# IMPROVING THE DEVELOPMENT PROCESS FOR EUKARYOTIC CELL CYCLE MODELS WITH A MODELING SUPPORT ENVIRONMENT

Nicholas A. Allen Clifford A. Shaffer Naren Ramakrishnan Marc T. Vass Layne T. Watson

Department of Computer Science Virginia Polytechnic Institute and State University Blacksburg, VA 24061-0106, U.S.A.

October 3, 2003

# ABSTRACT

Theoretical molecular biologists attempt to describe cellular processes and regulatory networks with continuous and discrete mathematical models. Previous practice has been to develop models largely by hand, and then to validate them primarily by comparing time series plots versus the observed experimental results. This paper reports our experiences in designing and building a modeling support environment (MSE) for cell cycle models. We describe improvements to the development process for molecular network models by (a) identifying the key elements of the existing modeling process, (b) incorporating simulation methodology into a revised modeling process, and (c) building and testing software that supports the revised modeling process.

## **1 INTRODUCTION**

Mathematical models of biochemical control systems attempt to derive observed physiological properties of cells from the underlying molecular regulatory networks. Examples include growth and division, chemotaxis, secretion and circadian rhythms. The hope is that creating such models will lead to a higher-level understanding of the biological processes involved. The common form of such models is (at some point in the process) systems of differential equations with discrete switching. A number of independent groups have been developing tools to support aspects of the modeling process, a sampling includes [24, 27, 33]. BioSPICE [8] is a major effort by DARPA to provide a new generation of interoperable modeling and simulation tools. It seeks to improve the quality of pathway modeling by providing the community with common languages for expressing models [15, 21] and interoperability between various model description editors, simulators, and analysis tools.

However, the recent state of pathway modeling has been

largely ad hoc and labor intensive as most modelers have not tried existing tools, or those tools have proved inadequate for their needs. Many modelers still work by hand-sketching their ideas (see the discussion of wiring diagrams below) and then manually converting those sketches to differential equations. Analysis often involves visual comparisons between time series plots and experimentally collected results.

This paper describes how introducing appropriate modeling tools can improve the speed and accuracy of the model development process (which directly permits the creation of larger, more complex models), and also can lead to a more disciplined approach to the model lifecycle.

We present a brief description of typical molecular regulatory network models in Section 2. The original modeling process observed in a representative portion of the community is described in Section 3. Section 4 describes how the conical methodology [29] relates to the observed modeling process, and can be used to improve the modeling lifecycle. Section 5 briefly describes the JigCell modeling support environment (MSE) developed for cell cycle modeling and related problems, and compares JigCell to existing pathway modeling software. Section 6 describes the impact that the Systems Biology Markup Language (SBML) standardization effort and the BioSPICE project, of which JigCell is a part, have had on the modeling community. Section 7 presents our conclusions.

# 2 MOLECULAR REGULATORY NETWORKS

A simple example of a regulatory network is the set of reactions controlling the activity of MPF (mitosis promoting factor) in frog eggs (Figure 2). Such networks are often represented as graphs where vertices represent substrates and products (collectively referred to as species), and labeled directed edges connecting vertices represent the reactions. Chemical reactions cause the concentrations of the chemical species  $(C_i)$  to change in time according to the equation

$$\frac{dC_i}{dt} = \sum_{j=1}^R P_{ij}v_j, \ i = 1, \dots, N$$

where R is the total number of reactions,  $v_j$  is the velocity of the *j*th reaction in the network, and  $P_{ij}$  is the stoichiometric coefficient of species *i* in reaction *j* ( $p_{ij} < 0$  for substrates,  $p_{ij} > 0$  for products,  $p_{ij} = 0$  if species *i* does not take part in reaction *j*). The full set of rate equations is a mathematical representation of the temporal behavior of the regulatory network. These equations are then solved numerically [35] and the model behavior is interpreted from time-course output of the species concentrations. Modelers are faced with many computational problems: accurately and efficiently solving equations when velocities are characterized by widely varied time constants, finding steady state solutions, estimating rate constants by fitting numerical solutions to experimental data, and identifying bifurcation points in the multi-dimensional parameter space.

For example, a current model of the budding yeast cell cycle consists of about 30 differential equations containing 100 rate constants. The parameters are estimated from the cellcycle behavior of more than 100 mutants defective in the regulatory network. Simulating the entire set takes from a few minutes to an hour on a desktop PC for one choice of kinetic constants. To fit the model to the data by nonlinear regression will likely require thousands of repetitions of the full calculations. A model of such complexity (10-100 equations) represents the upper limit of what a dedicated modeler can produce "by hand" with a good numerical integrator like LSODE [16]. To adequately describe fundamental physiological processes (such as the control of cell division) in mammalian cells will require models of at least 100-1000 equations. To handle this next generation of dynamical models will require sophisticated software to automate the modeling cycle: network specification, equation generation, simulation and data management, and parameter estimation.

Ongoing efforts such as the DARPA BioSPICE initiative [8] aspire to support the necessary increase in model size. This project supports approximately fifteen research groups, including our own. The goals of BioSPICE are to advance the broad efforts of biochemical pathway modeling by supporting modeling efforts, experimentation efforts, and software toolbuilding efforts together. Software efforts are not just to provide sets of tools to modelers, but to make the tools developed by the various groups interoperable. Additionally, tools developed as part of BioSPICE are made freely available under an open source license. While it is not known yet whether the interoperability goals will be met, the BioSPICE project is generating the potential for interaction through the definition of APIs for communication between tools, and language definitions for data such as model definitions and simulation outputs.



Figure 1: Original Modeling Process

#### **3 ORIGINAL MODELING PROCESS**

Figure 1 shows a modeling process that has successfully developed several biochemical pathway models [36]. The process evolved over more than 10 years of developing models. It is not based on formalisms nor documented; new modelers learn the process through demonstration and mentoring. Until recently, the modelers' tools were primarily off-the-shelf for solving and analyzing differential equations, and not specialized for pathway modeling.

Before a model can be developed, there must be a problem that the model intends to solve. Problem formulation includes analysis of requirements, identification of a solution method, and specification of modeling objectives [4]. Without a formulated problem, the modeler risks inadequately solving the problem or solving the wrong problem. A notable feature of the original modeling process is that it deals solely with model development and does not contain problem formulation. Over the past two years of our observation, this particular group of modelers has not formulated a completely new problem (in the sense of wishing to develop a model for a new organism, or apply new solution techniques to a previously developed model). Instead, they expand existing models by attempting to match additional experimental observations. Is this infrequent reformulation an inherent property of the problems these modelers are attempting to solve or a side effect of the current modeling process? Observations with different modeling processes might answer this question in the future.

The original modeling process has four primary stages: design, translate, evaluate, and accept. Models are created and refined in the design and translate stages. Testing occurs during the evaluate stage. The accept stage produces a presentable model from the information the modeler has recorded. This may be the simplest process that could successfully build a model.

In the ensuing text, a stage labeled "x" in a diagram will be



Figure 2: Wiring Diagram

denoted by  $\xrightarrow{x}$ , and such symbols are used to mark the paragraph where those stages are discussed. Stages drawn with solid lines in diagrams indicate successful completion of a process. Stages drawn with dashed lines indicate error-recovery activities.

 $\xrightarrow{design}$  The design stage begins with creating a wiring diagram from an idea of how a biological process is carried out. The wiring diagram is a graph that captures the chemical species (also known as products and reactants) at the nodes, and represents interactions that create, destroy, and convert these species at the arcs. Additionally, the wiring diagram may note the kinetic information for a reaction: describing the mechanics of the reaction and the rate at which the reaction occurs. Figure 2 is a wiring diagram depicting post-translational modification of cyclin, an important regulation process, in frog egg extracts based on the model in [25].

