

Modeling networks of signaling pathways

Upinder S. Bhalla, National Centre for Biological Sciences,

Bangalore 560065

INDIA

Version of 11 November 1999

Modeling networks of signaling pathways

Introduction.

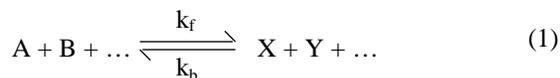
The inside of a cell is a remarkable environment in which to find computation. At first glance it is as if a thick soup somehow manages not only to keep track of many unique signals but also to have them interact in specific ways to give rise to computation. This feat is accomplished through a mapping of signal identity onto molecular identity. In this mapping lies the power of the cellular computer. Every chemical reaction is a molecular transformation, and computation is born through this manipulation of chemical 'symbols'.

In a chemical plant, or for that matter in the conventional 'Harvard' computer architecture, the instructions are neatly separated from the raw materials or data respectively. In the cell, however, computation merges seamlessly with all aspects of cellular chemistry. The cellular machinery of the cytoskeleton and protein trafficking are as much a part of the signaling apparatus as they are substrates for the signaling to act upon. This all-encompassing nature of cellular signaling means that it touches upon most disciplines of biology. It also makes it hard to separate the wood from the trees, that is, understand the overall biological function in terms of the individual enzymatic reactions. A central theme of this chapter is how simulations may help to bridge this gap. As is done in this book, we can identify three main levels of analysis of biological signaling: 'well stirred' (i.e., point or non-diffusive) systems, spatially structured reaction-diffusion systems (Chapter 3), and stochastic models dealing with probabilities of individual molecular events (Chapter 4). One could think of these as beginner, intermediate and advanced topics in the subject. The 'well-stirred' analysis is reasonably well characterized experimentally, and is comparatively straightforward to specify and solve. It is also necessary to first specify the basic reactions before introducing additional details such as space and probability. The emphasis in this chapter is therefore to develop methods for building empirically accurate models of signaling at the level of well-stirred cells. Test-tube biochemistry is relatively simple to simulate, and programs for modeling enzyme kinetics have existed since the days of punch-cards. The difficulties in scaling up to tens of signaling pathways and thousands of reactions are not computational, but to do with interface design and most of all with converting experimental data into kinetic parameters in a model. The general approach outlined in this chapter is to modularize the problem in terms of individual signaling pathways, usually involving just one key signaling enzyme. The grand picture is built up by drawing upon a library of models of individual pathways. This fits well with standard experimental approaches, and also lends itself to plug-and-play upgrades of complex models as more complete individual enzyme models are developed.

Methods.

Equations

Signaling pathways are fundamentally based on chemistry. A description of signaling in terms of a series of chemical reactions is therefore about as general as one can get in the domain of the well-stirred cell. Fortunately, the numerical methods for handling such reaction systems are simple and fast enough to embrace this generality without having to make compromises. Furthermore, the reaction formalism appears to work well even for empirical and approximate models, as we shall see below. As experimental data are frequently incomplete, this is important in adopting this formalism for biochemical models. Our starting point for such models is the basic reaction (Example 1):



This can generalize to any number of terms on either side, and any stoichiometry. Reactions of order greater than two are uncommon, and when they are encountered they are often a reflection of incomplete mechanistic data. At the molecular level, higher order reactions would require a simultaneous collision between multiple molecules. The probability of this is extremely low. Instead, apparent high-order reactions typically proceed through fast intermediate reaction steps.

Two rules govern equation 1:

$$d[A]/dt = -k_f.[A].[B]... + k_b.[X].[Y]... \quad (2)$$

and

$$d[A]/dt = d[B]/dt = -d[X]/dt = -d[Y]/dt ... \quad (3)$$

or equivalently,

$$[A_0] - [A] = [B_0] - [B] = [X] - [X_0] = [Y] - [Y_0] ... \quad (3a)$$

$$[A] + [B] + .. + [X] + [Y] + ... = \text{constant} \quad (3b)$$

Equation 2 is the differential equivalent of the rate equation. The various forms of equation 3 express the stoichiometry of the reaction and are equivalent to conservation of mass.

