C++ class library for nonlinear optimization. Using this approach, we have developed a general framework in which new algorithms can be easily developed and tested on various application problems. This approach also has the advantage of generating re-usable code for future applications. We completed the implementation of a set of libraries called OPT++ [1] that can be used on a variety of UNIX platforms including SGI, SUN, HP, IBM, and DEC workstations. In addition, the OPT++ code has been compiled and tested on the Intel Paragon, and PC compatible platforms running Linux.

Milestone 2: The second milestone was to develop a set of problems for testing the new optimization algorithms. This milestone was broken down into 3 activities: 1) discussions with customers, 2) development of an optical scattering problem simulator for the test data base, and 3) development of a finite-element modal analysis simulator for the test data base.

Discussions with Customers. The first phase involved meeting with potential customers of our optimization package, including Mike McGlaun, 1431 and Bob Kee, 8745. The discussions centered on the features needed by optimization users. In addition, several new test problems were suggested which we are investigating for future inclusion into our test data base.

Optical scattering problem. Quick, accurate monitoring of semiconductor patterning processes is central to the efficient manufacture of integrated circuit computer chips. In most conventional methods of monitoring, semiconductor wafers are removed from the production line following key steps and inspected using standard white light microscopy - a technique that is straightforward but which yields only one-dimensional information. Other more rigorous inspection techniques used in the IC industry (for example, scanning electron microscopy) can give 3D information but are very time consuming and destructive. When an integrated circuit is illuminated by a laser, an interference pattern is created in the diffracted light. Geometric properties of the chip surface can, in principle, be inferred from computational analysis of the scattered radiation; this information can then be used to monitor and improve process quality. With this new laser method, a chip could be tested rapidly, automatically, and non-destructively at all stages of processing. This would decrease costs, speed process flow and increase yield, and hence be of substantial economic significance.

The practical difficulty with the laser method has been that it requires the computational solution of an inverse problem. The unknown geometry must be discovered by an iterative process in which a guess is refined by correcting a previous guess and recomputing the scattered radiation pattern. This computation had proven prohibitively expensive. As part of this activity, we completed a massively parallel implementation of the serial forward scattering code on Sandia's 1840 node Intel Paragon supercomputer. The code runs at over 5 GFlops on the full machine and improvements in speed are projected as the code is tuned.

Finite-element Mode Analysis. Integrating a commercial Finite Element (FE) code like NASTRAN into an optimization toolkit presents significant challenges. The task we initially proposed compared results from FE modals with experimental test results, thus data from both types of processes must be evaluated. The optimization strategy we decided upon used the design sensitivity data available from the FE code for the estimates of the first order derivatives of the estimated response. However, because the objective function is not simply related to the output, additional post-processing steps are required. This processing is non-trivial, and significantly impacts the design of the function calls. The need to cast the process into the form of a function call required that interface routines be written to spawn new processes and to interpret the results of the computations based on the results in the data files. A single function evaluation consists of the following seven tasks:

- 1) a call to modify the input data deck according to the specified input parameters,
- 2) a call to the FE analysis package,
- 3) determination of success of the FEA,
- 4) interpretation of the output data files,
- 5) optional computation of Modal Assurance Criteria (MAC) to correlate modes of the analysis and test,
- 6) computation of an objective function, and

- 7) computation of gradients using the partial derivatives of the FE output.

Milestone 3: The third milestone for the first year was to implement a set of parallel algorithms for noisy optimization. We implemented one such algorithm based on the parallel direct search (PDS) method that scales well on massively parallel machines. We also started collaborations with Dr. Virginia Torczon at the Center for Research on Parallel Computing at Rice University to develop new parallel algorithms based on pattern search methods.

3.2 FY95 Accomplishments

In the second year of the LDRD there were two major milestones:

Date	Milestone
3/95	Application and comparison of new methods to the test problem data base.
9/95	Implementation of derivative-free optimization methods.

These milestones were broken down into five major tasks:

- 1) Complete the construction of the test data base
- 2) Complete the implementation of the PDS method on the Intel Paragon
- 3) Run tests of PDS on the problems from the test data base
- 4) Develop new derivative-free optimization methods
- 5) Implement a graphical user interface to the optimization package

Task 1: The second year of the project concentrated on developing interfaces between OPT++ and several analysis codes within Sandia, including NASTRAN, TACO, TWAFER, QTRAN, and CCEMD. All of these interfaces were completed and tested on various applications. One important design problem involved a chemical vapor deposition (CVD) reactor being studied by SEMATECH (see Figure 1). In this application, we used a combination of OPT++ and TACO to compute optimal heater powers to minimize the temperature variation in the wafer stack. The optimal solution generated a temperature profile with less than 0.1 degree variation, providing an order of magnitude improvement over the hand-optimized case (see [2-3]).