The notation for wiring diagrams is not currently standardized. Modelers often invent ad hoc notation to express abstractions, replication, and unusual processes. Kinetic information is frequently presented separately from the wiring diagram or must be inferred from figures. Without full information about the rate laws and constants, the model can be structurally analyzed but not simulated.

Since the wiring diagram often lacks full details of the kinetics, it is typical for our modelers to first rewrite the model as a series of chemical reaction equations, along with appropriate kinetic functions to describe the reaction.

 $\xrightarrow{translate}$  The translate stage is the process of converting the reaction equations to systems of ordinary differential equations. For each species in the system, the modeler creates a differential equation. Reactions that involve the species determine the right hand side of the differential equations. Parameters for the differential equations are set according to the kinetic information for the reaction. Frequently, exact values for these parameters are not known. In this case, estimates are made for the parameter values and updated as the model is developed.

In some cases, a protein is never created nor destroyed, but is converted between different forms. When this occurs, the quantity is said to be conserved, and one of the differential equations is replaced with an algebraic expression called a conservation relation. In addition to the continuous differential Table 1: Equations from Wiring Diagram

$$\begin{array}{lll} \displaystyle \frac{d\mathbf{C25}_a}{dt} &=& \displaystyle \frac{k25_f*(\mathbf{MPF}_a+\epsilon_1)*\mathbf{C25}_i}{J25_f+\mathbf{C25}_i} - \displaystyle \frac{k25_r*\mathbf{C25}_a}{J25_r+\mathbf{C25}_a} \\ \\ \displaystyle \mathbf{C25}_i &=& 1.0-\mathbf{C25}_a \\ \\ \displaystyle \frac{d\mathbf{CycB}}{dt} &=& k_1-\mathbf{CycB}*(\mathbf{Cyclosome}+k_3*\mathbf{Cdk1}) \\ \\ \displaystyle \mathbf{Cyclosome} &=& \displaystyle (k_2'*\mathbf{Csom}_i+k_2''*\mathbf{Csom}_a) \\ \\ \displaystyle \frac{d\mathbf{MPF}_a}{dt} &=& k_3*\mathbf{CycB}*\mathbf{Cdk1} + \displaystyle (k_c'*\mathbf{C25}_i+k_c''*\mathbf{C25}_a)*\mathbf{MPF}_i \\ &-& \displaystyle (k_w'*\mathbf{Wee}_i+k_w''*\mathbf{Wee}_a+\mathbf{Cyclosome})*\mathbf{MPF}_a \\ \\ \displaystyle \frac{d\mathbf{MPF}_i}{dt} &=& \displaystyle (k_w'*\mathbf{Wee}_i+k_w''*\mathbf{Wee}_a) \\ &-& \displaystyle (k_c'*\mathbf{C25}_i+k_c''*\mathbf{C25}_a+\mathbf{Cyclosome})*\mathbf{MPF}_i \\ \\ \displaystyle \mathbf{Cdk1} &=& 1.0-\mathbf{MPF}_a-\mathbf{MPF}_i \\ \\ \displaystyle \frac{d\mathbf{Wee}_a}{dt} &=& \displaystyle \frac{k\mathbf{Wee}_r*\mathbf{Wee}_i}{J\mathbf{Wee}_r+\mathbf{Wee}_i} - \displaystyle \frac{k\mathbf{Wee}_f*(\mathbf{MPF}_a+\epsilon_2)*\mathbf{Wee}_a}{J\mathbf{Wee}_f+\mathbf{Wee}_a} \\ \\ \displaystyle \mathbf{Wee}_i &=& 1.0-\mathbf{Wee}_a \\ \\ \displaystyle \frac{d\mathbf{Csom}_a}{dt} &=& \displaystyle \frac{k\mathbf{cy}_f*(\mathbf{MPF}_a+\epsilon_3)*\mathbf{Csom}_i}{J\mathbf{cy}_f+\mathbf{Csom}_a} - \displaystyle \frac{k\mathbf{cy}_r*\mathbf{Csom}_a}{J\mathbf{cy}_r+\mathbf{Csom}_a} \\ \\ \\ \displaystyle \mathbf{Csom}_i &=& 1.0-\mathbf{Csom}_a \end{array}$$

C

equation model, there is also a discrete event model. Certain cellular processes, such as cell division, are modeled by discrete events that set species values, alter rate laws or constants, and switch between sets of differential equations driving the continuous model. Table 1 shows a system of ordinary differential equations and conservation relations produced from Figure 2.

 $\xrightarrow{evaluate}$  The evaluate stage begins by generating time series plots of important species concentrations from the model. These plots correspond to experimental observations of the process in the lab. The modeler compares the time series plots with the experimental observations and judges whether the model adequately represents the biological process. In addition to determining if a time series matches observed concentrations, the modeler might also seek to determine if gross physical behavior has been reproduced, such as an appropriate mass at cell division or death of the cell at the appropriate stage in the cell cycle.

 $\xrightarrow{accept}$  The accept stage is an assertion that the model adequately represents the biological process and consists of final preparations for archiving and disseminating the model.

The remaining stages in the original modeling process are error recovery stages. Errors are detected by informal examination of time series plots. The modeler must infer the nature and location of the error from experience. Because the developed models are typically underspecified, the modeler cannot always accurately identify the cause of an error. Laboratory experiments can test hypotheses, but are extremely expensive. In general, the goal of the modelers is to match the existing collection of laboratory results.

 $\xrightarrow{redesign}$  The redesign stage corrects errors in the wiring diagram. Reactions are added and deleted based on the modeler's developing intuition about what mechanisms must be included in the model to adequately reproduce the desired experimental behavior.  $\xrightarrow{repair}$  The repair stage corrects errors made in the translation between wiring diagram and differential equations. Manually creating differential equations is time consuming and prone to error. Tedious checking between the wiring diagram and differential equations is required to detect and correct errors in translation.

 $\xrightarrow{refit}$  The refit stage corrects errors made in the assignment of differential equation parameters. New estimates are made for the kinetic rate constants based on comparison with known experimental results. The modeler typically changes only a small number of rate constants at each iteration due to their potential interactions.

 $\xrightarrow{restart}$  The restart stage is the termination of a particular model and marks the start of the next idea of how a biological process is carried out.

## 4 APPLYING A METHODOLOGY

The original modeling process has successfully developed models that define the current state of the art. However, the modeling community recognizes that they are at the limit of the complexity their current methodology can support, which is driving many new efforts in tool development such as BioSPICE. Methodologies assist in understanding the model development process, and indicate requirements for supporting that process [6]. Formal methodological approaches provide well-defined and tested techniques for the model development process.

Based on our experience with the original modeling process, we enumerate some capabilities that a methodology needs to support this modeling community. Modelers have an ultimate goal of producing models that are validated and accepted. Demonstrating that their models are valid and should be accepted requires performing verification, validation, and testing (VV&T). Modelers should employ VV&T frequently to minimize wasted effort on bad models. Models developed by the original modeling process have proved extremely long-lived and are repeatedly adapted to meet changes in specification. We expect this to continue and so require a modeling process that is capable of introducing change at any stage without undue cost. Computational technology is also expected to change significantly during the lifetime of a model; models and the modeling process need to be insensitive to the runtime host and adaptable to high performance computing techniques of the next 10 years. Using the terminology of [30], our modeling process must primarily support correctness and testability, secondarily support adaptability, maintainability, and portability, and test throughout the model lifecycle.

We select the conical methodology [29], as a methodology supportive of our requirements and with sufficient adaptability to capture our original and desired modeling processes. [4] describes a model lifecycle compatible with the conical methodology; we will use similar terminology to describe our modeling process. The primary objectives of the conical methodology are correctness, testability, adaptability, reusability, and maintainability [30]. This is a good match with our primary and secondary requirements. The conical methodology prescribes a top-down model definition phase followed by a bottom-up model specification phase. As we are creating a domain-specific MSE, we significantly reduce the amount of work required in the definition phase by predefining constructs in our tools.