For small systems of equations it is useful to use the mass conservation relationship directly to reduce the number of differential equations:

$$[B] = [B_0] - ([A_0] - [A]). \quad (4)$$

For large, extensively coupled systems, however, keeping track of the conservation relationships among many diverging reaction pathways is cumbersome and does not save much computation. In this case it may be simpler to write out the full differential equation for each reaction component and solve each independently, secure in the knowledge that the differential equations embody the stoichiometry.

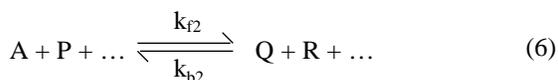
$$d[B]/dt = -k_f.[A][B] ... + k_b.[X].[Y] \quad (5a)$$

$$d[C]/dt = +k_f.[A][B] ... - k_b.[X].[Y] \quad (5b)$$

$$d[D]/dt = +k_f.[A][B] ... - k_b.[X].[Y] \quad (5c)$$

...

An added advantage of this approach is that it lends itself to modular and object-oriented simulation schemes. Suppose we have an additional reaction

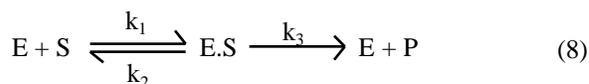


The equation for A alone would need to change to incorporate the additional terms:

$$d[A]/dt = -k_f [A][B]... -k_{f2}[A][P] + k_b.[X].[Y]... + k_{b2}.[Q].[R] \quad (7)$$

If one relies on the differential equations to maintain stoichiometry, it is possible for numerical errors to introduce embarrassing problems with stoichiometry and mass conservation. This can be turned to our advantage, by using conservation as a test for numerical accuracy and as a quick empirical way of deciding if the timesteps need to be shorter.

The prototypical reaction scheme for a Michaelis-Menten enzyme (Michaelis and Menten, 1913, Stryer 1995) crops up frequently in signaling interactions (example 2):



This is readily represented by two reactions in sequence. The final reaction is conventionally assumed to be unidirectional. This is a good approximation for most systems with low product concentrations, though

strictly speaking it too is reversible. The Michaelis-Menten scheme is so common that the modeler may wish to represent it as a self-contained module.

Integration methods.

Three characteristics of biological signaling reactions help select between integration methods:

1. Reaction loops
2. Rates spanning at least three orders of magnitude
3. Very large numbers of reactions

Reaction loops are a very common motif: even enzyme reactions contain such a loop (Figure 1). Other situations include interconversion between states of a system, and when an enzyme can be reversibly bound to multiple activators.

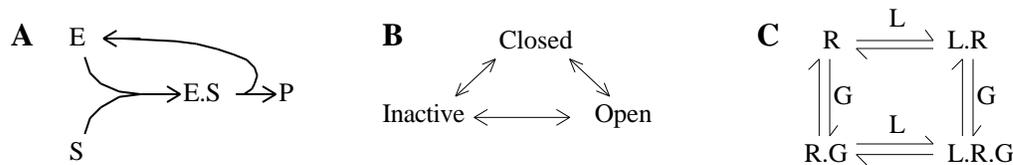


Figure 1: Reaction loops. **A:** Enzyme reaction. E: Enzyme, S: Substrate, P: Product. **B:** States of a voltage-gated ion channel. **C:** Activation scheme for a G-protein coupled receptor. R: Receptor, L: Ligand, G: G-protein.

These loops make it difficult to apply implicit integration methods efficiently to reaction systems. As described above, reaction loops can be eliminated through conservation relationships but this may give rise to other computational problems. For non-looping systems it is sometimes possible to order the implicit solution matrix in a manner that can be solved using backward elimination in order (N) operations. This was described for branching neuronal models by Hines (1984) and is utilized in chapter 8. In the presence of loops, however, one needs to solve the matrix of reactions, which takes order (N³) operations in the general case (Press et. al. 1988). This effectively rules out implicit integration for large reaction systems. This is unfortunate, as the wide range of time-constants (point 2 above) means that one needs to use timesteps of the order of the fastest rate in the system unless one is using implicit methods (Mascagni and Sherman 1998). Higher-order integration methods, especially with variable timestep, may be useful in solving reaction systems under many conditions. Runge-Kutta and Bulirsch-Stoer are among the techniques that have been examined (Press et al 1988). However, both methods are non-local, that is, they are difficult to decompose in a modular manner. They may also suffer from instability or drastic speed reduction when confronted with abrupt stimuli typical of experimental manipulations. The wide range of reaction rates also slows down such methods. The exponential Euler method (MacGregor 1987) is a very simple integration scheme predicated on the assumption that variables will follow exponential decay time-courses. This happens to fit well with most reaction systems, and the simplicity of the method makes it easy to implement in a modular fashion. It is also extremely robust and can handle abrupt stimuli gracefully. It is, of course, not as efficient as the higher-order methods operating at their best, and does not have built-in variable time-step calculations. These can sometimes be imposed 'from above' if the modeler knows when a particular simulation is expected to be in a slowly changing state.

