A Hierarchical Parallel Scheme for Global Parameter Estimation in Systems Biology

J. He*, M. Sosonkina††, C. A. Shaffer*, J. J. Tyson†, L. T. Watson*, J. W. Zwolak*

Department of Computer Science*, Department of Biology†

Virginia Polytechnic Institute and State University

Blacksburg, Virginia 24061

Ames Laboratory, Iowa State University††

236 Wilhelm Hall, Ames, IA 50011.

Contact E-mail: jihe@vt.edu

Abstract—This paper presents a sophisticated and efficient parallel scheme for the DIRECT global optimization algorithm of Jones et al. (1993). Although several sequential implementations for this algorithm have been successfully applied to large scale MDO problems, few parallel versions of the DIRECT algorithm have addressed well algorithm characteristics such as a single starting point, an unpredictable workload, and a strong data dependency. These challenges engender many interesting design issues including domain decomposition, data access and management, and workload balancing. In the present work, a hierarchical parallel scheme has been developed to address these challenges at three levels. Each level is supported by parallel and distributed data structures to access shared data sets, distribute workload, or exchange messages. Parameter estimation problems in systems biology provide an ideal application context for the present work. Global nonlinear parameter estimation results obtained on a 200 node Linux cluster are given for a cell cycle model for frog eggs.

Index Terms—DIRECT (DIviding RECTangles) algorithm, global optimization, GPSHMEM (generalized portable shared memory), load balancing strategy, multidisciplinary design optimization, parallel and distributed data structures, parameter estimation, systems biology

1. Introduction

The optimization algorithm DIRECT (DIviding RECTangles) is a global search algorithm proposed by Jones et al. [11], designed as an effective approach to solve continuous optimization problems subject to simple constraints. In the past decade, DIRECT has been successfully applied to many modern large-scale multidisciplinary engineering problems ([2], [3], [9], and [20]). Recently, DIRECT has been used in global nonlinear parameter estimation problems in systems biology [14]. However, unnecessary overhead and complexity caused by inefficient implementation inside other software packages (e.g., Matlab)
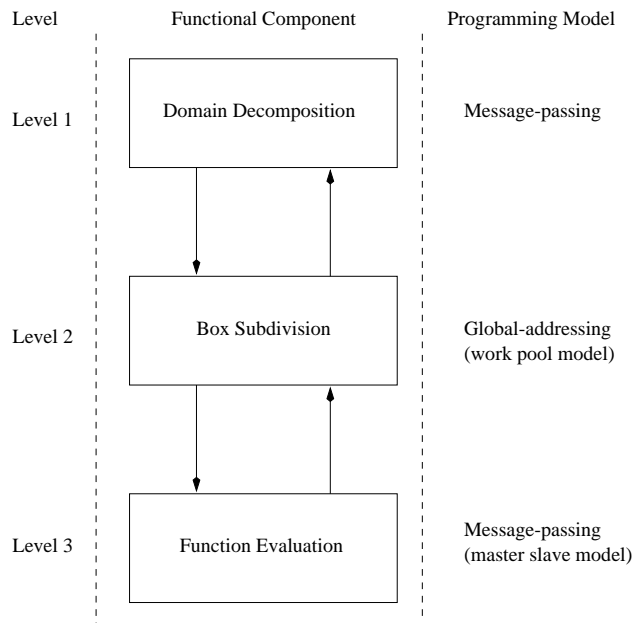

Figure 1.1. Three functional levels with a mixed programming paradigm.

may obscure DIRECT's advanced features. Some computational biologists are attracted by its unique strategy of balancing global and local search, its selection rules for potentially optimal regions according to a Lipschitz condition, and its easy-to-use black-box interface. Like other global optimization approaches of [4] and [6], DIRECT is being challenged by high-dimensional nonlinear models for parameter estimation. The present work applies DIRECT to a parameter estimation problem for a frog egg model from the JigCell systems biology project [1]. Modeling details and results obtained from a local method can be found in [21].