Figure 1. Model of the CVD reactor used for computing the optimal heater powers and the associated temperature profile before and after optimization.

In another application, we integrated OPT++ with NASTRAN for a structural dynamics problem. Two different commercial versions of NASTRAN (from either MSC or CSAR) can be called from OPT++. The problem is to update physical parameters, such as the thickness of a plate or a Young's modulus to improve agreement between the natural frequencies of test and analysis models. The semi-analytic derivatives computed within NASTRAN are available in the optimizer, which significantly reduces the number of analyses required. This enhances our ability to provide test validated models for highly accurate models. The same methods can be used for model based design optimization.

We also implemented a C++ interface to a molecular dynamics code to compute low energy conformations for molecule design. Using OPT++ we were able to easily compare various optimization methods for conformational searching and energy minimization (see [4]). These problems are being studied by Sandia to aid in the design of biosensors.

Tasks 2-3: In collaboration with Virginia Torczon and David Serafini at Rice University, we developed a new bound constrained parallel direct search algorithm and applied it to the CVD test problem. This code runs on various platforms including the Intel Paragon, the IBM SP2, and a distributed computer network using either PVM or MPI. The new algorithm was applied to the CVD reactor problem on both an Intel Paragon and an IBM SP2. Tables 1-2 show the times and speedups for this problem on these platforms. Both results show that these methods are parallelizable, with efficiencies around 90% easily achieved.

Nodes	Time (secs)	Speedup	Efficiency
1	81423	1.00	
2	41029	1.98	0.99
4	21334	3.82	0.96
8	10803	7.54	0.94
16	5286	15.40	0.96
32	2708	30.07	0.94
64	1458	55.72	0.87

Table 1. Speedups and efficiencies for CVD reactor problem using the bound constrained parallel direct search method on an Intel Paragon using 512 PDS search directions.

Directions	Time (secs)	Speedup	Efficiency
256	686	94	0.73
512	528	101	0.79
1024	615	107	0.84
2048	891	110	0.86
4096	1154	115	0.90
8192	1719	115	0.90

Table 2. Effect of the number of PDS search directions on the speedups and efficiencies using the bound constrained parallel direct search method on an IBM SP2 using 128 processors.

Task 4: We continued to improve the capabilities of OPT++ by including several new features for engineering design problems. We added 2 new classes of derivative-free algorithms, a genetic algorithm and a simulated annealing algorithm. In addition, we implemented a third method for nonlinear least-squares problems. These methods have already proved useful in several of the applications mentioned under Task 1.

We also worked with Stanford University to develop new optimization capabilities. This work led to the development of a new global optimization algorithm based on stochastic methods for use in cases where there is missing and/or noisy data (see [5,6,7,8]).

Task 5: We implemented a graphical user interface to aid in building interfaces between analysis codes and OPT++. This GUI can automatically generate a C++ main routine, compile it and link it with existing analysis codes so that the analyst need not learn C++ to use the optimization methods within OPT++.

3.3 Additional results

In many of our test problems, derivatives are not available or are expensive to compute. To address this problem, we worked on generating derivatives for analysis codes using an automatic differentiation code, ADIFOR. This capability allowed us to apply more powerful gradient-based

methods to problems where derivatives are not available and finite-difference calculations are too expensive.

3.4 FY96 Accomplishments

In the first two years, we encountered many engineering applications that had nonlinear constraints on the design parameters, for example, in the dynamic optimization and control of CVD reactors and in designing bio-molecules. In addition, the computation of derivatives continued to be a bottleneck in the solution of many of our problems since most simulation codes do not provide derivatives. As a result of these observations, we focused our efforts on 3 major areas:

- 1) the development of a general constrained optimization capability for OPT++
- 2) the use of automatic differentiation techniques to generate derivatives for analysis codes
- 3) the implementation of constrained parallel direct search methods

With respect to our first goal, we designed and implemented a new class hierarchy for bound constrained objects and incorporated them into the OPT++ package. This new feature allows for the general solution of optimization problems with simple upper and lower bounds. We also implemented an active set Quasi-Newton method that can handle bound constrained problems. This method was successfully used on an application that computes the deposition rate of a chemical vapor deposition furnace. As part of solving this application, we also developed a new interface to the nonlinear optimization package NPSOL that allows the user to use general nonlinear constraints. Finally, we applied for and were granted a copyright on the OPT++ software.

We also continued our research into techniques for computing optimal designs for chemical vapor deposition furnaces. We were able to create analysis codes that can compute sensitivities for both the TWAFER and TACO2D code using ADIFOR. These codes were then used to compute optimal heater settings that were better than the ones produced by the original codes. This work was documented and presented at the 6th AIAA/MSO/ISSMO conference on Multidisciplinary Analysis and Optimization (see [9-10]). In addition, we organized a mini-symposium on Optimization in Control and Design Applications at the SIAM Optimization meeting on May 20-22, 1996.

