

**Deterministic Parallel Global Parameter Estimation
for a Model of the Budding Yeast Cell Cycle**

Thomas D. Panning

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Computer Science

Layne T. Watson
John J. Tyson
Clifford A. Shaffer

May 25, 2006
Blacksburg, VA

Keywords: DIRECT (DIviding RECTangles) algorithm, Direct search,
MADS (Mesh Adaptive Direct Search) algorithm, Computational biology

Deterministic Parallel Global Parameter Estimation for a Model of the Budding Yeast Cell Cycle

Thomas D. Panning

ABSTRACT

Two parallel deterministic direct search algorithms are combined to find improved parameters for a system of differential equations designed to simulate the cell cycle of budding yeast. Comparing the model simulation results to experimental data is difficult because most of the experimental data is qualitative rather than quantitative. An algorithm to convert simulation results to mutant phenotypes is presented. Vectors of the 143 parameters defining the differential equation model are rated by a discontinuous objective function. Parallel results on a 2200 processor supercomputer are presented for a global optimization algorithm, DIRECT, a local optimization algorithm, MADS, and a hybrid of the two. A second formulation is presented that uses a system of smooth inequalities to evaluate the phenotype of a mutant. Preliminary results of this formulation are given.

Acknowledgments

First, thanks to everyone on my committee for your time and advice. Thank you Dr. Watson for the guidance and honesty throughout my time as a graduate student. Thank you Dr. Tyson for giving me an appreciation of biological modeling. Thank you Dr. Shaffer for easing the transition from undergraduate to graduate student.

Thanks to my parents for all of their encouragement, support, and guidance. Thanks to my wife for her immeasurable support.

This work was partly supported by Defense Advanced Research Projects Agency (DARPA) grant F30602-02-0572.

Contents

1	Introduction	1
1.1	Observed and Predicted Phenotypes	3
1.1.1	Rules of Viability	3
1.1.2	Initial Conditions	4
2	Formulation as a Discontinuous Minimization Problem	5
2.1	The Objective Function	5
2.2	Algorithms	8
2.2.1	DIRECT	8
2.2.2	MADS	9
2.3	Results	14
2.4	Analysis	18
3	Formulation as a Smooth Feasibility Problem	21
3.1	Phenotype Constraints	21
3.1.1	Viable Phenotype Constraints	21
3.1.2	Phenotypes for <i>pds1Delta</i> Mutants	23
3.1.3	Inviability Phenotype Constraints	23
3.1.4	G1 Arrest	23
3.1.5	G2 Arrest	24
3.1.6	Metaphase Arrest	24
3.1.7	Telophase Arrest	24
3.1.8	Evaluating All of the Mutants	25
3.2	Evaluating the Constraints	25
3.3	Optimization	35
4	Conclusion	37
5	Bibliography	38
6	Appendix A: Mutant Phenotypes	39
	Vita	42

List of Figures

1	The five stages of the cell cycle	5
2	DIRECT on an example problem	10
3	MADS on an example problem	13
4	The objective function value versus the number of evaluations for MADS and DIRECT ...	15
5	NOMAD's performance when started from points found by DIRECT	16
6	The performance of MADS, DIRECT with $\epsilon = 0$, and DIRECT with $\epsilon = 0.1$	16
7	The distribution of points when MADS ran	17
8	The distribution of points when DIRECT ran with $\epsilon = 0$	17
9	The distribution of points when DIRECT ran with $\epsilon = 0.1$	18
10	The sensitivity of the model to local disturbances of the parameter vectors	20

List of Tables

1	Constants used in the objective function	7
2	Results for the runs that started from the best known point	16
3	The MADS runs that started from points found during the DIRECT run	18
4	The numbering of viable constraints	26
5	The numbering of telophase arrest constraints	27
6	Mutants that had no violated constraints, manually obtained parameters	31
7	Mutants that had violated constraints, manually obtained parameters	32
8	Mutants that had no violated constraints, optimized parameters	33
9	Mutants that had violated constraints, optimized parameters	34

1. Introduction

Molecular cell biology describes how cells convert genes into behavior. This description includes how a cell creates proteins from genes, how those proteins interact, and how networks of interacting proteins determine physiological characteristics of the cell. The central biological question addressed here is how protein interactions regulate the cell cycle of budding yeast (*Saccharomyces cerevisiae*).

The budding yeast cell cycle consists of four phases (G1, S, G2, M), with cell division occurring in the final phase. A newborn cell starts in G1 phase (unreplicated DNA), during which time it grows to a sufficiently large size to warrant a new round of DNA synthesis (S phase). After DNA synthesis has completed, the cell passes briefly through G2 phase (replicated DNA) and then enters M phase (mitosis, where the two copies of each DNA molecule are separated and the cell divides, creating two new cells that are in G1 phase).

The protein interactions that govern these cell cycle events are modeled using differential equations that describe the rate at which each protein concentration changes. In general, the concentration of protein A, written as $[A]$, changes according to the differential equation

$$\frac{d[A]}{dt} = \text{synthesis} - \text{degradation} - \text{binding} + \text{dissociation} - \text{inactivation} + \text{activation},$$

where “synthesis” is the rate at which new protein A molecules are synthesized from amino acids (which depends on the concentration of active messenger RNA molecules for a particular protein), “degradation” is the rate at which protein A is broken down into amino acids and polypeptide fragments (which depends on the activity of specific proteolytic enzymes), “binding” is the rate at which protein A combines with other molecules to form distinct molecular complexes, “dissociation” is the rate at which these complexes break apart, “inactivation” is the rate at which certain post-translational modifications (e.g., phosphorylation) of protein A are made, and “activation” is the rate at which these modifications are reversed (e.g., dephosphorylation). Each of these rates is itself a function of the concentrations of the interacting species in the network. For example,

$$\text{synthesis} = k_1[\text{transcription factor}],$$

$$\text{degradation} = k_2[\text{proteolytic enzyme}][A],$$

$$\text{binding} = k_3[A][B], \text{ where } B \text{ is a binding partner,}$$

$$\text{dissociation} = k_4[AB],$$

$$\text{inactivation} = \frac{k_5[\text{kinase}][A]}{J_5 + [A]},$$

$$\text{activation} = \frac{k_6[\text{phosphatase}][A_p]}{J_6 + [A_p]}, \text{ where } A_p \text{ is the phosphorylated form of } A.$$

In these rate laws, the ks are rate constants and the Js are Michaelis constants. Other differential equations must be used to determine the temporal dynamics of the concentrations of the “transcription factor,” “proteolytic enzyme,” “kinase,” etc.

The budding yeast cell cycle model [5] consists of 36 such differential equations for two classes of variables: regulatory proteins and physiological “flags.” The regulatory proteins are triggers for specific events of the budding yeast cell cycle: Cln2 triggers budding, Clb5 triggers DNA synthesis, Clb2 drives cells into mitosis, and Esp1 drives cells out of mitosis and back to G1. The physiological “flags” are dummy variables that track the strength of these trigger proteins. For example, “BUD” is an integral of the activity of Cln2; when BUD = 1, a new bud is initiated. “ORI,” an integral of [Clb5], represents the state of “origins of replication.” When ORI = 1 (“fired” origins), DNA synthesis is initiated; at cell division, when [Clb2]+[Clb5] drops below a threshold level, ORI is reset to zero (“licensed” origins). Finally, “SPN” represents the alignment of replicated chromosomes on the mitotic spindle. SPN is driven by Clb2 activity; i.e., SPN is an integral of [Clb2].

In the budding yeast model there are 143 rate constant parameters (ks , Js , etc.). In some cases, these parameters can be calculated directly from laboratory experiments (e.g., apparent protein half-lives), but most parameters are difficult to obtain directly from experimentation. Normally, modelers determine the remaining parameters by making educated guesses, solving the differential equations numerically, comparing the simulation results with laboratory data, and then refining their guesses. (Modelers call this process “parameter twiddling” [1].) For the budding yeast cell cycle, the laboratory data consists of observed phenotypes of more than 100 mutant yeast strains constructed by disabling and/or over-expressing the genes that encode the proteins of the regulatory network.

Although parameter twiddling is extremely tedious, it was used to obtain a parameter vector $(s_1, s_2, \dots, s_{143})$ for which the model’s predictions are consistent with almost all of the budding yeast mutants being modeled. Obviously, the modelers would prefer a method that allows them to spend more time working on the model and less time twiddling parameters. In addition, a person can only keep track of a few parameters at one time, which makes it easy for him or her to unwittingly miss a portion of the parameter space. For these reasons, modelers would prefer to use a tool that determines “good” parameters automatically, quickly, and accurately.

This thesis describes two formulations that reflect the biological criteria for an acceptable model. After a description of the phenotypes, Chapter 2 presents a discontinuous objective function and the results of optimizing that function. Chapter 3 describes a system of smooth constraints and some preliminary results. Conclusions are presented in Chapter 4.

1.1. Observed and Predicted Phenotypes

Throughout this thesis, the *observed* phenotype refers to the phenotype that was recorded in a laboratory experiment. The *predicted* phenotype refers to the phenotype that the mathematical model (with its associated parameters) predicts. The *wild type* is the normal strain of an organism. The *mutant* strains have genetic changes that make them behave differently from the wild type in some way.

Experimental biologists have studied many budding yeast mutants to learn about the cell cycle regulatory system. Of these mutants, 115 were chosen to model (see Appendix A). A model of budding yeast can be considered acceptable only if it is able to duplicate the behavior of most of these mutants. (It would be too much to expect a model to account for *all* the “observations” because of lingering uncertainties about the reaction network and inevitable mistakes in phenotyping mutants.) When the model is used to simulate a mutant, the parameter vector can be changed only in ways that are dictated by the genetic changes in the mutant. Consider the hypothetical proteins A and B presented in the previous section: if a mutant had a modified form of B that did not bind to A, then in the parameter vector for that mutant, k_3 would be set to zero and all the other parameters would be kept at the wild type values.

When comparing the model to the experimental data, it is important to realize that much of the data from laboratory experiments is qualitative. Such data is of the form “the cell is viable but considerably larger than wild type cells” or “the cell arrests in G1 phase and eventually dies.” The quantitative data that is available (e.g., duration of G1 phase, cell mass at division) is generally imprecise. With all these uncertainties, there may be many, clustered parameter vectors that allow the model to reproduce the experimental data sufficiently well. The proteins mentioned below (e.g., Clb2) are well known in the cell cycle literature [5]. What matters for the model here is the structure of the cell cycle regulatory pathways, not the details of the biochemistry.

1.1.1. Rules of Viability

To compare solutions of the differential equations with experimental data, it is necessary to predict cell cycle properties from a simulation of regulatory protein dynamics. Viability is determined by four rules:

1. The modeled cell must execute the following events in order, or else the modeled cell is considered inviable:
 - (a) DNA licensed for replication (modeled by a drop in $[\text{Clb2}] + [\text{Clb5}]$ below K_{ez2} , resetting $[\text{ORI}]$ to zero);

- (b) start of DNA synthesis (due to a subsequent rise in $[Clb2] + [Clb5]$, causing $[ORI]$ to increase above one), signaling the end of G1 phase, before a wild-type cell in the same medium would divide twice;
 - (c) alignment of DNA copies (due to a rise in $[Clb2]$, causing $[SPN]$ to increase above one) while $[Esp1]$ is less than 0.1;
 - (d) separation of DNA copies (modeled by $[Esp1]$ increasing above 0.1, due to Pds1 proteolysis at anaphase);
 - (e) cellular division (modeled by $[Clb2]$ dropping below a threshold K_{ez}), which resets $[BUD]$ and $[SPN]$ to zero.
2. The cell is inviable if division occurs in an “unbudded cell” (i.e., if $[BUD]$ does not reach the value 0.8 before event (e) occurs).
 3. The cell cycle should be stable, i.e., the squared relative differences of the masses and G1 phase durations in the last two cycles should both be less than 0.05.
 4. Lastly, the modeled cell is considered inviable if cell mass at division is greater than four times or less than one-fourth times the steady-state mass at division of the wild type in the same medium.