#### 4.1 Goals for Improving the Process

After we documented the original process, examined methodological frameworks, and listened to the concerns of modelers, we identified four areas for which the modeling process needs to be improved: documentation, testing, standardization, and automation. These four areas are important for developing models quickly and accurately. Independent verification and validation are testing activities performed by someone other than the model developer with the goal of improving the quality of the model [2]. Independent testing reduces potential modeler bias in evaluation, promotes earlier error detection, reduces error cost, and enhances operational correctness. We want to introduce independence into the modeling process at each iterative cycle with the goal of supporting independence for all testing activities. This level of support requires significant advances in the four outlined areas. We believe that making these improvements will ultimately lead to an increased rate of model accreditation and acceptance.

The goal of documentation is to record critical information about the modeling process. Model documentation is needed at every stage of the modeling process and is critical for future planning of modeling tasks. We want to record the model itself each time the description of the model is transformed. We want to record the procedure used for testing to support automated testing and review of VV&T methods. We want to record the results from testing for presentation and for comparison against future tests.

Comparison with experimental data has been the main testing technique used for validating these models. However, the quantity and quality of experimental data available for a particular system may be limited. It is a major expense to conduct new laboratory experiments for further testing or expansion of the model. In contrast, modeler time is relatively cheap. We emphasize verification during model construction to prevent the introduction of errors that strain our limited testing resources. The goal of testing is to introduce VV&T activities as soon as possible into the modeling process and to continuously monitor for introduced errors. We codify several indicators of model credibility [3] as automated tests that can be performed continuously during model development. When working with the wiring diagram, we want to verify that the graph structure of the diagram corresponds to the modeler's understanding of the structure. We want to verify that the names of species, rate laws, and constants are used consistently across the diagram. When building the executable model, we want to verify that the simulator can properly execute our model and that all required information is available. We want to perform the primary testing activities from a recorded plan that can be defined by an agent independent of the design or specification teams.

The goal of standardization is to adopt uniform notations and processes that reduce burdens on communication and development. For our modelers, a wiring diagram is the initial abstract representation for the model. Unfortunately, no standards exist for the graphical language of wiring diagrams, though the representation of Kohn [23] is becoming increasingly popular. The pathway modeling community is currently involved in standardizing the Systems Biology Markup Language (SBML), an XML-based representation of models at the chemical reaction level [18, 21]. While it is hoped that SBML will apply to wiring diagrams in the sense that model editing tools should be able to convert between SBML and wiring diagram representations, SBML files are not meant to be directly edited by modelers. The main purpose of SBML is to facilitate model exchange between modeling groups, who will then load the models into ad hoc editing tools. Biological models can have operations repeated across multiple reactions and incorporate subcomponents developed as part of other models. Sufficient support for interchanging model fragments would allow replacing both with well-tested black-box subcomponents.

We can also employ standardization at each stage of the modeling process by using domain-specific information to construct uniform sequences of tasks. A uniform process reduces the developmental tasks required of modelers and can prevent some errors in the planning stage of the modeling process.

Most of the existing modeling process consists of work that is performed repeatedly. The goal of automation is to have the computer perform some of these repetitive tasks, and speed up other tasks substantially. VV&T activities exist as automatable tasks throughout the modeling process; supporting these activities with automated tools can significantly reduce the time and effort for model testing [5]. Modelers repeatedly modify parameter and initial condition sets, at each modification comparing the revision against experimental data. We want to perform regression testing as frequently as possible and repeat testing activities from previous iterative cycles to ensure that model quality is maintained after each model transformation. Our testing activities should be numerous and specific so that when an error is introduced, we can identify what stage the error was introduced at and in what component of the model the error is located. When the user is engaged in modifying the model, the testing process should be conducted automatically, and feedback on the relative performance of the modification should be supplied automatically.

### 4.2 Revised Modeling Process

The revised modeling process, Figure 3, begins with an already defined problem. This problem definition includes an analysis of requirements and an identification of modeling objectives. We must make an assumption here that our modeling tools are adaptable to the solution technique chosen as part of the problem definition. Domain specificity allows us to make this assumption: the tools were developed specifically to meet the needs of biological modelers who have a large class of problems of interest.

From the problem description, modelers begin to develop model ideas that they believe will satisfy the problem requirements. The process of realizing and testing these ideas is extended from the original modeling process.

 $\xrightarrow{design}$  Starting with a conceptual model for a biological process, the modeler must first produce a model that can be understood by others. In addition to the wiring diagram or reaction equations, a complete model requires rate laws, constants, and the discrete event model that will control switches in the differential equation model. Models at this stage can be structurally tested and checked for completeness and consistency of kinetic information.

 $\xrightarrow{translate}$  The model must then be translated from a humanunderstandable form to an executable form. For the domain of biological modeling, we possess a significant amount of information about this transformation process. A sufficiently described model translates by a mechanical process. The modeler only needs to tweak the control parameters for the evaluation process. However, the model must still be verified to ensure that the model was sufficiently described before translation, self-consistent, and tolerant of the numerical errors to which the chosen simulation process is susceptible.

 $e^{valuate}$  After a model is simulatable, the modeler tests it against the requirements and objectives in the problem definition. VV&T activities in the evaluate stage likely account for a significant portion of the model development time and should have computer assistance to automate the test process [5]. Entering the problem definition only once and automating the testing process from previous iterative cycles of the modeling process are important for reducing model development time. Additionally, there should be provisions for independent execution of the evaluate stage to prevent modeler bias in the testing process.

 $\xrightarrow{check}$  A model that meets the requirements and objectives stated in the problem might still be rejected. In the biological domain, two reasons this occurs are that the model is insufficiently based on established biological processes or that the model is not significantly better than an existing, simpler model. The check stage addresses these issues. Comparing the proposed model against accepted models [42] representing similar processes can test the first. Performing a statistical analysis between the proposed model and a collection of models for the same system can test the second. Both techniques test the model against other models of biological systems.  $\xrightarrow{accept}$  and  $\xrightarrow{reject}$  The accept stage is relatively unchanged from the original modeling process. We formalize the preparations in the original accept stage as the process of creating documentation and presentations to show that the model is sufficiently accurate for its intended purpose [4]. Additionally, we introduce the reject stage for models that pass all of our tests but are rejected by decision makers.

The testing phases of the modeling process have been considerably augmented from the original process. Models are tested at every stage that creates or transforms a recorded model description. Additionally, the results of model testing more specifically point to causes of errors and direct the modeler to an appropriate stage for correcting the error.

 $\stackrel{test}{\longrightarrow}$  Test stages represent VV&T activities that take place in real time during model development. Continuous verification is especially important when the amount or quality of experimental data is limited. A test stage operates concurrently with tool use and indicates errors found in model information entered in the tool. Errors found and corrected in a test stage do not propagate to other stages of the modeling process. This reduces the amount of time to correct the error and reduces unnecessary switching between tools.

 $\xrightarrow{redesign}$  and  $\xrightarrow{repair}$  We have condensed the error recovery stages to redesign and repair, which correspond to activities that correct errors detected by validation and verification, respectively. When describing the software that implements this revised modeling process, we will once again enumerate specific types of error recovery activities.

The location, scope, and frequency of testing activities are the most significant difference between the original and revised modeling processes. Modelers get immediate feedback about errors detected during the process of transforming the model from one form to another. For errors that cannot be automatically detected, we attempt to identify more specifically the source and type of error so that the modeler spends less time diagnosing the problem. We hope that increasing the specificity of error reporting leads to a smaller average error recovery time. Moreover, even though we cannot always automatically detect where and how an error was introduced, we can automatically perform modeler-defined diagnostics to determine that there definitely was an error at some point in the last iteration of the process.