As the scale of both problems and clusters of workstations grows, computational parallelism of

1

optimization algorithms has become a very active research area. However, the nature of the DIRECT algorithm presents both potential benefits and difficulties for a sophisticated and efficient parallel implementation. Gablonsky [5] and Baker et al. [2] are among the few parallel DIRECT implementations known in the public domain. In [5], Gablonsky adopts a master-slave paradigm to parallelize the function evaluations, but little discussion is given to the issue of parallel performance and potential problems, such as load-balancing and interprocessor communication, both of which raise many challenging design issues. A major contribution in [2] is a distributed control version equipped with dynamic load balancing strategies. Nevertheless, that work did not fully address other design issues such as a single starting point and a strong data dependency.

The present work proposes a hierarchical parallel scheme (shown in Figure 1.1) to address design issues by three function components in different levels. Level 1 splits the entire search space to start the processing at multiple points, detects the stopping conditions, and merges the results at the end. This level transforms the original single-start sequential algorithm to a multistart parallel algorithm. Below Level 1, Level 2 uses GPSHMEM [16] to establish a global addressing space to ease the strong data dependency problem occurring in the algorithm step (refer to Section 2) that identifies a set of potentially optimal boxes to be subdivided at the next iteration. This globally shared data structure also forms a work pool paradigm [7] to apply a dynamic load balancing strategy to adjust the workload among processors at Level 2. Similar to [2], a master-slave paradigm is used between Level 2 and 3 for distributing function evaluation tasks.

All three levels take advantage of dynamic process management in MPI-2 [8] so that processors are assigned to different levels at run time. As shown in Figure 1.1, a mixed programming paradigm is constructed with a global addressing model and a message passing model. A similar style of subgrouping processors for multiple level parallelism (MLP) was called GSPMD (group single program multiple data) in [18]. By contrast, in [12] and [17], a hybrid parallel programming model involving thread-level parallelism (OpenMP) and message passing (MPI) was used to vary the number of threads and CPUs in order to simplify the dynamic load balancing strategy in MLP. Mixed parallel programming models may

become a trend due to the recent deployment of many large scale clusters of shared memory multiprocessor workstations [12].

Section 2 describes the DIRECT algorithm and the parallel design issues. The proposed parallel implementation is described in Section 3, with the systems biology problem described in Section 4. Computational results and some performance analysis are included in Section 5.

2. Design Challenges

The sequential DIRECT algorithm can be described by the following six steps [11], given an objective function $f(x)$ and the n-dimensional design space $D = \{x \mid l \leq x \leq u\}$.

Step 1. Normalize the design space D to be the unit hypercube. Sample the center point c_i of this hypercube and evaluate $f(c_i)$. Initialize $f_{\min} = f(c_i)$, evaluation counter $m = 1$, and iteration counter $t = 0$.

Step 2. Identify the set S of potentially optimal boxes.

Step 3. Select any box $j \in S$.

Step 4. Divide the box j as follows:

  (1) Identify the set I of dimensions with the maximum side length. Let δ equal one-third of this maximum side length.

  (2) Sample the function at the points $c \pm \delta e_i$ for all $i \in I$, where c is the center of the box and e_i is the ith unit vector.

  (3) Divide the box j containing c into thirds along the dimensions in I, starting with the dimension with the lowest value of $w_i = \min\{f(c + \delta e_i), f(c - \delta e_i)\}$, and continuing to the dimension with the highest w_i. Update $f_{\min}$ and m.

Step 5. Set $S = S - \{j\}$. If $S \neq \emptyset$ go to Step 3.

Step 6. Set $t = t + 1$. If iteration limit or evaluation limit has been reached, stop. Otherwise, go to Step 2.

Steps 2 to 6 form a processing loop controlled by two stopping criteria—limits on iterations and function evaluations. Starting from the center of the initial hypercube, DIRECT makes exploratory moves across the design space by probing the potentially optimal hyperboxes. "Potentially optimal" is precisely defined in [11], but roughly a hyperbox is potentially optimal if, for some Lipschitz constant, the objective function is potentially smaller in that hyperbox than in any other hyperbox. It is observed in Step 4
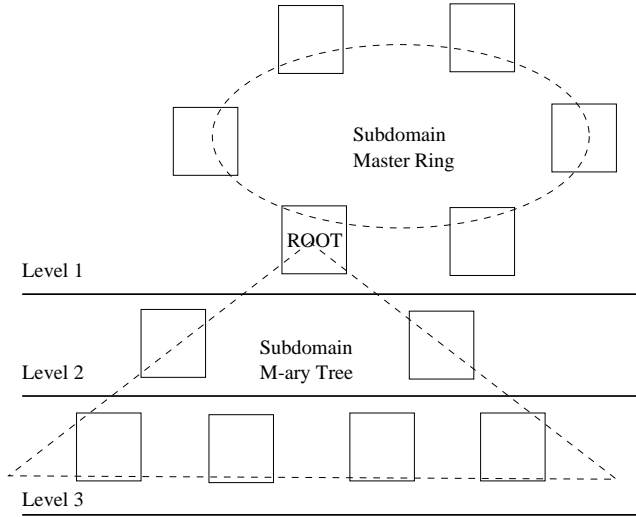
Figure 3.1. A 3-level hierarchical parallel scheme.