If the observed phenotype for a mutant does not complete one of the checkpoints (e.g., the mutant cells do not bud), then the predicted phenotype of that mutant must exhibit the same behavior. It is possible for a cell to complete all of the checkpoints in the first cycle and then become arrested somewhere in the second cycle. If a cell has this type of observed phenotype, then a correct model must predict the same number of cycles before arresting in the same manner. If the observed cell has a viable phenotype, then the model must predict a viable cell with a similar G1-phase length, and a similar mass at division.

1.1.2. Initial Conditions

In the experimental data set, many of the mutations are conditional, that is, the mutant cells when grown under “normal” conditions (say, glucose medium at room temperature) behave like wild-type cells, but when grown under “restrictive” conditions (say, galactose medium or elevated temperature) the cells express the genetic mutation and the aberrant phenotype. To model this situation at sample points in parameter space, start a “wild-type” simulation from arbitrary (but reasonable) initial conditions and integrate the differential equations for two full cycles, in order to wash out any effects of the initial conditions. Then record the state of the control system just after $[Clb2] + [Clb5]$ falls through K_{ez2} at the beginning of the third cycle. These recorded values are used as initial conditions for simulating a steady state wild-type cell and for simulating each of the mutants.

2. Formulation as a Discontinuous Minimization Problem

This formulation uses an objective function that rates the difference between two phenotypes. Section 2.1 presents this objective function and gives the mathematical definition of a “phenotype” as it is used in this chapter. Two deterministic algorithms, DIRECT and MADS, that are applicable to parameter estimation, are briefly described in Section 2.2. Numerical results on the supercomputer System X (a 2200 processor cluster) are given in Section 2.3. A short analysis of these results is given in Section 2.4. Parallel efficiency and scalability are important issues to be addressed separately—the emphasis here is on the biological problem, the discontinuous objective function formulation, and the practical applicability of DIRECT and MADS to such optimization problems.

The viability rules presented in the previous chapter are used by an algorithm [2] (called a *transform*) that outputs a phenotype from a solution of the differential equations. The transform keeps track of what stage the cell is in, where the stages are demarcated by the events (a-e) in the first rule of viability. The first stage is *unlicensed*, which ends when the first event, origin relicensing, occurs. The other four stages are, in chronological order, *licensed*, *fired*, *aligned*, and *separated*. When the simulated cell is in the *separated* stage, cellular division signals the transition back to the *unlicensed* stage. The relations among the stages, events, and biological phases of the cell are shown in Figure 1. If one of the rules of viability is broken, the transform sets an error flag and records the stage when the error occurred and the number of cycles (i.e., cell divisions) completed from the time when the mutation was expressed to the time when the cell arrested.




Figure 1. The five stages of the cell cycle.

2.1. The Objective Function

The objective function takes the observed phenotype and predicted phenotype for all of the mutants and computes a nonnegative score. Zero indicates a perfect match and larger numbers indicate increasingly worse matches. The ensuing discussion uses the symbol O for observed phenotype values and P for predicted phenotype values.

A budding yeast phenotype for a single mutant is represented by a six-tuple (v, g, m, a, t, c) , where the viability $v \in \{\text{viable}, \text{inviable}\}$, the real number $g > 0$ is the steady state length of the G1 phase in minutes, the real number $m > 0$ is the steady state mass at division expressed as a multiple of the wild type's steady state mass at division in the same medium (e.g., glucose or galactose), the stage when arrest occurs is

$$a \in A = \{\text{unlicensed}, \text{licensed}, \text{fired}, \text{aligned}, \text{separated}\},$$

the positive integer t is the arrest type (e.g., if events occur in improper order), and the nonnegative integer c is the number of successful cycles completed. Arrest types cannot be compared unless the stage of arrest is the same for both phenotypes. The observed and predicted phenotypes are written $O = (O_v, O_g, O_m, O_a, O_t, O_c)$ and $P = (P_v, P_g, P_m, P_a, P_t, P_c)$, respectively. Then

$$\mathcal{P} = \{(v, g, m, a, t, c)\} = \{\text{viable}, \text{inviable}\} \times (0, \infty)^2 \times A \times \{1, \dots, 10\} \times \mathcal{Z}_+$$

is the space of all budding yeast phenotypes, where \mathcal{Z}_+ denotes the nonnegative integers $\{0, 1, 2, \dots\}$.

The rating function, R , computes a non-negative real number that expresses the deviations between the observed and predicted phenotypes for a mutant. This rating function is a modified version of the one developed by N. Allen et al. [2]; the only difference is that if P_v is missing (if integration fails for some reason), then $R(O, P) = \omega_v$. The rating function is split into four cases depending on the viability of the observed and predicted phenotypes. If $O_v = \text{inviable}$, $P_v = \text{viable}$, and O_c is missing, then $R(O, P) = \omega_v$, the same as if $O_c = 0$. Otherwise, if a needed classifier is missing, the term is simply dropped and does not contribute to the objective function. In the case that classifiers are missing, this allows the objective function value to be at or near zero when viability is in agreement between the phenotypes, and forces larger objective function values when viability is not in agreement.

In what follows, the ω s and σ s are constants defined in Table 1. The rating function when all classifiers are present is given by

$$R(O, P) = \begin{cases} \omega_g \times \left(\frac{O_g - P_g}{\sigma_g} \right)^2 + \omega_m \times \left(\frac{\ln \frac{O_m}{P_m}}{\sigma_m} \right)^2, & \text{if } O_v = \text{viable} \text{ and } P_v = \text{viable}, \\ \omega_v \times \frac{1}{1 + P_c}, & \text{if } O_v = \text{viable} \text{ and } P_v = \text{inviable}, \\ \delta_{O,P} + \omega_c \times \left(\frac{O_c - P_c}{\sigma_c} \right)^2, & \text{if } O_v = \text{inviable} \text{ and } P_v = \text{inviable}, \\ \omega_v \times \frac{1}{1 + O_c}, & \text{if } O_v = \text{inviable} \text{ and } P_v = \text{viable}, \end{cases}$$

where δ is a real-valued discrete function used to assess a penalty for the arrest stage and type, given by

$$\delta_{O,P} = \begin{cases} \omega_a, & \text{if } O_a \neq P_a, \\ \omega_t, & \text{if } O_a = P_a \text{ and } O_t \neq P_t, \\ 0, & \text{if } O_a = P_a \text{ and } O_t = P_t. \end{cases}$$

The rating function is tuned by parameters that allow adjusting the relative importance of classifiers. The parameters given by Table 1 were set so that a rating of around ten indicates a critical error in the model's prediction of a phenotype.

Symbol	Definition	Value
ω_g	G1 length weight	1.0
σ_g	G1 length scale	10.0
ω_m	Mass at division weight	1.0
σ_m	Mass at division scale	$\ln 2$
ω_a	Arrest stage weight	10.0
ω_t	Arrest type weight	5.0
ω_c	Cycle count weight	10.0
σ_c	Cycle count scale	1.0
ω_v	Viability weight	40.0

Table 1. Constants used in the objective function.

Denote the real numbers by \mathcal{R} , the integers by \mathcal{Z} , the number of parameters by N_p , and the number of mutants by N_m . Let the domain of the objective function be the box

$$\Omega = \{x \in \mathcal{R}^{N_p} : \ln(s_i/u_i) \leq x_i \leq \ln(s_i \times u_i), i = 1, \dots, N_p\},$$

where $u \in \mathcal{R}^{N_p}$ is a vector of positive scale factors (most components of u are approximately 100) reflecting *a priori* limits on the rate constants, and $s \in \mathcal{R}^{N_p}$ is the modeler's best guess point. For any $x \in \Omega$, the simulated phenotype of mutant $j \in \{1, 2, \dots, N_m\}$ is $P_j(x) \in \mathcal{P}$. The objective function is the weighted sum of the rating function over all of the mutants, that is,

$$f(x) = \sum_{j=1}^{N_m} \mu_j R(O_j, P_j(x)),$$

where $\mu_i > 0$ indicates the relative importance of the i th mutant. The objective function value at the best previously known point [5] is 470.

2.2. Algorithms

This section describes two algorithms that show promise for optimizing the discontinuous objective function described in the previous section. Consider the problem of minimizing $f : \Omega \rightarrow \mathcal{R}$, where $\Omega = [l, u] \subset \mathcal{R}^n$ is a box.

2.2.1. DIRECT

The DIRECT (Dividing Rectangles) global minimization algorithm [13] requires the objective function to be Lipschitz continuous to guarantee convergence. Even though the objective function used here is discontinuous, the DIRECT algorithm seems to be an efficient and reasonable deterministic sampling strategy worth trying.

The DIRECT algorithm is one of a class of deterministic direct search algorithms that does not require gradients. It works by iteratively dividing the search domain into boxes that have exactly one function value at the box's center. In each iteration, the algorithm determines which boxes are most likely to contain a better point than the current minimum point—these boxes are called “potentially optimal”. It then subdivides the potentially optimal boxes along their longest dimensions. Intuitively, a box is considered potentially optimal if it has the potentially best function value for a given Lipschitz constant. The formal definition from [13] follows.

DEFINITION 1. Suppose that the unit hypercube has been partitioned into m (hyper) boxes. Let c_i denote the center point of the i th box, and let d_i denote the distance from the center point to the vertices. Let $\epsilon \geq 0$ be a positive constant. A box j is said to be *potentially optimal* if there exists some $\tilde{K} > 0$ such that for all $i = 1, \dots, m$,

$$\begin{aligned} f(c_j) - \tilde{K}d_j &\leq f(c_i) - \tilde{K}d_i, & \text{for all } i = 1, \dots, m, \\ f(c_j) - \tilde{K}d_j &\leq f_{\min} - \epsilon|f_{\min}|. \end{aligned}$$

The DIRECT algorithm is described by the following six steps [12].

Step 1. Normalize the design space Ω 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 i th 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.

For an illustration of how the DIRECT algorithm searches the domain on an example problem, see Figure 2. Both serial [12] and parallel [9] versions of DIRECT have been described in the literature.

2.2.2. MADS

A MADS (Mesh Adaptive Direct Search) algorithm, as defined by Audet and Dennis [4], minimizes a nonsmooth function $f : \mathcal{R}^n \rightarrow \mathcal{R} \cup \{+\infty\}$ under general constraints $x \in \Omega \subseteq \mathcal{R}^n$, $\Omega \neq \emptyset$. If $\Omega \neq \mathcal{R}^n$, the algorithm works with f_Ω , which is equal to f on Ω and $+\infty$ outside Ω . Using f_Ω in lieu of f is called a “barrier” approach to handling arbitrary constraints $x \in \Omega$.

In each iteration, a MADS algorithm evaluates the objective function f_Ω at a finite number of trial points. Central to these algorithms is the concept of a mesh, which is a discrete set of points in \mathcal{R}^n . Every previous trial point must lie on the current mesh, and in each iteration the algorithm may only generate new trial points on the current mesh. This is not as restrictive as it might sound because the algorithm changes the mesh after each iteration (with the restriction that all previously evaluated points remain in the new mesh).

