

# Parallelized direct search of a binary objective

P. Grant and D.G. Walker

Vanderbilt University

## Abstract

Many scientific and engineering problems are solved by utilizing simulations of computationally intensive mathematical models within massive design spaces. As a result, parametric studies of these models are cost prohibitive in terms of computational time. The focus of this work is a particular kind of parameter sweep problem where the simulation of the model, given real parameters, results in a binary value. The goal of this work is to design, develop, and implement a parallel algorithm for bounding a binary objective when the simulation is a computationally intensive mathematical model. A fully-functioning implementation is provided for a two-dimensional example using a client-server architecture. Results show that a straight bisection search is approximately 50% faster than a full parametric sweep for most continuous functions. With parallelization and load balancing, the simulation is remarkably faster, exhibiting near-linear speedup up to 16 processors for most functions.

## 1 Introduction

Engineering analysis is often concerned with identifying conditions responsible for catastrophic failure. For example, consider the load required to buckle a column. Using simple materials relations and static analysis, an estimate can be obtained for a range of column size and load. However, for complex systems, the prediction of failure can be difficult due to the interaction of many forces and the confounding interaction of many parameters. In these cases, detailed models that govern the fundamental behavior of systems are designed to predict the performance of systems immediately prior to failure. Depending on the application, models can be computationally expensive, and exploration of large design spaces can be cost prohibitive. Yet engineers often want to know where in the design space a system will fail.

The calculation of microelectronic device failure due to a single energetic particle strike [1, 2] is an example of an application that requires significant computational time for a single design point. Failure of a device is determined by the particle energy, biasing conditions, and strike location and direction, among other parameters. Simulations for a single design point require hours of computer time. As a result, to locate all combinations of bias, particle energy and direction that result in failure would require years of computer time.

Another example that does not deal directly with failure but is a similar class of engineering problem is ignition prediction. In diesel engines, ignition is dependent on a variety of parameters such as temperature, pressure, crank angle, cylinder charge and so forth. To predict whether ignition will occur, a detailed chemistry and thermodynamic simulation must be performed allowing the states to evolve in time as the gas is compressed by the cylinder [3]. Failure in this context would actually be ignition, but the results of the simulation indicate whether the fuel mixture burns or not. These models require expensive computations.

These examples fall into a class of problems that are characterized by a binary objective: the diesel fuel ignited or it didn't; the ion destroyed the device or it didn't; the column buckled or it didn't. Furthermore, simulation results on the 'un-failed' side are not terribly useful for determining how close to failure we might be; we simply now the system did not fail. Results obtained from design points of the failed side are often unavailable because the simulation entered an error condition or the quantities are non-physical (due to failure). Therefore, we can't depend on the simulation results to provide any more information other than failure or not-failure. In essence we are attempting to find the location of the discontinuity ( $x_d$ ) in a Heaviside step function:

$$H(x) = \int_{-\infty}^x \delta(x' - x_d) dx'. \quad (1)$$

Due to the lack of information from a function evaluation, binary objective problems are amenable only to direct search methods much like those problems without numerical objectives [4]. Problems without numerical objectives still usually have a ‘better’ or ‘worse’ designation between any two points even though the objective is not quantified. This qualitative evaluation still gives a search direction. Our binary objective is similar except that we don’t have enough information to determine a direction to proceed. We are assuming that gradients are non-existent or meaningless except across the interface where the objective changes by a unit step. Nevertheless, direct search methods can be applied to these problems.

The binary objective problem somewhat resembles that of an edge detection problem in image processing [5]. However, traditional approaches still rely on higher order derivatives to smooth the image and locate where in an image the intensity changes dramatically. The Canny approach [6] assumes noise is inherent to the data and that all function evaluations have already been acquired (pixel color). These two assumptions are contrary to our problem and make the standard edge detection algorithms intractable for our approach. Moreover, the type of edge being detected is usually quite different than our design space surface. Nevertheless, the interpolation line segment approach used by some edge detection routines [7] is partially responsible for the interpolation scheme used in our approach to speed up serial analysis. Finally, edge detection problems do not scale to multiple dimensions, making these approaches only peripherally related.