that multiple new hyperboxes are generated for each potentially optimal hyperbox. The multiple function evaluation tasks at each iteration give rise to a natural functional parallelism used both in [5] and [2].

In addition, a few design challenges are also observed here. First, the algorithm starts with one normalized domain, which produces simply one evaluation task for all the acquired processors. With a single starting point, load balancing is always an issue at an early stage, even though the situation will be improved as the algorithm progresses by subdividing the domain and generating multiple evaluation tasks. Second, the number of boxes subdivided at each iteration is unpredictable depending on the result of identifying the potentially optimal boxes. For iterations that generate fewer new boxes, a load imbalance occurs with some processors sitting idle. Third, a strong data dependency exists throughout the algorithm steps. Only Step 2 and Step 4 can be parallelized, respectively, in terms of functional parallelism and data parallelism. Both involve shared data sets with a considerable growth rate. Efficient data decomposition and access methods are the major issues here.

3. Parallel Scheme and Implementation

A hierarchical parallel scheme is proposed here to address the above mentioned design challenges. It consists of three logical levels as shown in Figure 3.1. Each level deals with different design aspects including domain decomposition, load balancing, and task or data parallelism.

Level 1 holds a ring of processors, each of which is responsible for a DIRECT search in an assigned subdomain. A root processor is noted as the first node to start the processing and spawn other processes, one per processor, on the ring. The problem domain D is decomposed by the root processor into $S^2 \leq N_1 = \lfloor \sqrt{P} \rfloor$ subdomains, where P is the total number of processors. The parallel scheme described here requires $P \geq 16$. When the number of available processors is below 16, P should be set as the total number of processes, some of which may run on the same physical processor. The decomposition is accomplished in two phases. In phase one, the root processor finds the longest dimension of the domain and subdivides it into $S = \lfloor \sqrt{N_1} \rfloor$ partitions, each of which will be processed by S processors. In phase two, inside each of the resulting partitions, the currently longest dimension is subdivided into S parts. The total number of subdomains S^2 equals the number of processors needed for Level 1. As a result, exactly one processor is in charge of one subdomain.

For each subdomain, a process is dynamically spawned by the root processor to be the subdomain master starting a DIRECT search. Subsequently, a logical ring of processes is formed by the root and the newly spawned processes to adopt a termination detection strategy depending on different stopping criteria, including the total number of iterations, function evaluations, and the degree of subdivision. The overall termination condition is to keep every subdomain active until all subdomains have satisfied the specified stopping criteria. This rule may result in more computational cost. Thus, the stopping condition for the proposed parallel scheme is effectively a lower bound on the computational cost, while for the sequential algorithm, it is an exact limit for computation cost.

A ring topology naturally fits the equal relationship among subdomain masters at Level 1. Moreover, it represents the dependency of the stopping process of each subdomain on other subdomains on the ring. A subdomain terminates search activities only after all other subdomains have reached their specified stopping criteria. This discentralizes the termination control among the ring, thus avoiding the bottleneck at the root subdomain master. On the other hand, the communication latency on a ring is higher than on some other topologies, such as a tree. To reduce the communication time on the ring and improve the performance at Level 1, future work can consider a

tree topology instead of a ring. The stopping process for the entire domain is controlled by a token passed in the ring that consists of subdomain masters. It is originated with value zero from the root subdomain master R_0 and passed around the rest of the ring. After each iteration of DIRECT, each subdomain master R_i checks if a token has arrived. If not, DIRECT proceeds. If the token was received, and the stopping condition obtained from M_0, the root mini master at Level 2 (in this subdomain), is not satisfied, the token value v is reset to zero and the token is passed along in the ring, and DIRECT continues. If the token was received and the local stopping condition is met, the token value v is incremented by one. If $v = S^2$, a termination message is sent to all R_i. If $v < S^2$, the token is passed along and DIRECT continues (so this subdomain is being explored more than required by the stopping criteria). Lastly, R_0 will collect the final results and report the total number of evaluations and the range of number of iterations as well as minimum diameters.