To further define the mesh, three entities— Δ_k^m , D , S_k —must be introduced. First, the *mesh size parameter* $\Delta_k^m > 0$ controls the granularity of the mesh at iteration k ; after the k th iteration, Δ_{k+1}^m is adjusted from Δ_k^m depending on the success of that iteration. The second entity is an $n \times n_D$ matrix D , where each column $D_{\cdot j} = Gz_j$ (for $j = 1, 2, \dots, n_D$) for some fixed nonsingular generating matrix $G \in \mathcal{R}^{n \times n}$ and nonzero integer vector $z_j \in \mathcal{Z}^n$. The columns of D must also be a positive spanning set, $\text{Pos}(D) = \mathcal{R}^n$ (i.e., the cone generated by nonnegative combinations of columns of D spans \mathcal{R}^n). Lastly, S_k is the set of points where the objective function has been

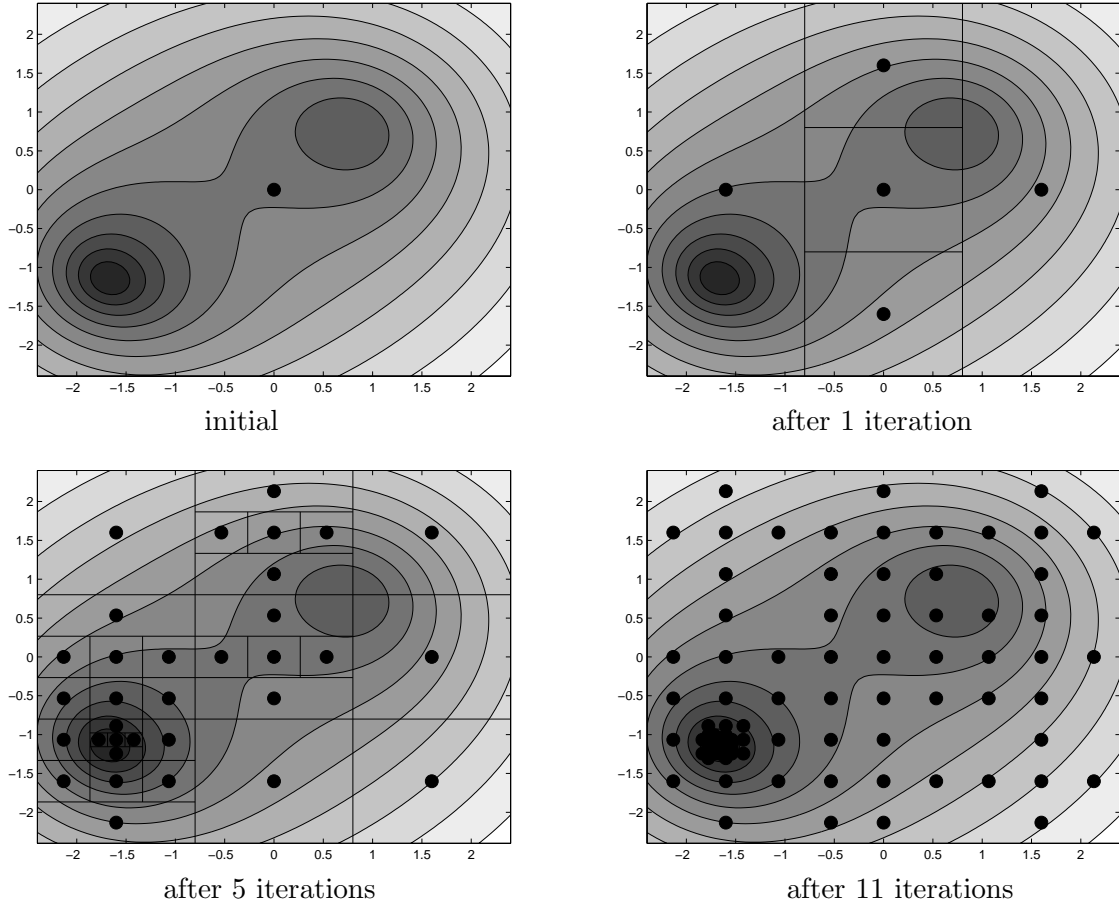


Figure 2. These graphs show the function evaluations that DIRECT performed after zero, one, five, and eleven iterations. Comparing the first and second graphs shows how DIRECT divides a two-dimensional box. The second and third graphs include the rectangles that DIRECT had created. After five iterations, DIRECT has found the global minimum at $(-1.8, -1.1)$. After the fifth iteration DIRECT has explored the domain, subdividing most of the larger boxes. After eleven iterations, DIRECT has evaluated the function at points near the local minimum.

evaluated by the start of iteration k . Now that those entities have been introduced, the current mesh can be precisely defined.

DEFINITION 2. At iteration k , the current *mesh* is defined to be

$$M_k = \bigcup_{x \in S_k} \{x + \Delta_k^m Dz : z \in \mathcal{N}^{n_D}\}.$$

This definition ensures that all previously evaluated points are included in the mesh. It also shows that a smaller Δ_k^m will result in a more refined mesh, while a larger Δ_k^m will create a coarser mesh.

Now that the mesh has been defined, the iterations of a MADS algorithm can be described. Each iteration consists of two steps: the SEARCH step and the POLL step. The SEARCH step may

evaluate f_Ω at any finite number of mesh points. At which mesh points f_Ω is evaluated depends on the precise MADS algorithm in use. A MADS algorithm may even do zero evaluations in the SEARCH step; the SEARCH step is said to be empty when no points are considered. If the SEARCH step fails to find a mesh point at which f_Ω is less than $\min_{x \in S_k} f_\Omega(x)$, then the algorithm performs the POLL step by generating and evaluating f_Ω at new trial points around the current incumbent solution x_k , where $f_\Omega(x_k) = \min_{x \in S_k} f_\Omega(x)$. The *poll size parameter* Δ_k^p limits the distance between x_k and the new trial points. The set of new trial points is called a *frame*, and x_k is called the *frame center*. The MADS frame is constructed using x_k , Δ_k^p , Δ_k^m , and D to obtain a set D_k of positive spanning directions.

DEFINITION 3. At iteration k , the MADS *frame* is defined to be the set

$$P_k = \{x_k + \Delta_k^m d : d \in D_k\} \subset M_k,$$

where D_k is a positive spanning set such that $0 \notin D_k$ and for each $d \in D_k$,

- d can be written as a nonnegative integer combination of the columns of D : $d = Du$ for some vector $u \in \mathcal{N}^{n_D}$,
- the distance from the frame center x_k to a frame point $x_k + \Delta_k^m d \in P_k$ is bounded by a constant times the poll size parameter: $\Delta_k^m \|d\|_\infty \leq \Delta_k^p \|D\|_\infty$,
- limits (as defined in Coope and Price [6]) of the normalized sets D_k are positive spanning sets.

The algorithm evaluates f_Ω at points in the frame P_k until it encounters an improved point x^* ($f_\Omega(x^*) < f_\Omega(x_k)$) or it has evaluated f_Ω at all of the points in P_k .

After the algorithm has executed the SEARCH step and (conditionally) the POLL step, it sets the mesh size and poll size parameters, Δ_{k+1}^m and Δ_{k+1}^p , for the next iteration. If the iteration successfully found a better mesh point x_{k+1} such that $f_\Omega(x_{k+1}) < f_\Omega(x_k)$, then Δ_{k+1}^m will be larger than or equal to Δ_k^m ; otherwise, Δ_{k+1}^m will be smaller than Δ_k^m . The poll size parameter Δ_{k+1}^p must be set such that $\Delta_{k+1}^m \leq \Delta_{k+1}^p$, and it must satisfy

$$\liminf_{k \rightarrow \infty} \Delta_k^m = 0 \iff \liminf_{k \rightarrow \infty} \Delta_k^p = 0.$$

Exactly how Δ_{k+1}^m and Δ_{k+1}^p are generated is determined by the individual algorithm in use. A typical algorithm uses the following rules to set the mesh size parameter: $\Delta_0^m = 1$, and

$$\Delta_{k+1}^m = \begin{cases} \Delta_k^m/4, & \text{if } x_k \text{ is a minimizing} \\ & \text{frame center,} \\ 4\Delta_k^m, & \text{if an improved mesh} \\ & \text{point is found, and} \\ & \text{if } \Delta_k^m \leq \frac{1}{4}, \\ \Delta_k^m, & \text{otherwise.} \end{cases}$$

The same example algorithm then uses the simple rule $\Delta_k^p = \sqrt{\Delta_k^m}$ to determine the value of Δ_{k+1}^p . These rules ensure that Δ_k^m is always a power of 1/4 less than or equal to one, and Δ_k^m is always less than or equal to Δ_k^p .

In summary, the MADS class of algorithms is described by the following five steps.

Step 1. Let $x_0 \in \Omega$ and $0 < \Delta_0^m \leq \Delta_0^p$. Let D be an $n \times n_D$ matrix with the properties described earlier. Set the iteration counter $k := 0$.

Step 2. Perform the SEARCH step. This step varies among the individual algorithms; in all algorithms f_Ω is evaluated at a finite subset of points (called trial points) on the mesh M_k . If a trial point y is found such that $f_\Omega(y) < f_\Omega(x_k)$, then the algorithm may go to Step 4 with $x_{k+1} := y$.

Step 3. Perform the POLL step, evaluating f_Ω at points from the frame $P_k \subset M_k$ until a frame point x_{k+1} is found with $f_\Omega(x_{k+1}) < f_\Omega(x_k)$ or f_Ω has been evaluated at all of the points in P_k .

Step 4. Update Δ_{k+1}^m and Δ_{k+1}^p according to the specific algorithm's rules. In all algorithms,

- (1) Δ_{k+1}^m is greater than or equal to Δ_k^m if an improved mesh point is found,
- (2) Δ_{k+1}^m is less than Δ_k^m if an improved mesh point is not found,
- (3) Δ_{k+1}^p is greater than or equal to Δ_{k+1}^m , and
- (4) $\liminf_{j \rightarrow \infty} \Delta_j^m = 0$ if and only if $\liminf_{j \rightarrow \infty} \Delta_j^p = 0$.

Step 5. If an appropriate stopping criterion has been met, stop. Otherwise, set $k := k + 1$ and go back to Step 2.

The previous discussion presents the MADS class of algorithms. The following discussion describes a specific instance of the class, and Figure 3 shows how that algorithm behaves on an example problem. To emphasize the POLL step of the algorithm, there is no SEARCH step in the algorithm presented here.

In this MADS algorithm,

$$D = \begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -2 \end{pmatrix}.$$

Notice that a MADS mesh constructed using this matrix is identical to a mesh constructed using the matrix

$$B = \begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix}.$$

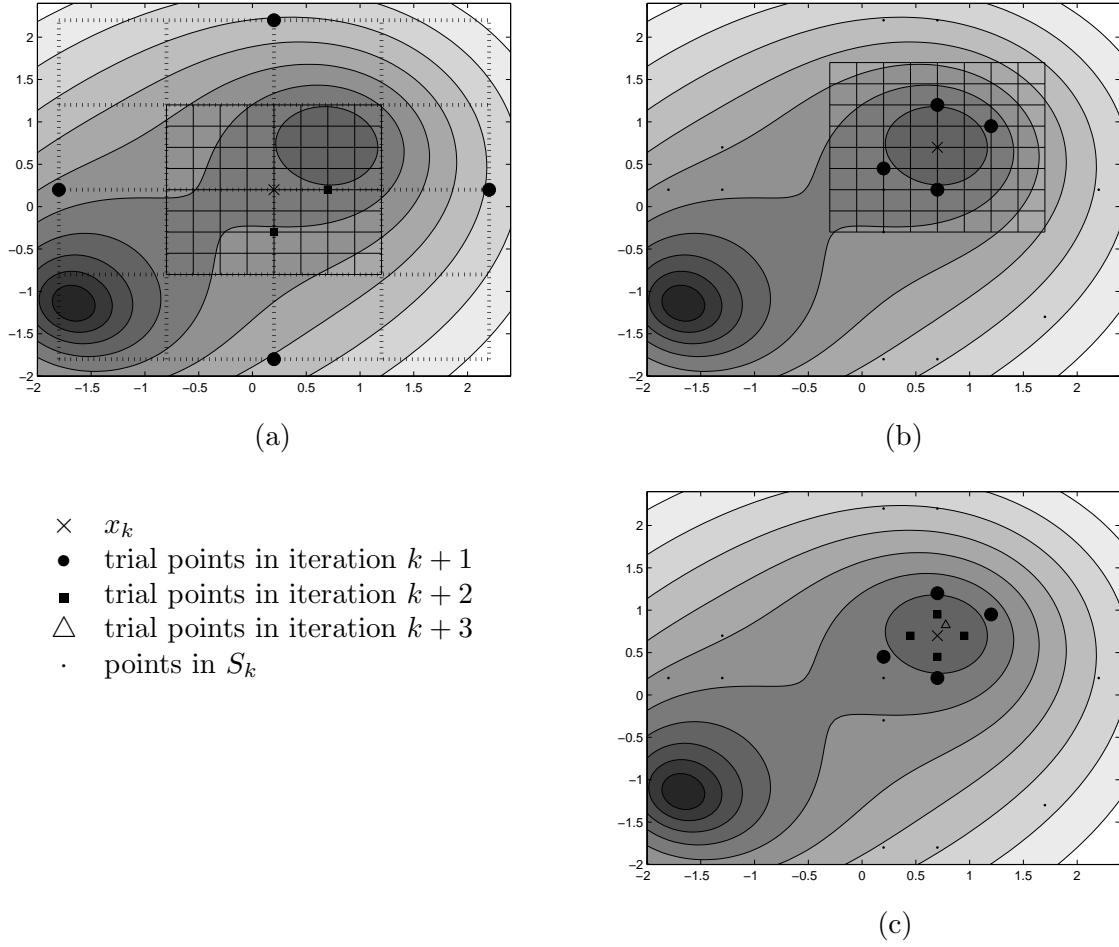


Figure 3. These three graphs show how a MADS algorithm can refine the mesh, choose different poll directions, and contract the search area. In Figure 3(a), the intersections of the dotted lines indicate points that met the criteria for frame points in the first iteration. From these possibilities, the algorithm chose the four points indicated by the large circles; these four points constitute P_1 . At all of these points, the function is higher than at x_0 , so the algorithm refined the mesh by setting $\Delta_2^m = \Delta_1^m/4$. The intersections of the solid lines in the same graph indicate possible frame points for the second iteration. The algorithm evaluated the function at only two points in P_2 because the function is lower at the second point than at x_0 . Figure 3(b) shows how the mesh allows a MADS algorithm to choose different poll directions at each iteration. Figure 3(c) highlights three consecutive iterations of the algorithm. In the first two iterations, the algorithm is unable to find an improved mesh point, therefore it restricts the search area to be closer to x_k . This can be seen by looking at the points evaluated in the three iterations; the circles are the furthest away from x_k , the squares are closer, and the triangle is the closest. There is only one triangle because the function value at that point is lower than $f(x_k)$, so the algorithm stopped the POLL step and went on to the next iteration.