## 5 JIGCELL

JigCell is a domain-specific MSE for biological pathway modeling, intended ultimately to become a problem solving environment (PSE) in the sense of [32] and [39]. JigCell's user workflow, Figure 4, corresponds closely with the modeling process we have identified. Table 2 lists support for our defined modeling goals in this MSE.

We have constructed a tailored environment rather than basing it on an existing, general-purpose MSE [7, 26]. We intend



Figure 3: Revised Modeling Process

to support users in biology and related fields who do not have significant experience in formal modeling (but who are domain experts). During the development of JigCell, we supplied the software to biologists and modelers for testing. As well as finding errors, they gave us feedback about new features and priorities for development.

We incorporate off-the-shelf components such as numerical libraries, visualization tools, and communications protocols where quality implementations exist and technical specifics about the component can be hidden from the user. This approach has not been a significant drawback: the majority of development work relates to domain-specific support rather than modeling infrastructure.

Development versions of JigCell are freely available at [1]. Official releases of JigCell are available as part of the BioSPICE releases [8].

## 5.1 JigCell Tools

The **Model Builder** creates a model specification that incorporates the wiring diagram, kinetic information, and discrete event model. A spreadsheet interface organizes the information of the wiring diagram and kinetics as a collection of chemical reaction equations. Each row of the Model Builder spreadsheet (Figure 5) specifies a chemical reaction equation including substrates, products, kinetic rate law, and kinetic rate constants. Chemical equations are a natural representation for many biological processes of interest and are applicable to a wide variety of fields outside biological modeling.

Restrictions are placed on the class of discrete event models: events can only be triggered based on algebraic conditions of

le <u>D</u> atabase <u>O</u> ptions	Help					
Mechanisms	#	Reaction	Name	Туре	Equation	Parameters
🛛 🗂 New Mechanism	1		٧	Function	((A1*A2)+(A3*A4))	
🛛 🗂 C:\JigCell\model	2	1	Michaelis-M	New	(((k1*S1)*M1)/(J1+S1))	
🗉 🗂 Model	3	1	Mass_Action	New	kf	
🗉 🗖 Compartmer	4		Mass_Action	New	(kf*S1)	
Parameters	5	1	Mass_Action	New	((kf*S1)*S2)	
Species	6	-> CycB		Mass_Action	k1	kf=k1
	7	CycB->		Mass_Action	(Cyclosome*CycB)	kf=Cyclosome
	8	CycB + Cdk1 -> MPFa		Mass_Action	((k3*CycB)*Cdk1)	kf=k3
1	9	MPFa -> MPFi		Mass_Action	(V(kwp,Weei,kwpp,Weea)*MPFa)	kf=V(kwp,Weei,kwpp,Weea)
	10	MPFi-> MPFa		Mass_Action	(V(kcp,C25i,kcpp,C25a)*MPFi)	kf=V(kcp,C25i,kcpp,C25a)
1	11	C25a -> C25i		Michaelis-M	(((k25r*C25a)*1.0)/(J25r+C25a))	M1=1.0; J1=J25r; k1=k25r
	12	C25i -> C25a		Michaelis-M	(((k25f*C25i)*((MPFa+eps1)))/(J25f+C25i))	M1=(MPFa+eps1); J1=J25f; k1=k25f
1	13	Weea -> Weei		Michaelis-M	(((kweef*Weea)*((MPFa+eps2)))/(Jweef+Weea))	M1=(MPFa+eps2); J1=Jweef; k1=kwee
	14	Weei -> Weea		Michaelis-M	(((kweer*Weei)*1.0)/(Jweer+Weei))	M1=1.0; J1=Jweer; k1=kweer
	15	Csoma -> Csomi		Michaelis-M	(((kcyr*Csoma)*1.0)/(Jcyr+Csoma))	M1=1.0; J1=Jcyr; k1=kcyr
	16	Csomi -> Csoma		Michaelis-M	(((kcyf*Csomi)*((MPFa+eps3)))/(Jcyf+Csomi))	M1=(MPFa+eps3); J1=Jcyf; k1=kcyf
	17	MPFa -> Cdk1		Mass_Action	(Cyclosome*MPFa)	kf=Cyclosome
	18	MPFi -> Cdk1		Mass_Action	(Cyclosome*MPFi)	kf=Cyclosome
Income in the local data	19	Cyclosome		Species	V(k2p,Csomi,k2pp,Csoma)	k2pp=k2pp; k2p=k2p

Figure 5: Model Builder Interface with Equations from Wiring Diagram

general	record model artifacts	Partial
	record model testing procedure	Full
	record model testing results	Full
design	verify consistent naming	Partial
	VV&T model structure	
	abstract repeated operations	Partial
	abstract compound operations	
	identify rate laws and constants	Full
	standardize notation	Full
translate	generate differential equations	Full
	find conservation relations	Full
	verify execution requirements	Full
	verify well-formedness	Partial
evaluate	make VV&T repeatable	Full
	support independent VV&T	Full
	identify causes of errors	Partial
	standardize process	Full
check	make covalidation repeatable	Partial
	support independent covalidation	
	identify causes of errors	Partial
	standardize process	Partial

species values and can only modify parameters, constants, and species values in the continuous model. However, these discrete models are sufficient for the biological systems we have studied and are easily and directly created by domain experts without modeling experience. For example, the discrete model can easily handle such processes as cell division: mass at time  $t + \Delta t$  is set to one-half the mass at time t when a particular species crosses a threshold value. The Model Builder both reads and writes its models in the form of SBML, which is becoming the standard interchange language for this modeling community.

Species names and kinetic information are checked continuously during model entry with color highlights indicating portions of the model that are not correctly specified. No mechanisms are included for testing overall model structure and limited support is provided for representing stochastic and spatial models, which represent specific sub-domains within the broader modeling community that we intend to support in the future. Division of the modeled cell into multiple topological compartments (volumes) is possible, but equations cannot contain spatial variables. Abstractions are possible in the sense that rate laws can be defined and reused. However, components in the form of black-box submodels are not currently supported since the community (in terms of the SBML standardization effort) has not yet defined mechanisms for this.

The Run Manager translates a model specification into an executable form. Each row in the Run Manager spreadsheet (Figure 6) specifies how to simulate a certain experiment including the model to use, parameter and initial condition sets, and the appropriate simulator settings. Parameter sets can contain a value for every parameter in the model or contain only the values changed in relation to another parameter set. For example, suppose we wish to describe a series of simulations for yeast cell mutants. On the first line, we describe how to simulate the wild-type yeast cell using a basal parameter set. On the second line, we describe a knockout mutant by using a parameter set that copies from the basal parameter set but changes the rate of synthesis of the knocked-out protein to zero. On the third line, we describe a mutant with a double copy of a gene by using a parameter set that copies from the basal parameter set and multiplies the synthesis rate associated with the gene

🛃 JigCell	RunManager - C:	¥JigCell¥run	s¥frogeggpaper.run					E	
File Opt	tions Help								
Runs	Basal Settings								
🗖 Run (	Groups lew Set of Runs	#	Name	Comments	Model	Parameters	Initial Conditions	Simulator Par	ame
• 🗖 N		1	Xenopus laevis oocyte		C:\JigCell\models\frogeggpaper.sbml	basal	active Wee	xpp 150	•
🕈 🗖 C	:UigCell\runs\frc	2	Xenopus laevis oocyte (extend	maxt=1000	C:UigCell\models\frogeggpaper.sbml	basal	active Wee	xpp 1000	
- C	Runs	3	Xenopus laevis oocyte (MPF st		C:\JigCell\models\frogeggpaper.sbml	basal	initial MPF 0.5	xpp 150	
		4	Xenopus laevis oocyte (eps fa		C:\JigCell\models\frogeggpaper.sbml	eps1	active Wee	хрр 150	
		5	Xenopus laevis oocyte (no kW		C:\JigCell\models\frogeggpaper.sbml	K/Veer=0	active Wee	xpp 150	
		6	Xenopus laevis oocyte (no KW		C:WigCell\models\frogeggpaper.sbml	K/Veer=0	initial MPF 0.5	xpp 150	
		7							
		8							
		9							
		10							
		11							
		12							
		13							
		14							
		15							
		16							
		17							-
	Þ	•	-			4			

Figure 6: Run Manager Interface for Model in Model Builder

by two. If the modeler then develops a better estimate for the rate constants, she only needs to change the basal parameter set. Modelers can more easily explore the parameter space because the Run Manager does the bookkeeping to adjust the simulation specification.