The binary objective problems described here can be treated as trivially parallel in that the results of each design point are independent of the results from other design points. Therefore, high-performance solutions can greatly reduce the total amount of time an engineer might have to wait for results. The simplest approach would be to discretize the design space and evaluate each point on a different processor. However, this approach does not limit the search in any way and can also become computationally prohibitive for multi-dimensional searches or when a limited number of processors is available. In general the brute force approach scales as  $N^p$ , where  $N$  is the number of parameters values to be tested for a single parameter and  $p$  is the number of parameters to be tested.

The present work focuses on a parallelized and load balanced bisection of the design space to locate a sudden transition at the discontinuity in the objective. We define the mathematical formulation of the problem and solution approach, explore how the solution can be improved serially, parallelize a basic bisection algorithm, and add load balancing. Performance of the serial solution is measured by the number of function evaluations required to obtain the threshold in the objective. Performance of the parallel algorithms is indicated by the speedup.

## 2 Mathematical Model and Solution Approach

The mathematical model of the problem is described in terms of adjustable parameters that govern the physics given by real variables represented by  $z_k \forall k \in [1, 2, \dots, n]$ . The parameter space is then the set of all possible design points where  $f(z_1, z_2, \dots, z_n) \rightarrow [0, 1]$  is a mapping from  $P$  to a binary variable. The problem can be defined as the following: Find sets  $V_0 \subseteq P \ni f(z_1, z_2, \dots, z_n) \rightarrow 0$  and  $V_1 \subseteq P \ni f(z_1, z_2, \dots, z_n) \rightarrow 1$ .

To make the problem more tractable, it is assumed that the boundary between the sets can be represented as a continuous  $(n - 1)$ -dimensional surface  $S$  in  $P$ . Because of this simplifying assumption, the solution may be represented by this surface and a binary value that indicates whether the value 1 lies above or below the surface.

To illustrate the solution process and provide metrics for evaluating the efficacy of the algorithm, a two-dimensional design space search is considered. Assume that the boundary  $S$  can be represented by a continuous curve in the  $x - y$  domain and that  $x \in [x_0, x_1]$  and  $y \in [y_0, y_1]$  defines the parameter space  $P$ . The goal is to find the boundary  $S$  to within  $x_{tol}$  along the  $x$ -direction and  $y_{tol}$  along the  $y$ -direction.

The solution proceeds by allowing  $x_i = x_0 + ix_{tol} \forall i \ni x_i < x_1$ . For each  $x_i$ , perform bisection along the  $y$ -direction to bracket the boundary.

Table 1: Number of function evaluations required to locate the surface  $S$  identified above to a tolerance of  $10^{-4}$ . The row labels identify the type of interpolating scheme and the column labels identify the type of interface function.

Type	line	parabola	poly-9	sine
Akima	9336	8153	3870	20243
Akima-periodic	10434	10278	3580	20859
Cspline	4421	3922	3574	20859
Cspline-periodic	21164	5313	3578	20985
Linear	4421	4694	3574	20859
None	21021	21002	2724	21021
Polynomial	341670	320794	3749	340084

The number of function evaluations can be estimated based on the tolerance required for each direction without having any knowledge of the shape of the surface  $S$ . The bisection algorithm gives

$$\mathcal{O} \left[ \frac{x_1 - x_0}{x_{tol}} \log \left( \frac{y_1 - y_0}{y_{tol}} \right) \right] \quad (2)$$

number of function evaluations. This is a significant improvement over a parametric sweep, which requires

$$\mathcal{O} \left( \frac{x_1 - x_0}{x_{tol}} \frac{y_1 - y_0}{y_{tol}} \right) \quad (3)$$

function evaluations. Assuming the boundaries and tolerances are the same in each direction, the total number of function evaluations required for finding the surface  $S$  is  $N \log(N)$  for the bisection instead of  $N^2$  for the full parametric sweep.