However, $\|D\|_\infty = 2$ while $\|B\|_\infty = 1$; thus, a MADS frame constructed using D instead of B will extend twice as far in every direction. From D , the matrix D_k is generated (using random

coefficients as described in [4]) at the beginning of the k th iteration so that it is a positive spanning set, and so that the (normalized) columns of D_i , for $i = 1, 2, \dots$, are dense in the unit circle \mathcal{S}^1 . The mesh size parameter Δ^m and poll size parameter Δ^p are updated using the same rules as presented earlier.

The example function in Figure 3 has two local minima in the chosen domain. The basin of attraction for the global minimum point is smaller than that of the other local minimum point. For the example problem, $\Omega = \{(x_1, x_2) : x_1 \in (-2, 2.4), x_2 \in (-2, 2.4)\}$.

2.3. Results

Because it takes approximately 17 seconds to evaluate the objective function on a 2.3 GHz PowerPC G5 processor, parallel versions of both algorithms were used. All computation took place on a cluster of 1,100 dual-processor Mac G5 nodes, called System X, at Virginia Tech.

NOMAD [4] is a C++ implementation of the MADS class of algorithms. The specific MADS algorithm used here is the same as the example algorithm given in Section 4, with $\Delta_0^m = 1/4$. To take advantage of System X, NOMAD’s implementation of the POLL step was parallelized using a master/worker paradigm. The master ran the MADS algorithm as presented in Section 4 and sent requests to the workers whenever objective function values were needed. NOMAD, started from the modeler’s best point, evaluated the objective function 135,000 times over 813 iterations using 128 processors, converging at a point for which the objective function value was 299.

pVTDirect [9] is a parallel implementation of DIRECT written in Fortran 95. While the DIRECT algorithm does not have a traditional “starting point”, the first sample in each subdomain is always taken at the center of the subdomain bounding box. For this problem, the bounding box was designed so that the modeler’s best point [5] would be at the center and therefore would be evaluated before any other points. pVTDirect (with only one subdomain and $\epsilon = 0$) ran for 473 iterations using 1024 processors and evaluated the objective function 1.5 million times, finding a point at which the objective function value was 212.

The first set of runs (the results of which are illustrated in Figures 4 and 5) use the same initial conditions simulating all of the points, not different initial conditions as described in Section 2.2. (These runs led to the discovery of the necessity of the procedure in 2.2.) Figure 4 shows the progress that each algorithm was able to make in minimizing the objective function. While NOMAD was able to quickly find a better point than the modeler’s best point, pVTDirect eventually found an even lower point. In a later run, NOMAD was started from pVTDirect’s lowest




Figure 4. The objective function value at the best point found versus the number of evaluations for MADS and DIRECT. (The computations on which Figures 4 and 5 are based used a standard set of initial conditions for every simulation, not the more accurate updating of initial conditions described in Section 2.2.)

point, but NOMAD was unable to make any further progress. After looking at Figure 4, it is tempting to believe that pVTDirect could have been stopped earlier (for instance, after 200,000 evaluations), and NOMAD started at pVTDirect’s last best point could have found a point at which the objective function value was 212 or less. To test this, NOMAD was started at the best point at the 54th, 157th, and 239th iterations of pVTDirect. These points correspond to the beginning, middle, and end of the second-lowest plateau in Figure 4. As shown in Figure 5, NOMAD started from the middle point converged to a point at which the objective function value was 210. However, the NOMAD runs started at the beginning and end plateau points converge to worse points than pVTDirect’s best point. These four extra NOMAD runs (including the one starting from pVTDirect’s best point) show that using NOMAD to get more value out of a pVTDirect run is not simple.

When the initial conditions were generated (as described in Section 2.2) for the best points from the previously described runs, the points received considerably worse objective function values, leading to the realization that the initial conditions have to be chosen per Section 2.2 for each parameter vector. To remedy this, the objective function was modified to generate proper initial conditions (see Section 2.2) for each point before evaluating the point. NOMAD and pVTDirect were then rerun on the corrected objective function. For this set of runs, pVTDirect was run twice, once with $\epsilon = 0$ and again with $\epsilon = 0.1$. Table 2 gives basic information for each method, and the results of these three runs are shown in Figure 6.

Figures 7, 8, and 9 give an idea of what areas of parameter space the different runs explored. The distance measure used for these figures was $\sum_{i=1}^{N_p} |s_i - p_i|/s_i$, where s is the modeler’s best




Figure 5. The performance of NOMAD when started from the best point at pVTDirect's 54th, 157th, and 239th iterations. The plots are shown as if the NOMAD runs started as soon as the respective pVTDirect iterations completed.

Run	lowest f	# evaluations	# CPUs	CPU hours
MADS	325	77,221	64	384
DIRECT, $\epsilon = 0$	233	1,243,429	400	8,895
DIRECT, $\epsilon = 0.1$	365	1,452,597	400	10,703

Table 2. Results for the runs that started from the best known point [5].




Figure 6. The objective function value at the best point found so far versus the number of evaluations for MADS, DIRECT with $\epsilon = 0$, and DIRECT with $\epsilon = 0.1$.

parameter vector and p is the parameter vector of the point being plotted. These figures confirm that DIRECT evaluates points further away from the starting point than MADS. What is not shown in Figures 8 and 9 is that DIRECT found only 77,752 points that evaluated to less than 480 when ϵ was set to 0.1, but it found 565,982 such points when ϵ was set to 0, even though both runs evaluated approximately the same number of points. This, combined with the information shown in Figures 8 and 9, shows that setting ϵ to 0.1 caused DIRECT to spend more time dividing large




Figure 7. The distribution of points that evaluated to less than 480 when MADS ran.




Figure 8. The distribution of points that evaluated to less than 480 when DIRECT ran with $\epsilon = 0$.

boxes and less time refining small boxes.

However, Figure 5 showed that MADS may be able to quickly improve on the best points that DIRECT has found. Unfortunately, it is difficult to guess which points from DIRECT will be good starting points for MADS. To find points that scored well but were reasonably far apart, the following algorithm was used. Define the closed ball $B(x, \delta) = \{y \in \mathcal{R}^n \mid \|y - x\| \leq \delta\}$, and denote the complement of the set S by $\mathcal{C}(S)$.

Step 1. Let P be all of the points that DIRECT evaluated when ϵ was set to 0. Set $k := 1$.

Step 2. Find $s_k \in P$ such that $f(s_k) \leq f(p)$ for all $p \in P \cap \mathcal{C}\left(\bigcup_{i=1}^{k-1} B(s_i, 2)\right)$.

Step 3. If $f(s_k) \geq 300$, then stop, else set $k := k + 1$ and go to Step 2.

This algorithm yielded 13 points to use as starting points for MADS runs. The results of those runs are given in Table 3. In this table, there does not seem to be any relationship between




Figure 9. The distribution of points that evaluated to less than 480 when DIRECT ran with $\epsilon = 0.1$.

$f(s_i)$ and the objective function value at the point to which MADS converges. It is also interesting that for 12 of the runs, MADS converges to a point that evaluates to 220–240, but the third run converges to a point that evaluates to a much lower score of 189. This is the lowest score found by any of the runs.

i	$f(s_i)$	final f	# evaluations	CPU hours
1	233	233	1,590	12
2	244	234	22,153	105
3	244	189	76,978	384
4	249	226	39,990	189
5	250	230	82,551	384
6	250	231	79,869	385
7	255	224	59,884	266
8	257	220	80,361	385
9	260	228	86,251	384
10	262	233	81,431	385
11	291	228	79,110	384
12	292	244	53,257	238
13	293	238	84,035	384

Table 3. The MADS runs that started from points found during the DIRECT (with $\epsilon = 0$) run. All MADS runs used 64 processors.

2.4. Analysis

Even with a discontinuous objective function and a 143-dimension search domain, both DIRECT and MADS performed well. When ϵ was set to zero, DIRECT explored the parameter space and refined the boxes near local minima. When ϵ was set to 0.1, DIRECT used most of its evaluations to explore the parameter space. Any inferences about the choice of ϵ must be made

with caution, since the box center here was already a very good point, and sizeable volumes of the 143-dimensional space still remain unsampled. MADS was almost always able to find a better point than its starting point, and it did so with far fewer evaluations than DIRECT. Using the two algorithms together yielded the lowest-scoring point.

The best previously known parameter vector from Chen et al. [5] has an objective function value of 470, and correctly models all of the mutants except (numbers from Appendix A) 7, 16, 35, 41, 45, 53, 76, 93, 97, 103, 104, and 110. The best point from DIRECT/MADS has an objective function value of 189, and correctly models all but mutants 35, 41, 45, 53, 76, 93, and 110.

How sensitive is the biological model (as reflected in the objective function) to local disturbances of these parameter vectors? If the DIRECT/MADS parameter values were rounded to two significant figures, would the model be such a good fit to the data? To investigate this question, random perturbations (up to $\pm 0.5\%$) were applied to all of the parameter values in both the Chen et al. parameter vector and the best DIRECT/MADS parameter vector, and histograms of the resulting objective function values were computed (Figure 10). For the Chen et al. parameter vector, about 25% of perturbations (white bars in Figure 10) give significantly worse fits to the data, a reflection of the fact that behavior of the model is quite sensitive to a small number of the parameters, as described in detail in [5]. The DIRECT/MADS parameter vector (grey bars in Figure 10) appears to be not only better but also its perturbations are less likely to produce inferior models than for the Chen et al. parameter vector. Thus the combination of DIRECT and MADS to perform parameter optimization on discontinuous objective functions in very high dimensional (> 100) parameter spaces is not only feasible computationally but also can find “good” parameter vectors that improve on the best estimates of expert modelers.




Figure 10. The best DIRECT/MADS parameter vector and the parameter vector from Chen et al. [5] were uniformly randomly perturbed by up to $\pm 0.5\%$ in all dimensions, and the objective function was evaluated at each of the perturbed vectors. This figure shows the percentage of vectors that scored in $[x, x + 100)$, for $x = 200, 300, \dots, 700$.

3. Formulation as a Smooth Feasibility Problem

The previous chapter uses an objective function that assigns a penalty to every discrepancy between the observed and predicted phenotypes. Simple bound constrained minimization is applied to the penalty function. Because of the nature of the data, some of the penalties are inherently discontinuous, while the ODE solution is a smooth function of the ODE parameters. Beyond the fact that discontinuous objective functions are difficult to minimize, either locally or globally, this penalty approach has deeper flaws. Biologists do not agree on what the penalty should be for a particular discrepancy, or even on which discrepancies should be penalized.