For Levels 2 and 3, there are $P - S^2$ processors available. Each subdomain master dynamically spawns $\lfloor M \rfloor$ mini master processors at Level 2 for box subdivision tasks, where

$$M = \sqrt{\frac{P - S^2}{S^2}},$$

derived from $M^2 \times S^2 = P - S^2$. Similarly, each mini master spawns $\lfloor \sigma \rfloor$ or $\lceil \sigma \rceil$ slaves for function evaluation tasks, where

$$\sigma = \frac{P - S^2(1 + \lfloor M \rfloor)}{S^2 \lfloor M \rfloor}.$$

A $\lfloor M \rfloor$-ary tree structure (in Figure 2.1) is rooted at a subdomain master, which is at the boundary of Levels 1 and 2. Therefore, a subdomain master plays two roles—one for communicating with the processors on the subdomain ring, and the other for managing the search in the subdomain $\lfloor M \rfloor$-ary tree. Subdomain management tasks here include updating current search results, and detecting local stopping conditions. Pseudo code 3.1 shows the interactions between the root mini master M_0 and other M_js $(j = 1, 2, \ldots, \lfloor M \rfloor - 1)$ in a subdomain i $(i = 0, 1, \ldots, S^2)$, which is managed by a subdomain master R_i.

M_0 receives DIRECT parameters (problem size N, domain D, and stopping conditions C_{stop}) from R_i;
broadcast DIRECT parameters to all M_j;
$done :=$ FALSE;
while TRUE
  if M_0 **then**
    receive a message from R_i;
    if not a termination message **then**
      send $done$ to R_i;
      send a message to keep M_js working;
      run one DIRECT iteration (reduce intermediate results);
      if C_{stop} satisfied **then**
        $done :=$ TRUE;
      end if
      continue;
    else
      send a termination message to all M_js;
      terminate workers at Level 3;
      store the reduced results;
      break;
    end if
  else
    M_js receive a message from M_0;
    if not a termination message **then**
      run one DIRECT iteration (reduce intermediate results);
    else
      break;
    end if
  end if
end while
M_0 sends the final results to R_i;

Pseudo code 3.1

At Level 2, mini masters collaborate on identifying the potentially optimal box set on a shared data structure in a global addressing space based on GPSHMEM [16]. This algorithm step is one of most interesting challenges in the parallelization of DIRECT. Two sets of global shared data structures OPTSET and OPTSET_INDEX (structures in solid lines shown in Figure 3.2) are used to hold potentially optimal boxes. The structures in dashed lines (SETLINK_INDEX and additional OPTSETs) will be implemented in the next version to enlarge OPTSET at run time. Basically, SETLINK_INDEX holds a list of pointers to dynamically allocated OPTSETs in the global addressing space. For the maximum flexibility in adding OPTSETs, SETLINK_INDEX can be

SETLINK_INDEX   OPTSET

M0   M1   M2   M3

OPTSET_INDEX

Figure 3.2. Data structures in GPSHMEM.

implemented as a local linked list on each processor at Level 2.

Both structures OPTSET and OPTSET_INDEX are allocated and distributed across all the mini masters, which use one-sided communication operations such as "put" and "get" to access shared data. These one-sided operations provide a direct access to remote memory with less interaction between communicating parties. At each iteration, M_j puts all the boxes with the smallest function values for different box sizes to its own portion in OPTSET and updates its index in OPTSET_INDEX. Next, M_0 gets all boxes in OPTSET and merges the boxes with the same size. After M_0 finds all potentially optimal boxes, it balances the number of boxes for each M_j's portion in OPTSET. Detailed algorithm steps are in Pseudo code 3.2. Finally, each M_j gets its portion of the workload, removes some boxes (if any) that are assigned to other mini masters, and starts processing each potentially optimal box. Each box is tagged with a processor ID and other indexing information to be tracked by its original owner. To minimize the number of local box removals and additions, M_0 restores the boxes back to their contributors before it starts load adjustment. This also guarantees maximum data locality on each processor. On each processor, a set of local data structures for storing and processing boxes are reused from the sequential DIRECT implementation described in [10]. GPSHMEM is a suitable tool to handle irregular data structures, and dynamic/unpredictable data access patterns characteristic of DIRECT.