### 3 Serial Improvements

To reduce the number of function evaluations even further, an approximation scheme can be used, which will leverage information already gathered about the location of the surface  $S$ . These approximations serve as a guide for the remaining bisections. Types of approximation used include all of the interpolation types supported by the GNU Scientific Library [8]: linear, polynomial, cubic spline, periodic cubic spline, akima spline, and periodic akima spline interpolation. Each type of interpolation was tested on a boundary represented by a line, sine, parabola, and polynomial. The results in terms of the number of function evaluations are presented in Table 1. The results suggest that simple interpolation (linear) is adequate for the vast majority of function types.

These functions examined are all smooth, and so the results are quite impressive. Cubic splines perform very well, but performance is worse than the regular bisection method in the case of the  $9^{th}$  degree polynomial, which wildly oscillates outside of the parameter space.

### 4 Parallel Algorithm

This problem can be parallelized easily because each function evaluation is assumed to be independent of all other design points. Therefore, little extra work is required to parallelize the algorithm: one must only assign a single  $x$ -value to each node and allow the bisection for that  $x$ -value to take place on that node. The number of function evaluations does not change when the algorithm is parallelized in this way.

A parallel version of the search algorithm was developed using MPI with ethernet network connectivity. No attempt to balance the load was made in this initial trial. Because the number of design points in the  $x$ -direction was fixed (based on the tolerance), each processor was assigned a fixed number of evaluations in that direction. For each  $x$  design point, a bisection was performed in the  $y$  direction until the required tolerance was reached. The test problem was to find the threshold given by  $y = 0.75x + 0.13$  on the domain  $x_0 = 0$ ,  $x_1 = 1$ ,  $y_0 = 0$ , and  $y_1 = 1$ . The model returned a one if the design point was above the line and 0 if the design point was below the line. The required tolerances for each direction were  $x_{tol} = 10^{-2}$  and  $y_{tol} = 10^{-3}$ . No interpolation was performed to reduce the number of iterations.

There are three possible sources of task distribution that could prohibit the parallelization from achieving a linear speedup.

1. The number of tasks does not divide easily among the available processors. For example, four function evaluations do not divide equally across three processors. While one processor is finishing the final function evaluation, the other three processors will sit idle.
2. Each processor may be assigned a different number of function evaluations due to predictive algorithms. For example, if one can provide a good guess where the interface is based on previous results, fewer iterations may be allowed to achieve the specified tolerance. Therefore, the processors where a predictive algorithm is effective may finish their tasks earlier than other processors.
3. Each function evaluation may require a different amount of time. For example, in the case of the ion-strike simulations described previously, the average time to determine whether a device failed given a single design point is four hours. Yet, some failure conditions can be estimated early in the simulation (of the order of 20 minutes of simulation time) and some simulations that come close to failure but ultimately survive can require ten to twenty hours of CPU time. Therefore, depending on the number of processors, it is conceivable that several processors may sit idle while the last few simulations finish the calculations.

In the foregoing tests, the number of function evaluations (1200) is much greater than the number of processors (up to 17), therefore, the load based on the number of function evaluations (point 1) should be similar. Concerning point 2, we are not using any predictive algorithm, so the number of function evaluations will not change based on previous results. Finally, if each function evaluation requires the same amount of time, then we expect linear speedup as shown in Figure 1 for the constant function evaluation time case. These results were obtained on a homogeneous cluster.

In real engineering systems, however, we can not depend on each function evaluation to be equivalent in terms of computational requirements (point 3). This wait time at the end of the simulation will reduce the speedup of the overall analysis as shown in Figure 1. In the random function evaluation time case, a random amount of sleep time was executed before the function returned with a one or a zero. This was implemented to simulate varying function evaluation times. This lack of linear speedup can be addressed through a load balancing scheme.

The strategy used to accomplish load balancing is that of a simple client-server architecture. The server maintains a queue of tasks to be completed. Initially the queue contains, in order, the low and high  $y$ -values for each  $x$ -value. The server doles out tasks from the front of the queue. When a task completes, a new task is generated and is placed at the end of the queue. A new task can be assigned in constant time. This form of load balancing addresses all three sources of degradation of linear speedup. However, it requires a dedicated server to dole out design points and requires a bit more communication overhead. Nevertheless, the results are dramatically improved. Figure 2 shows the maximum wait time of any processor in the group and the average wait time for all the processors in the group as a function of the number of processors in the group. The amount of wait time increases with the number of processors when no load balancing is performed. With load balancing, however, the wait time is decreased significantly. Therefore, the efficiency of this strategy is nearly perfect leading to a scalable speedup up to 17 processors.