Differential equations and code to handle the discrete event model are automatically generated from the specification. Modelers select the dependent variables of conservation relations, but the relations are automatically detected and generated. Automatic model translation insulates modelers from changes in simulation techniques and the runtime environment. New simulators are added by providing additional translators. This automation step represents a major improvement for our immediate modeling community, who previously converted reactions by hand in an error-prone and laborious process. This should also reduce the amount of model conversion work required of modelers to perform stochastic or spatial simulations in the future.

The **Comparator** and **Compare**<sup>2</sup> are tools for model testing and evaluation. Tests in the Comparator are assertions about a model or comparisons between model performance and experimental data. A test evaluates either operational accuracy or the accuracy in transforming the model. Performance on each test is scored according to a user-defined objective function that represents the goodness-of-fit between the expected result and the model result. We prefer objective functions that associate model performance with a degree of accuracy rather than a binary result. If the modeler is only told whether the result was acceptable, it is difficult to determine how robust the model is to parameter changes. Moreover, we prefer continuous objective functions, such as distance functions, as these are more amenable to automated optimization. The modeler chooses criteria for the objective functions based on the requirements and purpose of the model. Integrated editors support defining assertions, experimental data, procedures for transforming model results, and objective functions. Since we cannot predict all of the objective functions the modeler might wish to use, we allow the modeler (possibly with the assistance of a computer programmer) to add new functions. The modeler can automatically rerun a defined testing procedure in the Comparator.

The Comparator interface is a tabbed series of screens (Figure 7) that describe the testing procedure. First, the modeler enters the experimental, or expected, results that the model should reproduce. Then, the modeler associates each experimental result with a procedure for generating an equivalent result from the model. The Run Manager typically executes the model. Finally, the modeler describes how to compute an objective function and specifies what constitutes an acceptable function value. Unacceptable fits are highlighted in the display so that the modeler can quickly see where the model is having problems.

Tests in Compare<sup>2</sup> compare performance between the currently proposed model and a collection of other models. Models come from past revisions of the current model, independent models of the same system, and models with subsystems in common with the current model. Compare<sup>2</sup> performs ranking and selection among these models based on the same criteria defined in the Comparator. For each objective function defined in the Comparator, models are scored and ranked. We use rankings rather than absolute results as we have found that many objective functions were not designed to finely distinguish results on an absolute scale. Selection of the best model is made by a function combining the objective functions, such

III JigCell Comparator							o≝ Ø"	×
Comparator File Edit Run								
	86 49 D H							
Experiment Model Objective	Transform Editor	Transform Debugger	Objective Editor Comp	are^2				
Name	Experimental Val	Transform Value Objecti	ve Objective Result	Criterion	Acceptable	Value Type	Comment	
Xenopus laevis oocyte active MPF p	. ((50, 0.3), (85, 0 ((	).0, 0.0), (0.1, 7.0 WOSS	Infinity	Ignore		Time series	active MPF peaks high	
Xenopus laevis oocyte active MPF p	. ((55, 0.05), (90, 0 ((	).0, 0.0), (0.1, 7.0 WOSS	3.8578036721016306E-4	<	0.01	Time series	active MPF peaks low	
Xenopus laevis oocyte inactive MPF	((40, 0.4), (75, 0 ((	).0, 0.0), (0.1, 2.5 WOSS	3.1796366113578787E-4	×	0.01	Time series	inactive MPF	
Xenopus laevis oocyte active Wee	((45, 0.3), (80, 0 ((	).0, 1.0), (0.1, 0.9 WOSS	0.9364071984995537	<	0.5	Time series	active Wee	
Xenopus laevis oocyte active Wee (	((1, 0.0), (10, 0.0) ((	).0, 1.0), (0.1, 0.9 WOSS	0.8794764131245917	>=	0.5	Time series	active Wee (no KWeer)	
Xenopus laevis oocyte active MPF (	((50, 0.25), (55, 0 ((	).0, 0.0), (0.1, 7.0 WOSS	0.0037263235427044024		0.01	Time series	active MPF (no kWeer)	
Xenopus laevis oocyte active MPF (i	. ((50, 0.3), (85, 0 ((1	0.0, 0.5), (0.1, 0.4 WOSS	0.17419495751269864		0.01	Time series	active MPF (MPF=0.5)	1
Xenopus laevis oocyte active MPF (i	. ((50, 0.25), (55, 0 ((	).0, 0.5), (0.1, 0.4 WOSS	0.0037263235427044024	<	0.5	Time series	active MPF (MPF=0.5,no K	1
Xenopus laevis oocyte active Wee (	((1, 0.0), (10, 0.0) ((	).0, 0.5), (0.1, 0.1 WOSS	2.1781869437712366E-32	E	Equivalent	Time series	active Wee (MPF=0.5,no K	V.

Figure 7: Comparator Interface for Model in Model Builder

as a sum of rankings. Development work on Compare<sup>2</sup> is still ongoing. A drawback of this tool is that obtaining meaningful results requires building a significant collection of models. This is also a limiting factor on development since, without test models, it is unclear what tasks are most critical to automate. We hope to incorporate other automated model analyses that could reduce the startup costs of this tool.

By automating the comparison process, we make an additional task possible: parameter estimation. Some rate constants used in the continuous model are not experimentally determined or have a significant range of possible values. Without automated fitting, the modelers must calibrate the model parameters by manually searching for valid and optimal regions of the parameter space. An expert in the biology and mathematics of the model must repeatedly try parameter guesses, using experience to determine acceptability and to select the next guess. This activity has consumed a major part of the model development time in the past. Humans should not perform it at all. A significant obstacle is that inputting the domain expert's intuitive understanding of the model into a computer system is difficult. Additionally, when the model calibration procedure involves more than simple curve fitting, there are few general-purpose techniques. We seek to overcome these obstacles by customizing the domain-specific portions of our software to modelers, and by selecting mathematical software that has been tested on common problems in the domain. Experienced modelers can extend or replace the interface and computation engines when they encounter limitations with what we have preconstructed.

Although the stages involving parameter estimation included in Figure 4 are a subset of the evaluate and repair stages, we separate them because of their impact on the model

development process.

 $\xrightarrow{score}$  and  $\xrightarrow{report}$  The score stage defines an algorithm that determines whether one set of parameters produces a more acceptable model than another set. The algorithm requires experimental data, an executable model, the range of parameters for the executable model, and a user-defined objective function. The report stage injects the fitted parameters back into the modeling process for study and testing.

The **Parameter Estimator** finds unknown rate constants by fitting the model to experimental data. The data are typically not a solution to a differential equation, but rather a complicated, nonlinear functional of the differential equation solution. Furthermore, both the dependent and independent variables involved in these functionals are subject to experimental error. The Parameter Estimator performs both global and local searches during optimization.