This chapter takes a quite different approach to parameter estimation. The idea is to describe the mutant data as a system of smooth nonlinear inequalities derived from the (smooth) ODE model output and cell biology knowledge. These inequalities should be generally accepted by cell biologists, even though some of the details may be debatable. An ODE parameter vector that satisfies all of the inequalities thus defines an acceptable model, and is a *feasible* solution of the system of inequalities. The proposed approach to parameter estimation is to solve a feasibility problem defined by a system of smooth (continuous, piecewise C^∞) inequalities. One could surmise that what really is desired is the “most interior” point, defined by, e.g., constraint margins. The situation is not so simple, though, since a cell’s viability is robust with respect to environmental variations, and therefore it is really the “most robust” feasible point that is sought. Modeling such biological robustness is another research topic.

3.1. Phenotype Constraints

For each phenotype, this section presents a system of constraints that will be satisfied if and only if the simulation predicts the same phenotype. The constraints are written using as much biological notation as possible, so $[\text{Esp1}](t_6)$ refers to the concentration of the protein Esp1 at the time t_6 . The constraints include time variables t_i where i is a positive integer; although these variables are not part of the model, they must still be found by the parameter estimator. In the constraints, there are several constants which are set as follows: $\epsilon = 0.05$, $\delta_{G_1} = 20$, and $\delta_M = 1.4$.

3.1.1. Viable Phenotype Constraints

The constraints for a viable cell to go through one cycle are listed below, with annotations denoted by angle brackets.

$$t_1 < t_2 < t_3 < t_4 < t_5 < t_6 < t_8, \tag{1}$$

$$t_3 - t_1 < t_w, \quad \langle 2 \rangle$$

$$t_2 < t_7 < t_8, \quad \langle 3 \rangle$$

$$G_1 - \delta_{G_1} < t_3 - t_0 < G_1 + \delta_{G_1}, \quad \langle 4 \rangle$$

$$\min_{t_1 \leq t \leq t_2} ([\text{Clb2}](t) + [\text{Clb5}](t)) > K_{\text{ez2}} > [\text{Clb2}](t_2 + \epsilon) + [\text{Clb5}](t_2 + \epsilon), \quad \langle 5 \rangle$$

$$\max_{t_2 + \epsilon \leq t \leq t_3} [\text{ORI}](t) < 1 < [\text{ORI}](t_3 + \epsilon), \quad \langle 6 \rangle$$

$$\max_{t_1 \leq t \leq t_4} [\text{SPN}](t) < 1 < [\text{SPN}](t_5), \quad \langle 7 \rangle$$

$$\max_{t_4 \leq t \leq t_5} [\text{Esp1}](t) < 0.1 < [\text{Esp1}](t_6), \quad \langle 8 \rangle$$

$$[\text{BUD}](t_7) > 0.8, \quad \langle 9 \rangle$$

$$[\text{Clb2}](t_8 - \epsilon) > K_{\text{ez}} > [\text{Clb2}](t_8), \quad \langle 10 \rangle$$

$$\max_{t_1 \leq t \leq t_8 - \epsilon} [\text{mass}](t) < 4m_w, \quad \langle 11 \rangle$$

$$M - \delta_M < [\text{mass}](t_8 - \epsilon) < M + \delta_M. \quad \langle 12 \rangle$$

$\langle 1 \rangle$ These strict inequalities ensure the correct temporal ordering of the events defined by the times t_i . $\langle 2 \rangle$ $[\text{ORI}]$ must rise above one before a wild type cell would divide twice in the same medium (e.g., glucose or galactose); t_w is set to the amount of time a simulated wild type cell takes to divide twice with the same biological parameters. $\langle 3 \rangle$ t_7 (which marks $[\text{BUD}]$ rising above one) simply has to occur any time between t_2 and t_8 . $\langle 4 \rangle$ G_1 is the length of the G1-phase, as observed in experiments. This ensures that the simulated cell is within a predefined distance δ_{G_1} of the observed value. $\langle 5 \rangle$ $[\text{Clb2}] + [\text{Clb5}]$ drops, satisfying viability rule 1(a). $\langle 6 \rangle$ $[\text{ORI}]$ rises, satisfying rule 1(b). $\langle 7 \rangle$ $[\text{SPN}]$ rises, satisfying rule 1(c). $\langle 8 \rangle$ $[\text{Esp1}]$ rises, satisfying rule 1(d). $\langle 9 \rangle$ $[\text{BUD}]$ rises, satisfying rule 2. $\langle 10 \rangle$ $[\text{Clb2}]$ drops, satisfying rule 1(e). $\langle 11 \rangle$ Mass is always less than four times the mass of the wild type in the same medium; m_w is the mass of a simulated wild type cell with the same biological parameters. $\langle 12 \rangle$ M is the observed mass at division.

A cell that meets the above constraints is viable for one cycle. For the first cycle, remember that the starting conditions are just after $[\text{Clb2}] + [\text{Clb5}]$ dropped through K_{ez2} , so the fifth constraint must be omitted. For the rest of the cycles, repeat the constraints with variables $t_{1+7(n-1)}$, $t_{2+7(n-1)}$, \dots , $t_{8+7(n-1)}$ for cycle n . Note that the last time of the previous cycle is the first time in the current cycle. To enforce the stability requirements, add the constraints

$$\left(\frac{[\text{mass}](t_{8+7(N-2)}) - [\text{mass}](t_{8+7(N-1)})}{[\text{mass}](t_{8+7(N-2)})} \right)^2 < 0.05$$

and

$$\left(\frac{(t_{3+7(N-2)} - t_{1+7(N-2)}) - (t_{3+7(N-1)} - t_{1+7(N-1)})}{(t_{3+7(N-2)} - t_{1+7(N-2)})} \right)^2 < 0.05,$$

where N is the total number of cycles.

3.1.2. Phenotypes for $pds1\Delta$ Mutants

The $pds1\Delta$ mutants are incapable of synthesizing Esp1, but some of these mutants still manage to separate the DNA copies. It is suspected that these mutants use a different mechanism to separate the DNA, but that mechanism is not included in this model. So for the purposes of this model, $pds1\Delta$ mutants are not required to meet viability rule 1(d), and the corresponding constraints should be omitted when evaluating mutants 61, 62, 66, 67, 71, and 113 in Appendix A.

3.1.3. Inviability Phenotype Constraints

A cell that fulfills the above constraints is considered viable. Some of the mutants have an observed phenotype of inviable, so their constraints will be different. The constraints for an inviable mutant are determined by when and how the cell arrests. There are four major types of arrest stages: G1 arrest, metaphase arrest, G2 arrest, and telophase arrest. The following subsections present the inequality constraints for each of these arrest types.

3.1.4. G1 Arrest

In terms of the model, a G1 arrest means that none of [ORI], [SPN], or [BUD] should rise to one before the cell cycle is considered arrested. Whether or not $[Clb2] + [Clb5]$ drops below K_{ez} has no effect on whether the cell is G1 arrested, so it is not mentioned at all in the constraints. A cell arrests in G1 either because its mass has become greater than $4m_w$ (see viability rule 4) or because t_w time has passed (see checkpoint 1(b)). Constraints for a G1 arrested mutant are thus

$$\begin{aligned} \max\{t_1 - t_w, \max_{0 \leq t \leq t_1} [\text{mass}](t) - 4m_w\} &> 0, \\ \max_{0 \leq t \leq t_1} [\text{ORI}](t) &< 1, \\ \max_{0 \leq t \leq t_1} [\text{SPN}](t) &< 1, \\ \max_{0 \leq t \leq t_1} [\text{BUD}](t) &< 1. \end{aligned}$$

The first inequality ensures that t_1 is a time after the cell has arrested. Specifically, the quantity $t_1 - t_w$ will be greater than zero if a wild type cell could divide twice before t_1 . Similarly, the quantity $\max_{0 \leq t \leq t_1} [\text{mass}](t) - 4m_w$ will be greater than zero if the cell has grown to a mass greater

than $4m_w$ before t_1 . The maximum of these quantities is used because violating any of the viability rules causes a cell to be arrested.

3.1.5. G2 Arrest

For a cell to be arrested in G2, it must execute the first two checkpoints of a viable cell, but [SPN] must stay low.

$$\begin{aligned}
t_1 &< t_2 < t_3 < t_4, \\
t_3 - t_1 &< t_w, \\
\min_{t_1 \leq t \leq t_2} ([Clb2](t) + [Clb5](t)) &> K_{ez2} > [Clb2](t_2 + \epsilon) + [Clb5](t_2 + \epsilon), \\
\max_{t_2 + \epsilon \leq t \leq t_3} [\text{ORI}](t) &< 1 < [\text{ORI}](t_3 + \epsilon), \\
\max_{0 \leq t \leq t_4} [\text{SPN}](t) &< 1, \\
\max_{0 \leq t \leq t_4} [\text{mass}](t) &> 4m_w.
\end{aligned}$$

3.1.6. Metaphase Arrest

If a cell is arrested in metaphase, its chromosomes are aligned on its spindles (i.e., the [SPN] checkpoint must be reached) but the chromosomes have not separated (i.e., Esp1 has not activated). Metaphase arrested cells may be budded or unbudded. This means that the following constraints must be met.

$$\begin{aligned}
t_1 &< t_2 < t_3 < t_4 < t_5 < t_6, \\
t_3 - t_1 &< t_w, \\
\min_{t_1 \leq t \leq t_2} ([Clb2](t) + [Clb5](t)) &> K_{ez2} > [Clb2](t_2 + \epsilon) + [Clb5](t_2 + \epsilon), \\
\max_{t_2 + \epsilon \leq t \leq t_3} [\text{ORI}](t) &< 1 < [\text{ORI}](t_3 + \epsilon), \\
\max_{t_1 \leq t \leq t_4} [\text{SPN}](t) &< 1 < [\text{SPN}](t_5), \\
\max_{t_4 \leq t \leq t_5} [\text{Esp1}](t) &< 0.1, \\
\max_{0 \leq t \leq t_6} [\text{mass}](t) &> 4m_w.
\end{aligned}$$

3.1.7. Telophase Arrest

While the G1 phase is the earliest a cell can be arrested, telophase is the latest a cell can become arrested. A telophase arrested cell must complete all of the checkpoints except that [Clb2] can not drop below K_{ez} before the cell arrests. The first eight constraints for such a cell would

be the same as the constraints for a viable cell. The ninth constraint would be removed, and the tenth constraint would be changed to

$$\max_{t_2+\epsilon \leq t \leq t_7} [\text{Clb2}](t) < K_{\text{ez}} < \min_{t_7+\epsilon \leq t \leq t_8} [\text{Clb2}](t),$$

with the additional constraints

$$\begin{aligned} t_1 &< t_7 < t_8, \\ [\text{mass}](t_8) &> 4m_w. \end{aligned}$$

3.1.8. Evaluating All of the Mutants

As described earlier, a parameter vector must satisfy all of the constraints for all of the mutants before it can be considered feasible. Each mutant requires a separate simulation and will have its own set of time variables derived from its concentrations' trajectories. The mutants are numbered as in Appendix A, so the time for the i th mutant will be $t^{\{i\}}$. Also, because each mutant modifies the parameter vector slightly, the concentration of a substance at a specific time will vary among the mutants. Rather than explicitly specify which parameter vector is being used, the superscript of the time will indicate the parameter vector being used. For instance, mutant 13 is a G1 arrested mutant, so its constraints would be

$$\begin{aligned} \max\{t_1^{\{13\}} - t_w, \max_{0 \leq t^{\{13\}} \leq t_1^{\{13\}}} [\text{mass}](t^{\{13\}}) - 4m_w\} &> 0, \\ \max_{0 \leq t^{\{13\}} \leq t_1^{\{13\}}} [\text{ORI}](t^{\{13\}}) &< 1, \\ \max_{0 \leq t^{\{13\}} \leq t_1^{\{13\}}} [\text{SPN}](t^{\{13\}}) &< 1, \\ \max_{0 \leq t^{\{13\}} \leq t_1^{\{13\}}} [\text{BUD}](t^{\{13\}}) &< 1. \end{aligned}$$

3.2. Evaluating the Constraints

To test this formulation, the constraints for the telophase arrest and viable phenotypes were evaluated at two parameter vectors. The *pds1Δ* mutants were excluded from this test, leaving 61 mutants with an observed phenotype of viable and 15 mutants with an observed phenotype of telophase arrested. The first parameter vector used the best manually obtained biological parameters [5], and the second vector used the best biological parameters found in Chapter 2.