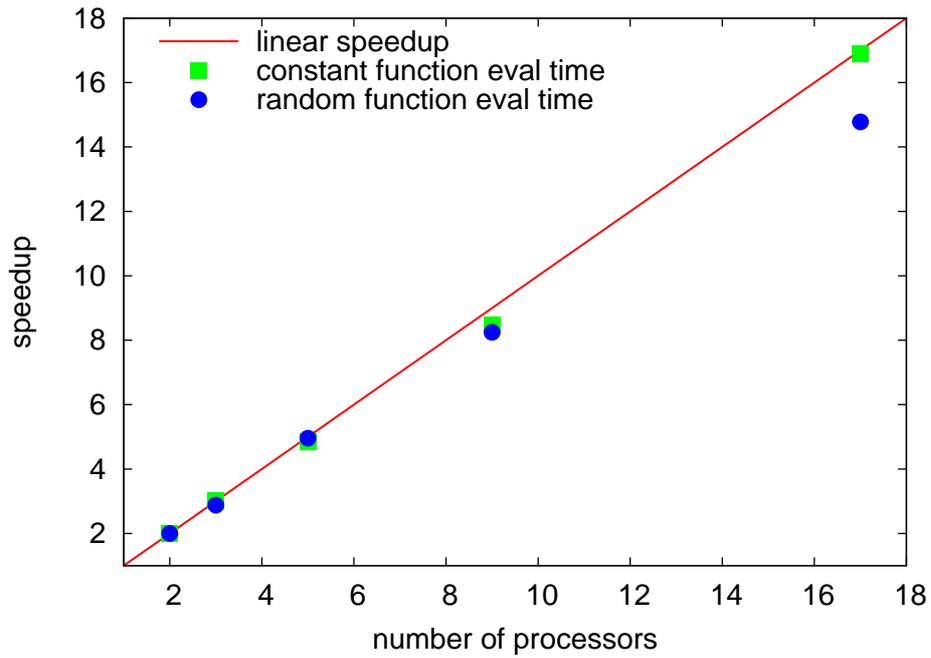


Figure 1: Speedup measured for a constant function evaluation time and a random function evaluation time (from 0 to 5 seconds of sleep).

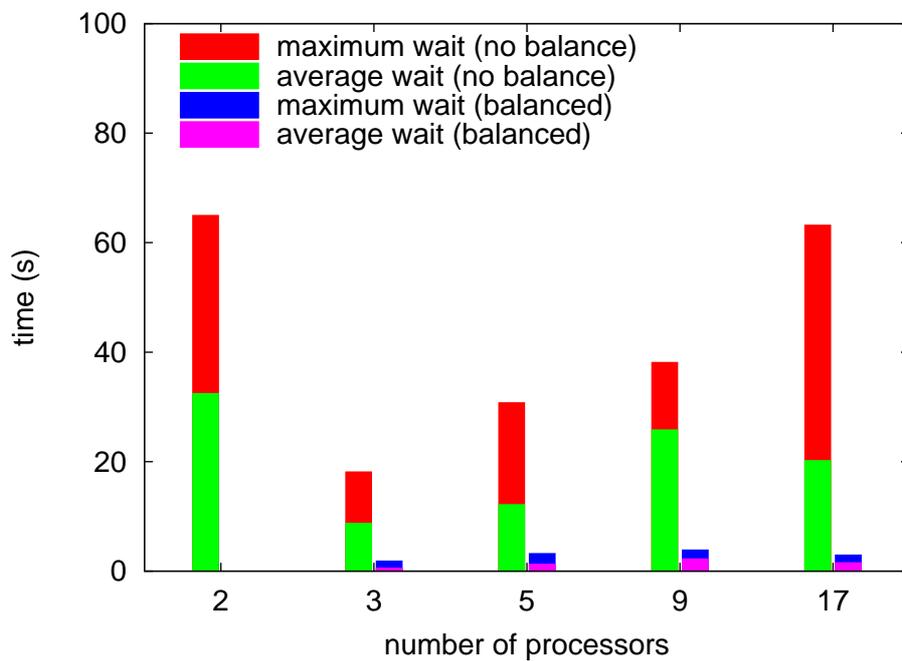


Figure 2: The maximum and average wait times for a non-balanced case and when load balancing is employed.