The global optimizer, named DIRECT [14], is a variant of Lipschitzian methods for constrained global optimization. We want to find the minimum value of our objective function,  $\min_{x \in D} f(x)$ , where  $D = \{x \in E^n | \ell \le x \le \mu\}$  is defined by bound constraints only and the objective function is Lipschitz continuous on D, satisfying  $|f(x_1) - f(x_2)| \le L ||x_1 - x_2||$  for all  $x_1, x_2 \in D$ . The Lipschitz optimization method has had many practical applications in science and engineering. Unlike some other methods, the Lipschitz method requires only a few parameters and does not rely on derivatives or other more analytical information about a system. However, the Lipschitz constant of a particular function is often unknown and difficult to estimate. The DIRECT method is guaranteed to converge to the global optimum without knowledge of the Lipschitz constant [22].

The algorithm takes its name from one of its key steps:



Figure 4: JigCell User Workflow

dividing rectangles (these rectangles are more commonly referred to as boxes). DIRECT is a pattern search method, which takes moves based on objective function values at a pattern of points. The points are the centers of the boxes. Center sampling is generally advantageous to corner sampling when the number of dimensions is expected to be large, as in the problems we are attempting to solve. A box is potentially optimal if there exists a value of the Lipschitz constant for which that box is the most likely to contain the global minimum. Each iteration subdivides all of the potentially optimal boxes. DI-RECT can operate in an exploratory mode, which emphasizes searching untested boxes, or in an exploitation mode, which emphasizes searching boxes with better objective function values.

DIRECT is robust to noise in objective function values [14]; convergence to the minimum is limited by the amount of noise in the objective function. This makes DIRECT well suited for stochastic models, which intentionally introduce noise into the model evaluation. Current implementations of DIRECT cannot handle integer variables or constraints other than simple bound constraints. Also, DIRECT is relatively inefficient for finding an accurate value of the minimum. Rather, we would expect to run it in the exploration mode and use the local optimizer to find the minimum from a collection of candidate starting points.

The local optimizer uses ODRPACK [9, 10, 11] as the underlying mathematical software. ODRPACK does not assume that the measurement errors are all in the dependent variables [12]. Rather, it seeks to minimize the weighted sum of orthogonal distances between the model and the data. The weighting factors scale the residuals and express the modeler's confidence in the reliability of particular observations. The output of ODRPACK gives a locally optimal parameter vector and a measure of the goodness-of-fit of the parameter vector. We can then compare the locally optimal solutions for the starting points picked by the global optimizer.

ODRPACK uses a trust region Levenberg-Marquardt method. The Levenberg-Marquardt method starts with the steepest descent method and smoothly changes to Newton's method when approaching the solution. The trust region implementation determines the step size based on the confidence in a local model of the objective function. At each step, the optimizer compares the expected improvement for taking the step with the actual improvement. This can cause slow convergence if the objective function is not differentiable near the optimum solution, since ODRPACK's expected improvement is based on estimated derivatives. Step functions in objectives can appear when matching categorical observations, such as whether a mutant is viable. However, we believe it is unlikely that optimum solutions will be located near such discontinuities in real biological systems. If that were the case, the organism would be sensitive to minute environmental changes, which is unfavorable for survival.

## 5.2 Evaluating JigCell

Several levels of evaluation are possible for MSEs. Microlevel studies employ formal usability testing [17, 34], which benchmarks performance for completing a task. Requirements of the modelers, frequency of use within the domain, and criticality of need determine the chosen tasks. An example of a critical task in the Model Builder is entering the kinetic information for a chemical reaction. The Model Builder could not function without supporting this task. Success or failure is determined by comparing results against a benchmark performance for the task. In formal usability testing, [31] shows that studying three to five people finds eighty percent of the usability problems.

A micro-level study [38] determined JigCell's effect on error rates converting the wiring diagram to a set of differential equations. For a collection of models, participants either constructed differential equations manually or using the software. The number of errors in the generated sets of differential equations was then measured. The results indicate a six-fold reduction in errors over the manual method of creating differential equations. Individual tools in JigCell as well as interactions between tools can be studied similarly.

Macro-level studies incorporate benchmarking and assess how well the MSE meets the specified needs of users. MSEs such as JigCell attempt to make knowledgeable users more productive and help them produce creative products. [37] suggests a methodology for evaluating MSEs using flow. Flow is an automatic, effortless, and focused state of consciousness. Creativity is more likely to result from flow states [13]. Detection of flow would indicate support for creativity.

We can evaluate flow in JigCell in the following manner. By specifying assistance for each workflow of [37], we classify support for problem solving and flow in JigCell. Users receive an instrumented version of JigCell; this version also sporadically prompts users to fill out a questionnaire targeting the characteristics of flow [41]. The collected data determines where flow is occurring in the user's work. If JigCell demonstrates flow in all workflows and meets the formative evaluation requirements, then JigCell indicates support for creativity.

### 5.3 Other Cell Cycle MSEs

We have reviewed several biochemical pathway MSEs and classify them as wizards, graphical editors, or textual editors. JigCell has primarily text-based editors. Table 3 lists support for our defined modeling goals in these MSEs. Modeling goals met by none of the MSEs are omitted; the evaluate and check stages are omitted entirely as none of the reviewed MSEs directly supports these processes.

Virtual Cell [24] is a graphical editor that directly mimics the wiring diagram. Virtual Cell also supports entering models directly as differential equations. The Virtual Cell simulator supports volume and spatial models. A server-side database records models and associated simulation results with the option to download data to the user's system. Bio Sketch Pad (BSP) [40] is also a graphical editor. The simulator for BSP supports volume models and cellular automata. Gepasi [27] uses a wizard interface. A series of dialog boxes lead the user through creating a reaction network. The simulator supports volume models and can perform parameter estimation. Jarnac [33] uses a textual interface. The textual editor is a programming environment for scripting simulations and other modeling tasks. Jarnac integrates with a graphical editor, JDesigner, to view wiring diagrams. The simulator supports volume models.

## 6 TOOL INTEROPERABILITY

Tools for doing biochemical pathway modeling have been around for many years, but until recently there has been little interoperability between them. We describe two projects that have taken as their goal support for interoperability between various biochemical pathway modeling groups. The first is SBML, and the second is DARPA's BioSPICE project (which funds the JigCell project team).

SBML [18] is an XML-based language for describing biochemical pathway models. SBML does not use a formal standards process (though the SBML community is beginning to investigate affiliating with a standards organization). The Software Platforms for Systems Biology Forum, which has been meeting semiannually since April 2000, guides development. The SBML project began with the Forum's first meeting, and the definition for SBML Level 1 was published in 2001 [20]. That version supported a relatively small number of biochemical pathway modeling tools. An updated version, SBML Level 2, was published in 2003 [19]. Another language effort similar to SBML is CellML [15]. CellML is also an XML-based markup language for describing biological systems. CellML's scope is broader than SBML. While it is somewhat more mature than SBML, it has not achieved much acceptance within the biochemical pathway modeling community in the way that SBML has, probably due to its broader focus.

Before SBML, there was virtually no interchange of models between working groups. Tools were primarily built for a particular research group with little intention that the models developed within a tool would be directly transferable to other groups with other tools. Models were exchanged only through publication. With the adoption of SBML, and the beginning stages of support for that language by many tools, the true exchange of models became possible. However, early tool builders primarily supported SBML through export capabilities, and were not able to import SBML models produced by other research groups. SBML initially suffered from an inability to express a number of existing models, leading tool builders to adopt naming and formatting conventions to handle their special cases (see for example [28]). Since these conventions were unique to each tool, few tools supported actual exchange of models. The SBML community (and by extension, the BioSPICE community) now actively promotes a model testbed and the exchange of models between working groups with the goal of more rigorous testing of both modeling tools and their SBML language support.