---

copy M_0's portion in OPTSET to LOCALSET;
merge boxes from M_j's portion in OPTSET
  to LOCALSET;
find N_{box} potentially optimal boxes in LOCALSET;
in LOCALSET, pick out boxes given by M_j to its

portion in OPTSET;
$avgload := \lceil (N_{box}/N_{proc}) \rceil$;
$i := 0$;
while TRUE
  if $i = N_{proc} - 1$ **break**;
  if OPTSET_INDEX$(i+1) < avgload$ **then**
    $i_1 := i$;
    while TRUE
    $underload := avgload -$ OPTSET_INDEX$(i+1)$;
    $i_1 := (i_1 + 1)\% N_{proc}$;
    if $i_1 = i$ **break**;
    if OPTSET_INDEX$(i_1+1) > avgload$ **then**
      $overload :=$ OPTSET_INDEX$(i_1+1) - avgload$;
      if $overload \geq underload$ **then**
        shift enough load over;
        OPTSET_INDEX$(i+1) := avgload$;
        OPTSET_INDEX$(i_1+1) :=$
          OPTSET_INDEX$(i_1+1) - underload$;
        break;
      else
        shift some and look for more;
        OPTSET_INDEX$(i+1) :=$
          OPTSET_INDEX$(i+1) + overload$;
        OPTSET_INDEX$(i_1+1) := avgload$;
      end if
    end if
    end while
  end if
  $i := i + 1$
end while

Pseudo code 3.2

As shown in Pseudo code 3.2, a centralized dynamic load balancing strategy is applied at Level 2 with shared data structures in GPSHMEM. At Level 3, workload balancing of processors is also centralized with a master-slave model. For Level 2, the workload is box subdivision on the subdomain assigned at Level 1. The root mini master adjusts the workload in the globally shared structure, and each mini master subdivides its share of potentially optimal boxes and distributes the function evaluation tasks down to its slaves at Level 3. Workload is spread from the subdomain master to mini masters at Level 2 and their slaves at Level 3. In some way, this is similar to a sender-initiated strategy in MLB (multilevel load balancing) defined in [13], where workload is sent down to a $\lfloor M \rfloor$-ary tree structure. Although the control mechanism is simple, this strategy suffers a common bottleneck problem with other centralized methods. A distributed control version will be considered in future work.

4. Systems Biology Modeling

The ultimate goal of molecular cell biology is to understand how the information stored in the genome is read out and put into action as the physiological behavior of a living cell. At one end of this continuum, genes direct the synthesis of polypeptide chains, and at the other end, networks of interacting proteins govern how a cell moves, feeds, responds, and reproduces. The model presented here is of the latter type and can be represented as a system of ordinary differential equations (ODE).

The model describes the activity of MPF in frog egg extracts. A frog egg extract is the cytoplasmic protein mix obtained by disrupting (or disintegrating) the membranes of hundreds of frog eggs and separating out the cytoplasm containing protein. The extract can be easily manipulated and assayed compared with intact frog eggs but it still has nearly the same chemical kinetics as an intact frog egg. MPF plays an important role in the chemical kinetics of frog eggs. Primarily the activity of MPF controls when the frog egg enters mitosis and subsequently divides. MPF is a dimer of Cdk1 and Cyclin and is at the center of the model.

The model includes two other proteins that regulate MPF activity: Wee1 and Cdc25. Wee1 inhibits MPF activation, and Cdc25 promotes MPF activation. In turn, MPF regulates Wee1 and Cdc25 creating feedback loops. These three proteins are the most important to cell division in frog eggs. The following system of ODEs describes their interactions:

$$\frac{dM}{dt} = \bigl(v'_d(1 - D) + v''_d D\bigr)(C_T - M)$$
$$- \bigl(v'_w(1 - W) + v''_w W\bigr)M, \tag{1}$$

$$\frac{dD}{dt} = v_d \left(\frac{M(1 - D)}{K_{md} + (1 - D)} - \frac{\rho_d D}{K_{mdr} + D} \right), \tag{2}$$

$$\frac{dW}{dt} = v_w \left(-\frac{MW}{K_{mw} + W} + \frac{\rho_w(1 - W)}{K_{mwr} + (1 - W)} \right), \tag{3}$$

where

$$M = \text{[MPF]}/\text{[total Cdk1]},$$
$$D = \text{[Cdc25P]}/\text{[total Cdc25]},$$
$$W = \text{[Wee1]}/\text{[total Wee1]},$$
$$C_T = \text{[total cyclin]}/\text{[total Cdk1]}.$$