Table 4. The numbering of the viable constraints for Tables 6, 7, 8, and 9.

Index	Constraint
1	$t_1 < t_2$
2	$t_2 < t_3$
3	$t_3 < t_4$
4	$t_4 < t_5$
5	$t_5 < t_6$
6	$t_1 < t_7$
7	$t_7 < t_8$
8	$t_6 < t_8$
9	$t_3 - t_1 < t_w$
10	$ t_3 - t_0 - G_1 < \delta_{G_1}$
11	G1 stability constraint
12	$\min_{t_1 \leq t \leq t_2} ([\text{Clb2}](t) + [\text{Clb5}](t)) > K_{\text{ez2}}$
13	$K_{\text{ez2}} > [\text{Clb2}](t_2 + \epsilon) + [\text{Clb5}](t_2 + \epsilon)$
14	$\max_{t_2 + \epsilon \leq t \leq t_3} [\text{ORI}](t) < 1$
15	$1 < [\text{ORI}](t_3 + \epsilon)$
16	$\max_{t_1 \leq t \leq t_4} [\text{SPN}](t) < 1$
17	$1 < [\text{SPN}](t_5)$
18	$\max_{t_4 \leq t \leq t_5} [\text{Esp1}](t) < 0.1$
19	$0.1 < [\text{Esp1}](t_6)$
20	$[\text{BUD}](t_7) > 1$
21	$[\text{Clb2}](t_8 - \epsilon) > K_{\text{ez}}$
22	$K_{\text{ez}} > [\text{Clb2}](t_8)$
23	$\max_{t_1 \leq t \leq t_8 - \epsilon} [\text{mass}](t) < 4m_w$
24	$ [\text{mass}](t_8 - \epsilon)/m_w - M < \delta_m$
25	Mass stability constraint

For both vectors, the time parameters were found using an algorithm that examines the ODE model simulation time series output and picks sensible values. In one pass through the simulation time series output, this algorithm attempts to pick the time variables so that they are in the proper order, and if possible, the constraints are satisfied. During the simulation, the algorithm keeps track of the maximum and minimum values for the model variables in the constraints (e.g., [BUD], [SPN]). When an event (cf. Section 2.1) occurs, the algorithm sets the respective time

Table 5. The numbering of the telophase arrest constraints for Tables 6, 7, 8, and 9.

Index	Constraint
1	$t_1 < t_2$
2	$t_2 < t_3$
3	$t_3 < t_4$
4	$t_4 < t_5$
5	$t_5 < t_6$
6	$t_1 < t_7$
7	$t_7 < t_8$
8	$\max_{t_2+\epsilon \leq t \leq t_3} [\text{ORI}](t) < 1$
9	$1 < [\text{ORI}](t_3 + \epsilon)$
10	$\max_{t_1 \leq t \leq t_4} [\text{SPN}](t) < 1$
11	$1 < [\text{SPN}](t_5)$
12	$\max_{t_4 \leq t \leq t_5} [\text{Esp1}](t) < 0.1$
13	$0.1 < [\text{Esp1}](t_6)$
14	$\max_{t_1+\epsilon \leq t \leq t_7} [\text{Clb2}](t) < K_{\text{ez}}$
15	$K_{\text{ez}} < \min_{t_7+\epsilon \leq t \leq t_8} [\text{Clb2}](t)$
16	$[\text{mass}](t_8) > 4m_w$

variables and then checks to see if all of the earlier time variables have been set. If there are earlier time variables that have not been set, the algorithm attempts to set each of them to a time that maintains the ordering of the variables and minimizes the violation of the constraints. A more formal description of this algorithm (for the viable phenotype) is given below (ϵ comes from Section 3.1). There is a similar algorithm for the telophase arrested phenotype.

*check_**: flags for checking earlier events

o: offset into the time parameters

n_c: number of cycles completed

reset_cycle: flag for resetting best times

s: maximum time that a previous event can be set to

t: current time

t_f: end of the simulation time

t_b: best time for [BUD] rising

t_c : best time for [Clb2] + [Clb5] dropping

t_e : best time for [Esp1] rising

t_r : best time for [ORI] rising

t_s : best time for [SPN] rising

$t := 0$; $n_c := 0$; $o := 1$; $s := 0$; $t_c := 0$; $t_r := 0$; $t_s := 0$; $t_e := 0$; $t_b := 0$

while ($t < t_f$) **do**

$check_Clb2_Clb5 := \mathbf{false}$; $check_ORI := \mathbf{false}$; $check_ORI := \mathbf{false}$; $check_SPN := \mathbf{false}$

$check_Esp1 := \mathbf{false}$; $check_BUD := \mathbf{false}$; $reset_cycle := \mathbf{false}$

if ($n_c > 1$) **then**

$o := 8 + (n_c - 1) \times 7$

else

$o := 1$

end if

Advance t to the next time in the ODE simulation time series output.

if ([Clb2] + [Clb5] dropped through K_{ez2}) **then**

$t_{o+1} := t - \epsilon/2$; $t_r := -1$; $t_s := -1$; $t_e := -1$

elseif ([ORI] rose through one) **then**

$t_{o+2} := t - \epsilon/2$; $s := t_{o+2} - \epsilon$; $t_s := -1$; $t_e := -1$; $check_Clb2_Clb5 := \mathbf{true}$

elseif ([SPN] rose through one) **then**

$t_{o+3} := t - \epsilon$; $t_{o+4} := t + \epsilon$; $s := t_{o+3} - \epsilon$; $t_e := -1$; $check_ORI := \mathbf{true}$

elseif ([Esp1] rose through 0.1) **then**

$t_{o+5} := t + \epsilon$; $s := t_{o+5} - \epsilon$; $check_SPN := \mathbf{true}$

elseif ([BUD] rose through one) **then**

$t_{o+6} := t + \epsilon$

elseif ([Clb2] rose through K_{ez}) **then**

$t_{o+7} := t - \epsilon/2$; $n_c := n_c + 1$; $s := t_{o+7} - \epsilon$

$check_Esp1 := \mathbf{true}$; $check_BUD := \mathbf{true}$; $reset_cycle := \mathbf{true}$

end if

```

if ( $check\_BUD = \mathbf{true}$  and  $t_{o+6}$  has not been set) then
    if ( $[BUD](s) < [BUD](t_b)$ ) then
         $t_{o+6} := t_b$ 
    else
         $t_{o+6} := s$ 
    end if
end if

if ( $check\_Esp1 = \mathbf{true}$  and  $t_{o+5}$  has not been set) then
    if ( $[Esp1](s) < [Esp1](t_e)$ ) then
         $t_{o+5} := t_e$ 
    else
         $t_{o+5} := s$ 
    end if

     $check\_SPN := \mathbf{true}$ 
end if

 $s := t_{o+5} - \epsilon$ 

if ( $check\_SPN = \mathbf{true}$  and  $t_{o+4}$  has not been set) then
    if ( $[SPN](s) < [SPN](t_s)$ ) then
         $t_{o+3} := t_s - \epsilon; t_{o+4} := t_s + \epsilon$ 
    else
         $t_{o+3} := s - \epsilon; t_{o+4} := s + \epsilon$ 
    end if

     $check\_ORI := \mathbf{true}$ 
end if

 $s := t_{o+3} - \epsilon$ 

if ( $check\_ORI = \mathbf{true}$  and  $t_{o+2}$  has not been set) then
    if ( $[ORI](s) < [ORI](t_r)$ ) then
         $t_{o+2} := t_r$ 

```

```

    else

         $t_{o+2} := s$ 

    end if

     $check\_Clb2\_Clb5 := \mathbf{true}$ 

 $s := t_{o+2} - \epsilon$ 

if ( $check\_Clb2\_Clb5 = \mathbf{true}$  and  $t_{o+1}$  has not been set) then

    if ( $[Clb2](s) + [Clb5](s) > [Clb2](t_c) + [Clb5](t_c)$ ) then

         $t_{o+1} := t_c$ 

    else

         $t_{o+1} := s$ 

    end if

end if

if ( $reset\_cycle = \mathbf{true}$ ) then

     $t_c := -1; t_r := -1; t_s := -1; t_e := -1; t_b := -1$ 

end if

if ( $[Clb2](t) + [Clb5](t) < [Clb2](t_c) + [Clb5](t_c)$  or  $t_c < 0$ ) then

     $t_c := t$ 

end if

if ( $[ORI](t) > [ORI](t_r)$  or  $t_c < 0$ ) then

     $t_r := t$ 

end if

if ( $[SPN](t) > [SPN](t_s)$  or  $t_c < 0$ ) then

     $t_s := t$ 

end if

if ( $[Esp1](t) > [Esp1](t_e)$  or  $t_c < 0$ ) then

     $t_e := t$ 

end if

if ( $[BUD](t) > [BUD](t_b)$  or  $t_c < 0$ ) then

```

$t_b := t$

end if

end while

For conciseness, the violated constraints for viable mutants are listed as “Cx-Ny”, where x indicates the cycle in which the violation occurred, and y is an index into Table 4 that indicates which constraint was violated. The violated constraints for telophase-arrested mutants are listed as “Ny”, where y is an index into Table 5. The results of evaluating the viable constraints on all of the mutants with an observed phenotype are shown in Tables 6, 7, 8, and 9. Tables 6 and 7 show the mutants that satisfied and did not satisfy the constraints, respectively, when the manually obtained biological parameters were used. Tables 8 and 9 show the same for the mathematically optimized biological parameters.

Table 6. Mutants that had no violated constraints, manually obtained parameters.

Wild type in glucose
 Wild type in galactose
cln1Δ cln2Δ
GAL-CLN2 cln1Δ cln2Δ
cln1Δ cln2Δ sic1Δ
cln1Δ cln2Δ cdh1Δ
cln3Δ
bck2Δ
 Multi-copy *BCK2*
cln3Δ bck2Δ GAL-CLN2 cln1Δ cln2Δ
cln1Δ cln2Δ cln3Δ GAL-CLN2
cln1Δ cln2Δ cln3Δ GAL-CLN3
cln1Δ cln2Δ cln3Δ sic1Δ
cln1Δ cln2Δ cln3Δ cdh1Δ
cln1Δ cln2Δ cln3Δ multi-copy CLB5
cln1Δ cln2Δ cln3Δ GAL-CLB5
cln1Δ cln2Δ cln3Δ multi-copy BCK2
sic1Δ
GAL-SIC1
GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ
GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ cdh1Δ
sic1Δ cdh1Δ GALL-CDC20
cdh1Δ
cdc6Δ2-49
sic1Δ cdc6Δ2-49
GAL-CLB2
 Multicopy *GAL-CLB2*
CLB2-dbΔ
CLB2-dbΔ in galactose

CLB2-dbΔ multicopy SIC1
CLB2-dbΔ GAL-SIC1
CLB2-dbΔ clb5Δ
CLB2-dbΔ clb5Δ in galactose
clb5Δ clb6Δ
GAL-CLB5
GAL-CLB5 cdh1Δ
CLB5-dbΔ
tem1Δ
GAL-TEM1
tem1-ts GAL-CDC15
tem1Δ net1-ts
tem1-ts multicopy CDC14
cdc15Δ
Multicopy CDC15
cdc15Δ net1-ts
cdc15-tsmulticopy CDC14
net1-ts
GAL-NET1
cdc14-ts
GAL-NET1 GAL-CDC14
TAB6-1 cdc15Δ
mad2Δ
bub2Δ
mad2Δ bub2Δ
APC-A
APC-A cdh1Δ
APC-A cdh1Δ GAL-SIC1
APC-A cdh1Δ GAL-CDC6
APC-A cdh1Δ multicopy CDC20
swi5Δ
sic1Δ cdc6Δ2-49 cdh1Δ GALL-CDC20
APC-A sic1Δ
APC-A GAL-CLB2

Table 7. Mutants that had violated constraints, manually obtained parameters.