An important standards-setting decision for BioSPICE was determining the model definition language to support interchange of models within the community. After examination of the potential for defining a language for BioSPICE, the decision was made to adopt SBML. The critical mass achieved by the BioSPICE community joining with the already existing SBML community has had a profound impact on the SBML development process. Initially, SBML could be characterized as defining a language that was the intersection of a small group of similar tools. Most users of SBML converted their models to systems of ODEs, perhaps with discrete events. Current efforts are moving in the direction of defining a language that is the union of a larger group of more disparate tools. This means support for a broader class of models, including stochastic simulations and spatial models. Structural model analyses such as flux balance analysis and bifurcation analysis are also anticipated. SBML is already at the stage where no single simulation tool will support the full range of features that SBML models can have. Some balance must be struck between the different types of modeling needs, the difficulties encountered by tool builders, and the risk of the language

		Virtual Cell	BSP	Gepasi	Jarnac
general	record model artifacts	Partial	Partial	Partial	Partial
design	verify consistent naming	Partial	Partial	Partial	Partial
	abstract repeated operations				Full
	abstract compound operations				Full
	identify rate laws and constants	Full	Full	Full	Full
	standardize notation	Full	Partial	Full	Full
translate	generate differential equations	Full	Partial	Partial	Partial
	find conservation relations	Full	Partial	Full	Full
	verify execution requirements	Full	Full	Full	Full
	verify well-formedness	Partial	Partial	Partial	Partial

Table 3: Support for Modeling Goals in Other MSEs

becoming so large that it is no longer truly a mechanism for common exchange.

The BioSPICE community charged itself with designing and building a toolset that could both support current modeling efforts and serve as a base for future modeling efforts. Since it is unlikely that any single tool, or small collection of tools, will satisfy the wide variety of modelers and modeling efforts, any specific task has numerous tools available. Thus, the BioSPICE community is engaged in two types of standardization efforts for tools: data formats and programmatic interfaces.

The exchange of formatted data within the BioSPICE community is difficult because a large number of data types are of interest to at least some part of the community. Popular data types, such as models and time series, were quickly identified and standardized by the community as high-priority items. However, there is also interest in exchanging data for gene expression, molecular interactions, imaging, protein mass-spectroscopy, kinetics, flow cytometrics, and flux balance. The standards creation process is driven by the community, and mostly distributed. There is a clearinghouse for archiving agreed-upon standards, but there is no corresponding authority for creating standards. Tools must have some internal format for representing data. Initially, this is their standard for that data format. If no one else is interested in using this type of data, there is no need to promulgate the standard. However, another group interested in using the same type of data will have their own internal format. At some point, there is a desire to interchange data between the tools. The groups then meet, and decide on a commonly acceptable interchange format. When a format achieves a plurality within its domain, it is recommended to the community as a standard for adoption.

Although data formats solve the problem of static interactions between tools, they cannot coordinate computations between tools. The number of possible tool-tool interactions is too great to define individual standards. Instead, standards are defined for the most commonly performed interactions, and a formal tool description language is proposed for describing tool capabilities not captured by one of the standards. Common interactions include tasks such as performing a simulation and parameter estimation. The tool description language is referred to as a meta-interface: it is intended that any particular tool be describable in terms of operations provided by the language. The operations are functions applied to the data types standardized by the community. However, this tool description language only specifies the syntax for tools, not the semantics. It is possible to connect tools that have compatible data types but a nonsensical meaning to their combined computation.

### 7 CONCLUSIONS

We have described our experiences documenting and improving the modeling process of a group of theoretical biologists. The revised modeling process is based on the observed process and incorporates a disciplined methodological approach along with proven techniques for reducing the cost of errors, reducing development time, and making iterations of the modeling process more consistent. The primary improvements come from introducing multiple forms of testing throughout the modeling lifecycle, and by making modelers explicitly address previously undocumented modeling tasks. The modeling support environment we have built for this revised process meets many of our defined modeling goals and is testable for its effects on the modeling process. Further case studies of modelers using this process will give clear guidance for future improvements, along with the unfulfilled modeling goals.

Further progress in pathway modeling desperately needs new tools and a better modeling process as described here. The BioSPICE community estimates that the size and complexity of existing models (which are about as large as the modeling process prior to BioSPICE could handle) must grow by two orders of magnitude to capture the control mechanisms of important processes in mammalian cells. Currently, manual parameter estimation and manual error diagnosis consume a significant portion of the model development time. We hope that by automating these tasks, tools such as JigCell and others developed as part of BioSPICE can supply one of the two needed orders of magnitude. Supporting the additional complexity required is an open question.

#### ACKNOWLEDGMENTS

Thanks are due to John Tyson and members of his modeling laboratory for making themselves available for study and reviewing our observations of their modeling process.

The work reported herein was partly sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory (AFRL), Air Force Materiel Command, USAF, under agreement number F30602-02-0572. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, or the U.S. Government.

#### REFERENCES

- [1] Nicholas A. Allen, Marc T. Vass, Jason W. Zwolak, et al. The eukaryotic cell cycle collaborative problem-solving environment group. Available at http://gnida.cs. vt.edu/~cellcyclepse/, 2003.
- [2] James D. Arthur and Richard E. Nance. Verification and validation without independence: a recipe for failure. In *Proceedings of the 32nd Conference on Winter Simulation*, pages 859–865, 2000.
- [3] Osman Balci. Credibility assessment of simulation results. In Proceedings of the 18th Conference on Winter Simulation, pages 38–44, 1986.
- [4] Osman Balci. Verification, validation, and accreditation. In Proceedings of the 30th Conference on Winter Simulation, pages 41–48, 1998.
- [5] Osman Balci, Richard E. Nance, James D. Arthur, and William F. Ormsby. Expanding our horizons in VV&A research and practice. In *Proceedings of the 34th Conference on Winter Simulation*, pages 653–663, 2002.
- [6] Osman Balci, Richard E. Nance, E. Joseph Derrick, Ernest H. Page, and John L. Bishop. Model generation issues in a simulation support environment. In *Proceedings of the 22nd Conference on Winter Simulation*, pages 257–263, 1990.
- [7] David W. Balmer and Ray J. Paul. Integrated support environments for simulation modelling. In *Proceedings of the 22nd Conference on Winter Simulation*, pages 243– 249, 1990.
- [8] BioSPICE. The BioSPICE development project. Available at http://www.biospice.org, 2003.

- [9] Paul T. Boggs, Richard H. Byrd, Janet R. Donaldson, and Robert B. Schnabel. Algorithm 676 - ODRPACK: Software for weighted orthogonal distance regression. ACM *Transactions on Mathematical Software*, 15(4):348–364, 1989.
- [10] Paul T. Boggs, Richard H. Byrd, Janet E. Rogers, and Robert B. Schnabel. User's Reference Guide for ODR-PACK Version 2.01: Software for Weighted Orthogonal Distance Regression. Center for Computing and Applied Mathematics, U.S. Department of Commerce, Gaithersburg, MD, 1992.
- [11] Paul T. Boggs, Richard H. Byrd, and Robert B. Schnabel. A stable and efficient algorithm for nonlinear orthogonal distance regression. *SIAM Journal on Scientific and Statistical Computing*, 8(6):1052–1078, 1987.
- [12] Ray J. Carroll, David Ruppert, and Lenny A. Stefanski. *Measurement Error in Nonlinear Models*. Chapman and Hall/CRC, New York, NY, 1995.
- [13] Mihaly Csikszentmihalyi. *The Psychology of Optimal Experience*. Harper & Row, New York, NY, 1990.
- [14] Jian He, Layne T. Watson, Naren Ramakrishnan, Clifford A. Shaffer, Alex Verstak, Jing Jiang, Kyung Bae, and William H. Tranter. Dynamic data structures for a direct search algorithm. *Computational Optimization and Applications*, 23:5–25, 2002.
- [15] Warren J. Hedley, Melanie R. Nelson, et al. CellML specification. Available at http://www.cellml. org/public/specification/20010810/, 2001.
- [16] Alan C. Hindmarsh. ODEPACK: A systematized collection of ODE solvers. In R. S. Stepleman, editor, *Scientific Computing*, pages 55–64. North Holland Publishing Company, 1983.
- [17] Debbie Hix and H. Rex Hartson. Developing User Interfaces: Ensuring Usability through Product & Process. John Wiley & Sons, Inc., New York, NY, 1993.
- [18] Michael Hucka, Andrew Finney, et al. SBML Systems Biology Markup Language. Available at http: //www.sbml.org/, 2003.
- [19] Michael Hucka, Andrew Finney, et al. Systems Biology Markup Language (SBML) Level 2: Structures and facilities for basic model definitions. Available at http://www.sbml.org/specifications/ sbml-level-2/version-1/sbml-level-2. pdf, 2003.
- [20] Michael Hucka, Andrew Finney, Herbert M. Sauro, Hamid Bolouri, et al. Systems Biology Markup