## 5 Discussion

The performance of the parallel algorithm can be improved by including the approximation scheme used in the serial algorithm. However, this effort introduces additional complexity. As processors are added to the simulation, the accuracy of the approximation decreases as illustrated in the following example. If all execution occurs on a single processor, then when evaluation  $k$  is assigned,  $(k - 1)$  other results may be used in the approximation. On the other hand, if 10 clients are used, then when evaluation  $k$  is assigned, at most  $(k - 10)$  other results may be used. Furthermore, the approximation in this case is an interpolator, meaning that the approximation must pass through each point. For these reasons, it is extremely difficult to place bounds on the number of evaluations required for a particular number of processors and type of approximation. As illustrated in the case of the  $9^{th}$  degree polynomial, the number of function evaluations does not always decrease when using approximation.

Although the example was cast in a two-dimensional space, the concepts here can be extended to multiple dimensions with a multi-dimensional bisection [9] algorithm. In fact, the two-dimensional problem may also be improved with a multi-dimensional bisection approach. However, the scalability is no longer a simple function of the problems size, and the parallelization is no longer trivial.

Improvements can be made to the load balancing strategy by introducing a list of

$$\log \left( \frac{y_1 - y_0}{y_{tol}} \right) \quad (4)$$

queues,  $q_i$ , such that  $q_i$  contains only tasks corresponding to an  $x$ -value that has had  $i$  tasks executed already. This strategy is commonly used in CPU schedulers and runs in constant time.

Other improvements could include mixing an approximation strategy with a load balancing strategy. This is difficult because the approximation, and thus the scalability, worsens as the number of processors increases. Also, the interpolators used should be replaced with better approximations, since the partial results are inexact at best.

## 6 Conclusions

Binary objectives occur naturally in many areas of science and engineering. Therefore, understanding these types of problems is important to efficient analysis and designs of engineering systems. Examining the problems by forming a mesh of the parameter space seems to be the accepted solution, which is certainly easy to extend to  $n$  dimensions. However, complicated physical models often require vast computing resources and an algorithm that scales better and reduces the computational cost is essential. We have shown that a simple bisection algorithm and simple parallelization with load balancing can reduce the computational cost by an order of magnitude. These results are limited to a boundary that is continuous in the parameter space. Moreover, we did not test the scheme for higher dimensions. Nevertheless, these results lay the framework for continued study in optimization problems with binary objectives.

## References

- [1] D. M. Fleetwood, P. S. Winokur, and P. E. Dodd. An overview of radiation effects on electronics in the space telecommunications environment. *Microelectronics Reliability*, 40(1):17–26, January 2000.
- [2] R. D. Schrimpf. *Radiation Effects on Embedded Systems*, chapter Radiation Effects in Microelectronics, pages 11–29. Springer Netherlands, 2007.

- [3] S.C. Kong and RD Reitz. Multidimensional modeling of diesel ignition and combustion using a multistep kinetics model. *Journal of Engineering for Gas Turbines and Power*, 115:781, 1993.
- [4] T. G. Kolda, R. M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45(3):385–482, 2003.
- [5] D. Marr and E. Hildreth. Theory of edge detection. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 207(1167):187–217, February 1980.
- [6] John Canny. A computational approach to edge detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-8(6):679–698, November 1986.
- [7] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, January 1988.
- [8] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi. *GNU Scientific Library Reference Manual*. Network Theory Limited, 3<sup>rd</sup> edition, 2001.
- [9] G. R. Wood. The bisection method in higher dimensions. *Mathematical Programming*, 55:319–337, 1992.