<i>GAL-CLN2 cln1Δ cln2Δ cdh1Δ</i>	C3-N20, C4-N19, C4-N20, C5-N15, C5-N17, C5-N20, C6-N15, C6-N17, C6-N20, C7-N15, C7-N17, C7-N20, C8-N15, C8-N17, C8-N20
<i>GAL-CLN3</i>	C6-N1, C8-N11, C5-N19, C6-N13, C6-N17, C7-N15, C7-N17, C8-N15, C8-N17
<i>cln1Δ cln2Δ bck2Δ</i>	C8-N24
<i>cdh1Δ cdc6Δ2-49</i>	C8-N10
<i>GAL-CLB2 sic1Δ</i>	N12, N16
<i>GAL-CLB2-dbΔ</i>	N12
<i>GAL-ESP1 cdc20-ts</i>	N5, N12

<i>cdc14-ts GAL-SIC1</i>	C1-N20, C2-N13, C2-N18, C2-N22, C3-N13, C3-N18, C3-N22, C4-N13, C4-N18, C4-N22, C5-N13, C5-N18, C5-N22, C6-N13, C6-N18, C6-N22, C7-N13, C7-N18, C7-N22, C8-N13, C8-N18, C8-N22
<i>TAB6-1</i>	C1-N8
<i>TAB6-1 CLB1 clb2Δ</i>	C1-N17, C2-N15, C2-N17, C3-N15, C3-N17, C4-N15, C4-N17, C5-N15, C5-N17, C6-N15, C6-N17, C7-N15, C7-N17, C8-N15, C8-N17
<i>APC-A cdh1Δ</i> in galactose	C1-N20, C2-N13, C2-N18, C2-N20, C2-N22, C3-N13, C3-N18, C3-N20, C3-N22, C4-N13, C4-N18, C4-N20, C4-N22, C5-N13, C5-N18, C5-N20, C5-N22, C6-N13, C6-N18, C6-N20, C6-N22, C7-N13, C7-N18, C7-N20, C7-N22, C8-N13, C8-N18, C8-N20, C8-N22
<i>APC-A cdh1Δ</i> multicopy <i>SIC1</i>	C1-N17, C2-N15, C2-N17
<i>APC-A cdh1Δ</i> multicopy <i>CDC6</i>	C8-N11, C7-N22, C8-N13, C8-N18, C8-N22

Table 8. Mutants that had no violated constraints, optimized parameters.

Wild type in glucose
 Wild type in galactose
cln1Δ cln2Δ
GAL-CLN2 cln1Δ cln2Δ
cln1Δ cln2Δ sic1Δ
cln1Δ cln2Δ cdh1Δ
GAL-CLN2 cln1Δ cln2Δ cdh1Δ
cln3Δ
GAL-CLN3
bck2Δ
 Multi-copy *BCK2*
cln1Δ cln2Δ bck2Δ
cln3Δ bck2Δ GAL-CLN2 cln1Δ cln2Δ
cln1Δ cln2Δ cln3Δ GAL-CLN2
cln1Δ cln2Δ cln3Δ GAL-CLN3
cln1Δ cln2Δ cln3Δ sic1Δ
cln1Δ cln2Δ cln3Δ cdh1Δ
cln1Δ cln2Δ cln3Δ multi-copy *CLB5*
cln1Δ cln2Δ cln3Δ GAL-CLB5
cln1Δ cln2Δ cln3Δ multi-copy *BCK2*
sic1Δ
GAL-SIC1
GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ
GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ cdh1Δ
sic1Δ cdh1Δ GALL-CDC20
cdh1Δ
cdc6Δ2-49
sic1Δ cdc6Δ2-49

GAL-CLB2
 Multicopy *GAL-CLB2*
CLB2-dbΔ
CLB2-dbΔ GAL-SIC1
CLB2-dbΔ clb5Δ
CLB2-dbΔ clb5Δ in galactose
clb5Δ clb6Δ
GAL-CLB5
GAL-CLB5 cdh1Δ
CLB5-dbΔ
tem1Δ
GAL-TEM1
tem1-ts GAL-CDC15
tem1Δ net1-ts
tem1-ts multicopy CDC14
cdc15Δ
 Multicopy *CDC15*
cdc15Δ net1-ts
cdc15-tsmulticopy CDC14
GAL-NET1
cdc14-ts
GAL-NET1 GAL-CDC14
TAB6-1
TAB6-1 cdc15Δ
TAB6-1 CLB1 clb2Δ
mad2Δ
bub2Δ
APC-A
APC-A cdh1Δ
APC-A cdh1Δ in galactose
APC-A cdh1Δ multicopy SIC1
APC-A cdh1Δ GAL-SIC1
APC-A cdh1Δ multicopy CDC6
APC-A cdh1Δ GAL-CDC6
APC-A cdh1Δ multicopy CDC20
swi5Δ
APC-A sic1Δ
APC-A GAL-CLB2

Table 9. The mutants that had violated constraints, optimized parameters.

<i>cdh1Δ cdc6Δ2-49</i>	C8-N10
<i>GAL-CLB2 sic1Δ</i>	N12, N16
<i>CLB2-dbΔ</i> in galactose	N16
<i>CLB2-dbΔ multicopy SIC1</i>	C8-N22, C8-N25
<i>GAL-CLB2-dbΔ</i>	N12
<i>GAL-ESP1 cdc20-ts</i>	N5, N12
<i>net1-ts</i>	C1-N8

<i>cdc14-ts GAL-SIC1</i>	C1-N20, C2-N13, C2-N18, C2-N22, C3-N13, C3-N18, C3-N22, C4-N13, C4-N18, C4-N22, C5-N13, C5-N18, C5-N22, C6-N13, C6-N18, C6-N22, C7-N13, C7-N18, C7-N22, C8-N13, C8-N18, C8-N22
<i>mad2Δ bub2Δ</i>	C1-N16
<i>sic1Δ cdc6Δ2-49 cdh1Δ GALL-CDC20</i>	C1-N16, C1-N18

The full inequality formulation for all mutants would have approximately 11,500 constraints (sum of constraints per phenotype accounting for multiple cycles) and 4000 variables (almost all of them times $t_j^{\{i\}}$). One ODE solution and the transforms necessary to match the ODE solution to experimental data take about 17 seconds on a 2.3 GHz G5 processor, so the need for parallel supercomputing is clear. While the nonlinear inequality formulation is definitely large scale, it is well within the range of problems being solved in industry. This work demonstrates the reasonableness of the inequality approach to parameter estimation for cell cycle modeling. Mechanically assembling all 11,500 constraints and then solving the feasibility problem will be a major undertaking requiring several man-years of effort and parallel supercomputing to evaluate the constraints, but is surely doable.

3.3. Optimization

Two simple test problems were created to test the ability of algorithms to find feasible parameter vectors. To create the test problems, time parameters were found that satisfied the constraints for the wild-type using the manually discovered model parameters. The starting point for each of the test problems was a small perturbation of the wild-type model parameters, with the time parameters kept at their previous values. For both starting points, the simulation output satisfied all of the rules of viability, but the constraints were not satisfied because the events occurred at different times. For the first test problem, a randomly picked model parameter was increased by 10%; for the second test problem, all model parameters were increased by 5%. It has already been shown that the wild-type is robust with respect to the model parameters [5], so it is unsurprising that these changes do not affect the wild-type's viability.

Two algorithms were used on both test problems. The FILTRANE algorithm [7] is designed to find a feasible point subject to smooth inequality constraints and simple bounds. The LANCELOT B algorithm [8] is designed to minimize an objective function subject to simple bound constraints and smooth equality constraints. To apply LANCELOT B to the constraints, the constraints were converted to an objective function. Let m be the number of constraints, and let $c_i(x) > 0$ for

$i = 1, \dots, m$ be the inequality constraints for the viable phenotype. Then the objective function (related to the Kreisselmeier-Steinhauser envelope function [15]) that was used in LANCELOT B is $\sum_{i=1}^m e^{-\omega \rho_i c_i(x)}$, where ω is a positive weight used to control the shape of the objective function, and ρ_i is used to scale the i th constraint.

Both algorithms were run three times: (1) optimizing only the model parameters, (2) optimizing the time parameters, and (3) optimizing both model and time parameters. Although both algorithms were able to reduce the constraint violations, neither algorithm was able to find a feasible parameter vector for either test problem. The FILTRANE algorithm seems to have inherent limitations that prevent it from making further progress. With LANCELOT B, the objective function that was constructed was probably responsible for the inability to find a feasible point. Several different values were used for ω , but LANCELOT B was still unable to find a feasible parameter vector. The ρ_i s for the constraints that were violated were increased, which caused LANCELOT B to satisfy those constraints, but the resulting parameter vectors violated other constraints.

4. Conclusion

Two algorithms were able to optimize the discontinuous formulation. The optimization results indicate that the robustness of a parameter vector must be considered in future optimization efforts. It is encouraging that the constraint formulation's classification of mutants was consistent with that of the discontinuous formulation. Although optimization efforts were unsuccessful with the constraint formulation, the constraints may have other strengths; their inherent smoothness may be useful for determining the robustness of parameters.

5. Bibliography

1. Allen, N. A., Calzone, L., Chen, K. C., Ciliberto, A., Ramakrishnan, N., Shaffer, C. A., Sible, J. C., Tyson, J. J., Vass, M. T., Watson, L. T., and Zwolak, J. W. (2003), Modeling regulatory networks at Virginia Tech, *OMICS* 7: 285–299.
2. Allen, N. A., Chen, K. C., Tyson, J. J., Shaffer, C. A., and Watson, L. T., Computer evaluation of network dynamics models with application to cell cycle control in budding yeast, *IEE Systems Biology* (to appear).
3. Allen, N. A., Shaffer, C. A., Ramakrishnan, N., Vass, M. T., and Watson, L. T. (2003), Improving the development process for eukaryotic cell cycle models with a modeling support environment, *Simulation* 79: 674–688.
4. Audet, C. and Dennis Jr., J. E., Mesh adaptive direct search algorithms for constrained optimization, *SIAM Journal on Optimization* (to appear).
5. Chen, K. C., Calzone, L., Csikasz-Nagy, A., Cross, F. R., Novak, B., and Tyson, J. J. (2004), Integrative analysis of cell cycle control in budding yeast, *Molecular Biology of the Cell*, 15: 3841–3862.
6. Coope, I. D. and Price, C. J. (2000), Frame based methods for unconstrained optimization, *Journal of Optimization Theory and Applications* 107: 261–274.
7. Gould, N. I. M. and Toint, Ph. L. (2003), FILTRANE, A Fortran 95 Filter-Trust-Region Package for Solving Nonlinear Feasibility Problems, The University of Namur, Belgium, Technical Report 03/17.
8. Gould, N. I. M., Orban, D., and Toint, Ph. L., (2004), GALAHAD, a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization, *ACM Trans. Math. Software* 29: 353–372.
9. He, J., Sosonkina, M., Shaffer, C. A., Tyson, J. J., Watson, L. T., and Zwolak, J. W. (2004), A hierarchical parallel scheme for a global search algorithm, in *Proc. High Performance Computing Symposium 2004*, J. Meyer (ed.), Soc. for Modeling and Simulation Internat., San Diego, CA, 43–50.
10. He, J., Sosonkina, M., Shaffer, C. A., Tyson, J. J., Watson, L. T., and Zwolak, J. W. (2004), A hierarchical parallel scheme for global parameter estimation in systems biology, in *Proceedings of the 18th International Parallel & Distributed Processing Symposium*, CD-ROM, IEEE Computer Soc., Los Alamitos, CA, 9 pages.
11. He, J., Sosonkina, M., Watson, L. T., Verstak, A., and Zwolak, J. W. (2005), Data-distributed parallelism with dynamic task allocation for a global search algorithm, in *Proc. High Performance Computing Symposium 2005*, M. Parashar and L. Watson (eds.), Society for Modeling and Simulation International, San Diego, CA, 164–172.
12. He, J., Watson, L. T., Ramakrishnan, N., Shaffer, C. A., Verstak, A., Jiang, J., Bae, K., and Tranter, W. H. (2002), Dynamic data structures for a direct search algorithm, *Computational Optimization and Applications* 23: 5–25.
13. Jones, D. R., Perttunen, C. D., and Stuckman, B. E. (1993), Lipschitzian optimization without the Lipschitz constant, *Journal of Optimization Theory and Applications*, 79: 157–181.
15. Kreisselmeier, G. and Steinhauser, R. (1979), Systematic Control Design by Optimizing a Vector Performance Index, *Proc. IFAC Symp. on Computer Aided Design of Control Systems, Zurich, Switzerland*, 113.
16. Murray, A. and Hunt, T. (1993), *The Cell Cycle: an Introduction*, Oxford University Press, New York, 1993.
17. Nasmyth, K. (1996) At the heart of the budding yeast cell cycle, *Trends in Genetics*, 12: 405–412.
18. Nurse, P. (2000), A long twentieth century of the cell cycle and beyond, *Cell*, 100: 71–78.