M, D, W, and C_T represent scaled concentrations of active MPF, active Cdc25, active Wee1, and total cyclin in the extract, respectively. The parameters v'_d, v''_d, v'_w, v''_w, v_d, K_{md}, ρ_d, K_{mdr}, v_w, K_{mw},


Figure 5.1. Comparison of optimization results generated by single-start DIRECT (dotted curve) and by multistart parallel DIRECT (solid curve).

ρ_w, and K_{mwr} are also scaled, making the system dimensionless.

The parameters of the model are unknown. They can be determined experimentally at a great cost to the experimentalist, and only to the theoretician's immediate benefit. The parameters can also be determined by comparing model simulations to experiments and adjusting the parameters until the model matches the experiments. In the present work, a formal objective function (see [21]) has been defined in terms of the parameters and given to the parallelized DIRECT optimization algorithm to perform a global parameter estimation (nonlinear orthogonal distance regression).

5. Simulation and Results

All the simulations were run on a 200-node Linux cluster with Portable Batch System (PBS). The list of assigned processors generated by PBS can be sorted to map the grouped processes (at each level) to adjacent processors. The LAM/MPI package was chosen for its support of dynamic process management in the MPI-2 standard.

Two groups of simulations were designed to evaluate the present work. The first group validated the optimization results obtained for the parameter estimation problem for the frog egg model described in the previous section. The second group measured the parallel performance for the frog egg parameter estimation problem and a synthetic test function on 16, 32, and 64 processors.

Figure 5.1 shows the optimization progress for the original single-start DIRECT (dotted line) and

**Figure 5.2. Single-start DIRECT result: time lag for MPF activation.**

Figure 5.3. Multistart DIRECT result: time lag for MPF activation.

**Figure 5.4. Single-start DIRECT result: phosphorylation of Wee1 during mitosis, when MPF is more active.**


Figure 5.5. Multistart DIRECT result: phosphorylation of Wee1 during mitosis, when MPF is more active.

for the transformed multistart parallel DIRECT (in solid line), where f_{min} was reduced from subdomain masters at each iteration. As the number of iterations grows, the objective function decreases faster in the case of the multistart parallel DIRECT algorithm.

In a second observation, different optimal parameter sets were found by the single-start and multistart DIRECT after the same number of iterations. Figures 5.2 through 5.5 compare the simulation results using the newly discovered parameter sets (solid lines) with the known set of experimental data (dots) as well as with the simulation results from the parameter sets in [15] and [19] (dotted lines). For the time lag plot, the parameter set by Moore [15] predicts the experimental data points better than that from the multistart DIRECT, although the latter gives a closer match in the case of phosphorylation of Wee1.

This suggests that the present work explores some previously unexplored regions of parameter space, which may lead to new optimal parameter sets with different biological interpretations.

Table 5.1 shows some preliminary parallel performance data for the frog egg 16-parameter estimation problem and a 100-dimensional Griewank function using $p = 3$, 6, and 15 processors in a 2-level subdomain tree, which is a part of $P = 16$, 32, and 64 processors under the 3-level parallel scheme. Here, speedup is computed relative to a *base*, the smallest number of processors required for the hierarchical parallel scheme. Because the 3-level parallel scheme does not have an exact limit for computation cost (as mentioned in Section 3), the time was measured

7

Table 5.1. Parallel timing (in seconds) and efficiency results for different numbers I of DIRECT iterations for the frog egg 16-parameter estimation problem and 100-dimensional Griewank function in a 2-level subdomain tree on 6 and 15 processors with $base = 3$ processors.