As part of our continuing collaborations with Rice University, Virginia Torczon delivered a bound constrained parallel direct search method. This method can be used on any parallel machine for which an MPI (message passing interface) library exists. In addition to the new software, Torczon was able to prove several new convergence results for the bound constrained version of PDS [11]. In particular, the analysis guarantees global convergence to a first-order constrained stationary point under standard assumptions on the problem. From a practical standpoint, the results in [11] show how to extend the original PDS software to solve problems with bound constraints.

4. Conclusions and Future Work

The results of this LDRD have shown that it is worthwhile to apply modern optimization techniques to many optimal design problems in engineering faced by Sandia. We have developed a library of optimization methods using object-oriented techniques to allow for ease of use and extendibility to other engineering and scientific applications. In addition, we have developed new parallel optimization methods and applied them to manufacturing problems on massively parallel supercomputers as well as networks of workstations.

There are many new research directions that suggest themselves at this point. Two areas that should be mentioned are mixed integer programs and algorithms for cluster computing. The first area includes problems for which some of the parameters take on discrete values while other parameters take on continuous values. This problem has received more attention recently, but there are few algorithms that can be used on practical problems such as the ones addressed in this report. The second area that should be addressed is that of parallel algorithms for clusters of shared memory processors (CLUMPS). This area has recently received attention as new computer centers such as those proposed for ASCI head in this direction. The major issue for these architectures is how to design algorithms that are latency tolerant or location insensitive.

References

1. J.C. Meza. OPT++: An Object-Oriented Class Library for Nonlinear Optimization, *Sandia Technical Report SAND 94-8225*, March 1994.
2. C.D. Moen, P.A. Spence, J.C. Meza. Optimal Heat Transfer Design of Chemical Vapor Deposition Reactors, *Sandia Technical Report SAND95-8223*, April 1995.
3. J.C. Meza, T.D. Plantenga. Optimal Control of a CVD Reactor for Prescribed Temperature Behavior, *Sandia Technical Report SAND95-8224*, April 1995.
4. J.C. Meza, R.S. Judson, T.R. Faulkner, A.M. Treasurywala. A Comparison of a Direct Search Method and a Genetic Algorithm for Conformational Searching, *J. Comp. Chem.*, Vol. 17, No. 9, 1142-1151, 1996. Also available as *Sandia Technical Report SAND95-8225*, April 1995.
5. E.H.S. Ip. Estimation in Second Order Dependency Model for Multivariate Binary Data, *Sandia Technical Report SAND 95-8100*, April 1995.
6. E.H.S. Ip and J. Diebolt. A Stochastic Difference Equation Approach to Inference with Missing Data: Some New Results, *Sandia Technical Report SAND 95-8101*, April 1995.
7. J. Diebolt and E.H.S. Ip. A Stochastic EM Algorithm for Approximating the Maximum Likelihood Estimate, *Sandia Technical Report SAND 95-8102*, April 1995.
8. E.H.S. Ip. A Stochastic EM Estimator in the Presence of Missing Data - Theory and Applications, *Sandia Technical Report SAND 95-8103*, April 1995.
9. A. Carle and M. Fagan. Automatic Differentiation of the TACO2D Finite Element Code Using ADIFOR, *Sandia Technical Report SAND96-8226*, April 1996.
10. C.D. Moen, P.A. Spence, J.C. Meza, T.D. Plantenga. Automatic Differentiation for Gradient-Based Optimization of Radiatively Heated Microelectronics Manufacturing Equipment, *Proceedings of 6th AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 1167-1175, Sept. 4-6, 1996, Bellevue, WA.
11. V. Torczon. On the convergence of pattern search methods, *SIAM J. Optimization*, Vol. 7, No. 1, 1-25, February, 1997.

DISTRIBUTION:

1	MS 0188	C. E. Meyers, 4523
1	MS 0439	G. M. Reese, 9234
1	MS 0441	R. W. Leland, 9226
1	MS 1111	B. A. Hendrickson, 9226
1	MS 9001	T. O. Hunter, 8000
1	MS 9003	D. L. Crawford, 8900
10	MS 9011	J. C. Meza, 8950
1	MS 9214	T. D. Plantenga, 8117
1	MS 9214	C. H. Tong, 8117
1	MS 9042	C. D. Moen, 8345
1	MS 9042	P. A. Spence, 8345
3	MS 9019	Central Technical Files, 8940-2
4	MS 0899	Technical Library, 4916
1	MS 9021	Technical Communications Dept., 8815/Technical Library, MS 0899, 4916
2	MS 9021	Technical Communications Dept., 8815 for DOE/OSTI