6. Appendix A: Mutant Phenotypes

Listed below are the phenotypic characteristics of all the mutants used to construct the objective function. The data are expressed as a six-tuple (v, g, m, a, t, c) as described in Section 3. The meanings of the values for t are described in [2]. For all fields other than viability, a dash (-) means that the data is either not available or not applicable for that mutant.

	Mutant name	v	g	m	a	t	c
1.	Wild type in glucose	viable	35.2	1	-	-	-
2.	Wild type in galactose	viable	109	1	-	-	-
3.	<i>cln1Δ cln2Δ</i>	viable	-	2	-	-	-
4.	<i>GAL-CLN2 cln1Δ cln2Δ</i>	viable	-	0.5	-	-	-
5.	<i>cln1Δ cln2Δ sic1Δ</i>	viable	-	-	-	-	-
6.	<i>cln1Δ cln2Δ cdh1Δ</i>	viable	-	-	-	-	-
7.	<i>GAL-CLN2 cln1Δ cln2Δ cdh1Δ</i>	viable	-	1.7	-	-	-
8.	<i>cln3Δ</i>	viable	-	1.7	-	-	-
9.	<i>GAL-CLN3</i>	viable	-	0.44	-	-	-
10.	<i>bck2Δ</i>	viable	-	1.4	-	-	-
11.	Multi-copy <i>BCK2</i>	viable	-	0.8	-	-	-
12.	<i>cln1Δ cln2Δ bck2Δ</i>	viable	-	1.7	-	-	-
13.	<i>cln3Δ bck2Δ</i>	inviable	-	-	licensed	5	0
14.	<i>cln3Δ bck2Δ GAL-CLN2 cln1Δ cln2Δ</i>	viable	-	-	-	-	-
15.	<i>cln3Δ bck2Δ</i> multi-copy <i>CLN2</i>	inviable	-	-	licensed	5	-
16.	<i>cln3Δ bck2Δ GAL-CLB5</i>	inviable	-	-	-	-	-
17.	<i>cln3Δ bck2Δ sic1Δ</i>	inviable	-	-	-	-	-
18.	<i>cln1Δ cln2Δ cln3Δ</i>	inviable	-	-	licensed	5	0
19.	<i>cln1Δ cln2Δ cln3Δ GAL-CLN2</i>	viable	-	-	-	-	-
20.	<i>cln1Δ cln2Δ cln3Δ GAL-CLN3</i>	viable	-	-	-	-	-
21.	<i>cln1Δ cln2Δ cln3Δ sic1Δ</i>	viable	10	3.5	-	-	-
22.	<i>cln1Δ cln2Δ cln3Δ cdh1Δ</i>	inviable	-	-	separated	3	-
23.	<i>cln1Δ cln2Δ cln3Δ</i> multi-copy <i>CLB5</i>	viable	-	-	-	-	-
24.	<i>cln1Δ cln2Δ cln3Δ GAL-CLB5</i>	viable	-	-	-	-	-
25.	<i>cln1Δ cln2Δ cln3Δ</i> multi-copy <i>BCK2</i>	viable	-	-	-	-	-
26.	<i>cln1Δ cln2Δ cln3Δ GAL-CLB2</i>	inviable	-	-	licensed	5	0
27.	<i>cln1Δ cln2Δ cln3Δ apc-ts</i>	inviable	-	-	aligned	3	0
28.	<i>sic1Δ</i>	viable	15	1	-	-	-
29.	<i>GAL-SIC1</i>	viable	135	2	-	-	-
30.	<i>GAL-SIC1-dbΔ</i>	inviable	-	-	licensed	5	0
31.	<i>GAL-SIC1 cln1Δ cln2Δ</i>	inviable	-	-	licensed	5	-
32.	<i>GAL-SIC1 cln1Δ cln2Δ cdh1Δ</i>	inviable	-	-	licensed	5	-
33.	<i>GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ</i>	viable	-	-	-	-	-
34.	<i>GAL-SIC1 GAL-CLN2 cln1Δ cln2Δ cdh1Δ</i>	viable	-	-	-	-	-
35.	<i>sic1Δ cdh1Δ</i>	inviable	-	-	unlicensed	1	1
36.	<i>sic1Δ cdh1Δ GALL-CDC20</i>	viable	-	-	-	-	-
37.	<i>cdh1Δ</i>	viable	-	0.6	-	-	-
38.	Cdh1 constitutively active	inviable	-	-	fired	3	-
39.	<i>cdc6Δ2-49</i>	viable	-	-	-	-	-
40.	<i>sic1Δ cdc6Δ2-49</i>	viable	-	-	-	-	-
41.	<i>cdh1Δ cdc6Δ2-49</i>	viable	20	2.0	-	-	-
42.	<i>clb1Δ clb2Δ</i>	inviable	-	-	fired	3	0
43.	<i>GAL-CLB2</i>	viable	-	-	-	-	-

44.	Multicopy <i>GAL-CLB2</i>	inviable	-	-	separated	3	0
45.	<i>GAL-CLB2 sic1Δ</i>	inviable	-	-	separated	3	0
46.	<i>GAL-CLB2 cdh1Δ</i>	inviable	-	-	-	-	-
47.	<i>CLB2-dbΔ</i>	inviable	-	-	separated	3	-
48.	<i>CLB2-dbΔ</i> in galactose	inviable	-	-	separated	3	-
49.	<i>CLB2-dbΔ</i> multicopy <i>SIC1</i>	viable	-	-	-	-	-
50.	<i>CLB2-dbΔ GAL-SIC1</i>	viable	-	-	-	-	-
51.	<i>CLB2-dbΔ clb5Δ</i>	inviable	-	-	separated	3	0
52.	<i>CLB2-dbΔ clb5Δ</i> in galactose	viable	-	-	-	-	-
53.	<i>GAL-CLB2-dbΔ</i>	inviable	-	-	separated	3	-
54.	<i>clb5Δ clb6Δ</i>	viable	65	-	-	-	-
55.	<i>cln1Δ cln2Δ clb5Δ clb6Δ</i>	inviable	-	-	licensed	5	0
56.	<i>GAL-CLB5</i>	viable	-	-	-	-	-
57.	<i>GAL-CLB5 sic1Δ</i>	inviable	-	-	unlicensed	-	1
58.	<i>GAL-CLB5 cdh1Δ</i>	inviable	-	-	-	-	-
59.	<i>CLB5-dbΔ</i>	viable	-	-	-	-	-
60.	<i>CLB5-dbΔ sic1Δ</i>	inviable	-	-	-	-	1
61.	<i>CLB5-dbΔ pds1Δ</i>	viable	-	-	-	-	-
62.	<i>CLB5-dbΔ pds1Δ cdc20Δ</i>	inviable	-	-	separated	3	0
63.	<i>GAL-CLB5-dbΔ</i>	inviable	-	-	-	-	1
64.	<i>cdc20-ts</i>	inviable	-	-	aligned	3	0
65.	<i>cdc20Δ clb5Δ</i>	inviable	-	-	aligned	3	0
66.	<i>cdc20Δ pds1Δ</i>	inviable	-	-	separated	3	0
67.	<i>cdc20Δ pds1Δ clb5Δ</i>	viable	-	-	-	-	-
68.	<i>GAL-CDC20</i>	inviable	-	-	fired	10	0
69.	<i>cdc20-ts mad2Δ</i>	inviable	-	-	aligned	3	0
70.	<i>cdc20-ts bub2Δ</i>	inviable	-	-	aligned	3	0
71.	<i>pds1Δ</i>	viable	-	-	-	-	-
72.	<i>esp1-ts</i>	inviable	-	-	aligned	1	0
73.	<i>PDS1-dbΔ</i>	inviable	-	-	aligned	1	0
74.	<i>GAL-PDS1-dbΔ</i>	inviable	-	-	aligned	1	0
75.	<i>GAL-PDS1-dbΔ esp1-ts</i>	inviable	-	-	aligned	1	0
76.	<i>GAL-ESP1 cdc20-ts</i>	inviable	-	-	separated	3	0
77.	<i>tem1Δ</i>	inviable	-	-	separated	3	0
78.	<i>GAL-TEM1</i>	viable	-	-	-	-	-
79.	<i>tem1-ts GAL-CDC15</i>	viable	-	-	-	-	-
80.	<i>tem1Δ net1-ts</i>	viable	-	-	-	-	-
81.	<i>tem1-ts</i> multicopy <i>CDC14</i>	viable	-	-	-	-	-
82.	<i>cdc15Δ</i>	inviable	-	-	separated	3	0
83.	Multicopy <i>CDC15</i>	viable	-	-	-	-	-
84.	<i>cdc15-ts</i> multicopy <i>TEM1</i>	inviable	-	-	-	-	-
85.	<i>cdc15Δ net1-ts</i>	viable	-	-	-	-	-
86.	<i>cdc15-ts</i> multicopy <i>CDC14</i>	viable	-	-	-	-	-
87.	<i>net1-ts</i>	viable	50	-	-	-	-
88.	<i>GAL-NET1</i>	inviable	-	-	separated	3	0
89.	<i>cdc14-ts</i>	inviable	-	-	separated	3	0
90.	<i>GAL-CDC14</i>	inviable	-	-	licensed	5	0
91.	<i>GAL-NET1 GAL-CDC14</i>	viable	-	-	-	-	-
92.	<i>net1Δ cdc20-ts</i>	inviable	-	-	aligned	1	-
93.	<i>cdc14-ts GAL-SIC1</i>	viable	-	-	-	-	-

94.	<i>TAB6-1</i>	viable	-	-	-	-	-
95.	<i>TAB6-1 cdc15Δ</i>	viable	-	-	-	-	-
96.	<i>TAB6-1 clb5Δ clb6Δ</i>	inviable	-	-	licensed	5	0
97.	<i>TAB6-1 CLB1 clb2Δ</i>	viable	-	-	-	-	-
98.	<i>mad2Δ</i>	viable	35	1	-	-	-
99.	<i>bub2Δ</i>	viable	35	1	-	-	-
100.	<i>mad2Δ bub2Δ</i>	viable	-	-	-	-	-
101.	<i>APC-A</i>	viable	20	1.5	-	-	-
102.	<i>APC-A cdh1Δ</i>	inviable	-	-	separated	3	-
103.	<i>APC-A cdh1Δ</i> in galactose	viable	-	-	-	-	-
104.	<i>APC-A cdh1Δ</i> multicopy <i>SIC1</i>	viable	-	-	-	-	-
105.	<i>APC-A cdh1Δ GAL-SIC1</i>	viable	-	-	-	-	-
106.	<i>APC-A cdh1Δ</i> multicopy <i>CDC6</i>	viable	-	-	-	-	-
107.	<i>APC-A cdh1Δ GAL-CDC6</i>	viable	-	-	-	-	-
108.	<i>APC-A cdh1Δ</i> multicopy <i>CDC20</i>	viable	-	-	-	-	-
109.	<i>swi5Δ</i>	viable	20	-	-	-	-
110.	<i>sic1Δ cdc6Δ2-49 cdh1Δ</i>	inviable	-	-	fired	3	1
111.	<i>sic1Δ cdc6Δ2-49 cdh1Δ GALL-CDC20</i>	viable	-	-	-	-	-
112.	<i>APC-A cdh1Δ clb5Δ</i>	inviable	-	-	-	-	-
113.	<i>APC-A cdh1Δ pds1Δ</i>	inviable	-	-	-	-	-
114.	<i>APC-A sic1Δ</i>	viable	-	-	-	-	-
115.	<i>APC-A GAL-CLB2</i>	inviable	-	-	separated	3	-

Vita

Thomas D. Panning was born in Tulsa, Oklahoma on August 19th, 1981. He earned a bachelor's degree in Computer Science, with minors in Mathematics and Biology, in December 2003 from Virginia Polytechnic Institute and State University. In January 2004, he started graduate study in computer science at Virginia Polytechnic and State University. In May 2006, he received the Master of Science degree in computer science from Virginia Polytechnic and State University.