| **Frog Egg Model** | | | | | |
I	3	6	$E(6)$	15	$E(15)$
10	1432	1152	0.62	372	0.77
20	3939	2267	0.87	1045	0.75
40	8657	4936	0.88	1960	0.88

| **Griewank Function** | | | | | |
I	3	6	$E(6)$	15	$E(15)$
20	42	34	0.62	27	0.31
60	132	98	0.67	81	0.33
100	719	283	**1.27**	138	**1.04**

for the lower two levels of a single subdomain tree, neglecting any interaction with the ring at Level 1. The formulas for S^2, M, and σ in Section 3 still apply, but only the processors assigned to one subdomain tree (working on the entire problem domain) are used for performance evaluation. When total processor number $P \in [16, 64]$, the 3-level framework splits the domain into four parts. Different ways of splitting the problem domain result in different search problems, so using 16, 32, and 64 processors guarantees the same search problem for the efficiency test. Therefore, the 2-level subdomain tree has $p = 3$, 6, and 15 processors accordingly. The smallest number of processors possible for a subdomain tree is $base = 3$. The parallel efficiency E is defined as

$$E = \frac{S_r}{p/base},$$

where $S_r = \text{Time}_{base}/\text{Time}_p$ is the relative speedup.

In Table 5.1, except for the case with 15 processors in the frog egg parameter estimation problem, the parallel efficiency E increases as the maximum number of iterations I grows. The frog egg parameter estimation problem, while certainly nontrivial, does not generate a large amount of work (function evaluations) per iteration, and so too much processor time is spent waiting and communicating rather than productively computing. Also, for $P < 256$, the number of subdomains $S = \lfloor P^{1/4} \rfloor$ does not greatly accelerate the exploration by DIRECT of the feasible parameter space. Much larger test problems (that generate thousands of function evaluations per


Figure 5.6. One-dimensional Griewank function with parameter $d = 500$.

iteration) are required to fairly assess the performance of the multilevel algorithm proposed here. A budding yeast cell cycle model with 154 parameters is being developed and will become an ideal testing case for the present work. As a synthetic high dimensional test case, a Griewank function was used here. The n-dimensional Griewank function

$$f(x) = 1 + \sum_{i=1}^{n} \frac{x_i{}^2}{d} - \prod_{i=1}^{n} \cos\left(\frac{x_i}{\sqrt{i}}\right),$$

where $d > 0$ is a constant to adjust the noise, has a unique global minimum at $x = 0$, and numerous local minima (one-dimensional instance is plotted in Figure 5.6). The larger the value of d, the deeper the minima values are. The numerical results are for an initial box $[-10, 90]^n$ and $d = 500$. In this test, superlinear speedups ($E > 1$) (highlighted values) observed in Table 5.1 imply that the data generated by DIRECT may be too large to fit into the memory of the *base* number of processors. This certainly degrades its performance due to paging operations from disk. When $P = 16$, a single mini master holds all the data on a processor and sends the function evaluations to two slave processes. In the case of $P = 32$ and $P = 64$, more mini masters are spawned to hold the data, thereby sharing the huge memory burden.

6. Conclusion and Future Work

Key contributions of the present work are i) the transformation from single-start to multistart, ii) the mixed programming model, and iii) the dynamic processor assignment.

In future work, a tree topology will be considered in place of the current ring implementation at Level 1. At Level 2, a new set of data structures SETLINK_INDEX and its associated OPTSETs will be implemented to add flexibility for enlarging the global

addressing space dynamically, as the number of box sizes is unpredictable. Second, a distributed control version for load balancing, desirable to eliminate the bottleneck at the master, will be developed. Third, a queue with adjustable size entries can be used to hold multiple function evaluation tasks on processors at Level 3 to reduce communication overhead. The size of the queue entries depends on the problem granularity. When the ratio of computation to communication is higher, the entry size in the queue is smaller.

Acknowledgments

References

[1] N. A. Allen, C. A. Shaffer, M. T. Vass, N. Ramakrishnan, and L. T. Watson , "Improving the development process for Eukaryotic cell cycle models with a modeling support environment", in *Winter Simulation Conference*, to appear, December, 2003.

[2] C. A. Baker, L. T. Watson, B. Grossman, R. T. Haftka, and W. H. Mason, "Parallel global aircraft configuration design space exploration", in *High Performance Computing Symposium 2000, A. Tentner (Ed.), Soc. for Computer Simulation Internat*, San Diego, CA, 2000, pp. 101–106.