Language (SBML) Level 1: Structures and facilities for basic model definitions. Available at http://www.sbml.org/specifications/ sbml-level-1/version-2/sbml-level-1v2.pdf, 2001.

- [21] Michael Hucka, Andrew Finney, Herbert M. Sauro, Hamid Bolouri, et al. The Systems Biology Markup Language (SBML): A medium for representation and exchange of biochemical network models. *Bioinformatics*, 19(4):524–531, 2003.
- [22] D. R. Jones, Cary D. Perttunen, and Bruce E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.
- [23] Kurt W. Kohn. Molecular interaction map of the mammalian cell cycle control and DNA repair system. *Molecular Biology of the Cell*, 10(8):2703–2734, 1999.
- [24] Leslie M. Loew and Jim C. Schaff. The Virtual Cell: A software environment for computational cell biology. *Trends in Biotechnology*, 19:401–406, 2001.
- [25] Gabor Marlovits, Christopher J. Tyson, Bela Novak, and John J. Tyson. Modeling M-phase control in Xenopus oocyte extracts: the surveillance mechanism for unreplicated DNA. *Biophysical Chemistry*, 72:169–184, 1998.
- [26] S. C. Mathewson. Simulation support environments. In Computer Modelling for Discrete Event Simulation, pages 57–100. John Wiley & Sons, Inc., 1989.
- [27] Pedro Mendes. Biochemistry by numbers: Simulation of biochemical pathways with Gepasi 3. *Trends in Biochemical Sciences*, 22:361–363, 1997.
- [28] Pedro Mendes. Gepasi and SBML. Available at http: //www.gepasi.org/gep3sbml.html, 2002.
- [29] Richard E. Nance. The conical methodology and the evolution of simulation model development. *Annals of Operations Research*, 53:1–45, 1994.
- [30] Richard E. Nance and James D. Arthur. The methodology roles in the realization of a model development environment. In *Proceedings of the 20th Conference on Winter Simulation*, pages 220–225, 1988.
- [31] Jakob Nielsen. Heuristic evaluation. In Jakob Nielsen and Robert L. Mack, editors, *Usability Inspection Meth*ods, pages 25–62. Wiley and Sons, New York, NY, 1994.
- [32] Naren Ramakrishnan, Layne T. Watson, Dennis G. Kafura, Cal J. Ribbens, and Clifford A. Shaffer. Programming environments for multidisciplinary Grid communities. *Concurrency and Computation: Practice and Experience*, 14:1241–1273, 2002.

- [33] Herbert M. Sauro. Jarnac: A system for interactive metabolic analysis. In *Animating the Cellular Map: Proceedings of the 9th International Meeting on Bio-ThermoKinetics*, pages 221–228. Stellenbosch University Press, 2000.
- [34] Ben Schneiderman. Designing the User Interface: Strategies for Effective Human-Computer-Interaction. Addison Wesley Longman, Inc., Reading, MA, 1998.
- [35] Lawrence F. Shampine and Marilyn K. Gordon. Computer Solution of Ordinary Differential Equations, The Initial Value Problem. W. H. Freeman, San Francisco, CA, 1975.
- [36] John J. Tyson and Bela Novak. Regulation of the eukaryotic cell cycle: Molecular antagonism, hysteresis, and irreversible transitions. *Journal of Theoretical Biology*, pages 249–263, 2001.
- [37] Marc T. Vass, John M. Carroll, and Clifford A. Shaffer. Supporting creativity in problem solving environments. In *Proceedings of the Fourth Creativity & Cognition Conference*, pages 31–37, 2002.
- [38] Marc T. Vass and Pete Schoenhoff. Error detection support in a cellular modeling end-user programming environment. In *IEEE 2002 Symposia on Human Centric Languages and Environments*, pages 104–106, 2002.
- [39] Layne T. Watson, Vinod K. Lohani, David F. Kibler, Randy L. Dymond, Naren Ramakrishnan, and Clifford A. Shaffer. Integrated computing environments for watershed management. *Journal of Computational Civil Engineering*, 16:259–268, 2002.
- [40] Johnathan Webb. Bio Sketch Pad. Available at http:// bio.bbn.com/biospice/biosketchpad, 2002.
- [41] Jane Webster, Linda K. Trevino, and Lisa Ryan. The dimensionality and correlates of flow in human computer interactions. *Computers in Human Behavior*, 9(4):411– 426, 1993.
- [42] Sammuel A. Wright and Kenneth W. Bauer. Covalidation of dissimilarly structured models. In *Proceedings of the* 29th Conference on Winter Simulation, pages 311–318, 1997.

# **AUTHOR BIOGRAPHIES**

NICHOLAS A. ALLEN is a PhD candidate in the Department of Computer Science at Virginia Tech (VPI&SU). He received BS degrees (magna cum laude) in Mathematics and Computer Science from Virginia Tech in 1999, and MS degrees in Mathematics and Computer Science from Virginia Tech in 2001. His research interests include software architecture and design, distributed computing, graph theory, and mathematical software. His email address is <nallen@ acm.org>.

**CLIFFORD A. SHAFFER** is an associate professor in the Department of Computer Science at Virginia Tech since 1987. He received his PhD from University of Maryland in 1986. His current research interests include problem solving environments, bioinformatics, component architectures, visualization, algorithm design and analysis, and data structures. His email address is <shaffer@vt.edu>.

NAREN RAMAKRISHNAN is an assistant professor of computer science at Virginia Tech. His research interests include problem solving environments, mining scientific data, and personalization. Ramakrishnan received his Ph.D. in computer sciences from Purdue University. Contact him at naren@cs.vt.edu.

MARC T. VASS is a PhD candidate in the Department of Computer Science at Virginia Tech. He received the BS degree in Computer Science from Virginia Tech in 2000 and the MS degree in Computer Science from Virginia Tech in 2001. His research interests include creativity and cognition, problem solving environments, flow, and theoretical human computer interaction. His email address is <mvass@vt.edu>.

LAYNE T. WATSON is a professor of computer science and mathematics at Virginia Tech. His research interests include fluid dynamics, structural mechanics, numerical algorithms, parallel computation, mathematical software, and image processing. He has worked for USNAD Crane, Sandia National Laboratories, and General Motors Research Laboratories and has served on the faculties of the University of Michigan and Michigan State University, East Lansing, before coming to Virginia Tech. He received his BA (magna cum laude) in psychology and mathematics from the University of Evansville, Ind., and his PhD in mathematics from the University of Michigan, Ann Arbor. His email address is <ltw@cs.vt.edu>.