[3] M. C. Bartholomew-Biggs, S. C. Parkhurst, and S. P. Wilson, "Using DIRECT to solve an aircraft routing problem", *Computational Optimization and Applications*, vol. 21, no. 3, pp. 311–323, 2002.

[4] W. R. Esposito and C. A. Floudas, "Global optimization in parameter estimation of nonlinear algebraic models via the Error-In-Variables approach", *Ind Eng. Chemistry and Research*, vol. 37, pp. 1841–1858, 1998.

[5] J. M. Gablonsky, "Modifications of the DIRECT algorithm", PhD thesis, Department of Mathematics, North Carolina State University, Raleigh, NC, 2001.

[6] C. Gau and M. A. Stadtherr, "Nonlinear parameter estimation using interval analysis", in *AIChE Symposium*, vol. 94, no. 320, pp 445-450, 1999.

[7] A. Grama, A. Gupta, G. Karypis, and V. Kumar, *Introduction to Parallel Computing*, Pearson Education Limited, 2nd Edition, 2003.

[8] W. Gropp, E. Lusk, and R. Thakur, *Using MPI-2: Advanced features of the message-passing interface*, The MIT Press, Cambridge, Massachusetts, London, England, 1999.

[9] J. He, A. Verstak, L. T. Watson, T. S. Rappaport, C. R. Anderson, N. Ramakrishnan, C. A. Shaffer, W. H. Tranter, K. Bae, and J. Jiang, "Global optimization of transmitter placement in wireless communication systems", in *Proc. High Performance Computing Symposium 2002*, A. Tentner (ed.), Soc. for Modeling and Simulation International, San Diego, CA, pp. 328–333, 2002.

[10] J. He, L. T. Watson, N. Ramakrishnan, C. A. Shaffer, A. Verstak, J. Jiang, K. Bae, and W. H. Tranter, "Dynamic data structures for a direct search algorithm", *Computational Optimization and Applications*, vol. 23, pp. 5–25, 2002.

[11] D. R. Jones, C. D. Perttunen, and B. E. Stuckman, "Lipschitzian optimization without the Lipschitz constant", *J. Optim. Theory Appl.*, vol. 79, no. 1, pp. 157–181, 1993.

[12] D. J. Marriplis, "Parallel performance investigation of an unstructured mesh Naiver-Stokes solver", *Internat. J. High Performance Computing Appl.*, vol. 16, no. 4, pp. 395–407, 2002.

[13] V. Kumar, A. Y. Crama, and N. R. Vempaty, "Scalable load balancing techniques for parallel computers", *J. Parallel Distributed Computing*, vol. 22, pp. 60–79, 1994.

[14] C. G. Moles, P. Mendes, and J. R. Banga, "Parameter estimation in biochemical pathways: a comparison of global optimization methods", *Genome Res.*, vol. 13, pp. 2467–2474, 2003.

[15] J. Moore, "Private Communication", *Aug.*, 1997.

[16] K. Parzyszek, J. Nieplocha, and R. A. Kendall, "A generalized portable SHMEM library for high performance computing", in *12th IASTED International Conference Parallel and Distributed Computing and Systems (PDCS)*, pp. 401–406, 2000.

[17] R. Rabenseifner, and G. Wellein, "Communication and optimization aspects of parallel programming models on hybrid architectures", *Internat. J. High Performance Computing Appl.*, vol. 17, no. 1, pp. 49–62, 2003.

[18] M. Ruiz, O. Lopera, and C. de la Plata, "Component-based derivation of a parallel stiff ODE solver", *Internat. J. Parallel Programming*, vol. 30, no. 2, pp. 99–148, 2002.

[19] Z. Tang, T. R. Coleman, and W. G. Dunphy, "Two distinct mechanisms for negative regulation of the Wee1 protein kinase", *EMBO J.*, vol. 12, no. 9, pp. 3427–36, 1993.

[20] L. T. Watson and C. A. Baker, "A fully-distributed parallel global search algorithm", *Engineering Computations*, vol. 18, no. 1/2, pp. 155–169, 2001.

[21] J. W. Zwolak, J. J. Tyson, and L. T. Watson, "Parameter estimation in a cell cycle model for frog egg extracts", in *High Performance Computing Symposium 2002*, A. Tentner (ed.), Soc. for Modeling and Simulation Internat., San Diego, CA, pp. 67–74 2002.