Software

Given that chemical rate equations entail the simplest kind of differential equations, it is not surprising that there is a plethora of useful software available for carrying out the numerical modeling (e.g., Matlab, Mathematica, others). Assuming that all packages do a reasonable job with the numerical integration, the choice of method should be governed more by availability and familiarity than by small differences in efficiency. For most models and modern computers, simulation times are likely to be much faster than real time, and one has the luxury of using whatever method is most convenient. Therefore the real challenge in specifying complex networks of signaling pathways is user-interface and data management related. There are rather fewer packages designed specifically for handling large numbers of chemical reactions in a friendlier format than simply entering rates and differential equations. (Xpp, Vcell, GENESIS/kinetikit). Given the interest in the field, such packages are likely to evolve rapidly.

Parameterization

Parameterization is by far the most difficult aspect of modeling signaling pathways (and most other biological systems as well). In developing a model one has to draw a balance between known biological detail, and the applicability of such detail to the problem at hand. This is constrained further by the availability of data, and the sheer scale of the problem of modeling extremely complex systems.

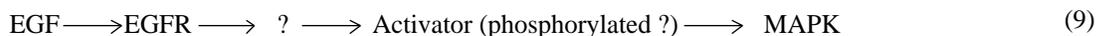
From test-tube data to models.

Cellular signaling is complex both in terms of the number of players, and the properties that emerge from their interactions. A systematic reductionist approach to cataloguing and characterizing the major pathways may seem daunting. Nevertheless, the admirably quantitative traditions of biochemistry have led to the accumulation of a remarkably complete base of data for many signaling molecules. Within the constraints of test-tube chemistry, it is possible to derive faithful models of many individual pathways. For example, the major protein kinases PKC (Nishizuka 1992), PKA (Døskeland and Øgreid 1984), CaMKII (Hanson and Schulman 1992) and MAPK (Huang And Ferrell Jr., 1996) have each been well characterized. Such data are essential for developing internally consistent models of pathways. It is significant, however, that the data are complete enough that *independently* developed models of the pathways are also consistent (Huang and Ferrell Jr. 1996; Bhalla and Iyengar 1999). Further validation of these models can be done by examining their predictive capacity in composite models. For example, the MAPK response to EGFR stimulation involves the receptor and its immediate signaling cascade, the Ras pathway as well as the MAPK cascade (See equation 10 below). Replication of such complex signaling sequences suggests that one can reasonably scale up the reductionist approach to larger problems.

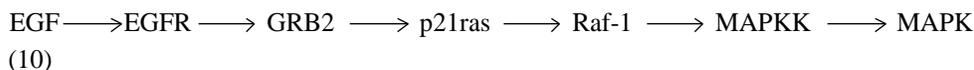
In order to scale up from a collection of rate constants to a framework which can tackle interesting biological problems, a necessary first step is the development of a library of models of individual pathways (Bhalla and Iyengar 1999, URLs: <http://www.ncbs.res.in/~bhalla/itploop/> or <http://piris.pharm.mssm.edu/urilab/>). The next stage is to merge relevant pathway models into a composite simulation of the system of interest, inserting system-specific interactions and parameters where necessary. This is a multistep process. Where experimental results involving combinations of a few of models are available, these can provide excellent constraints on model parameters. Finally (and this is the fun part) the composite model is explored with a variety of stimuli, emergent properties are examined, and explanations for known phenomena are sought.

Data sources.

Primary data for model building come in three main flavors. The first category of paper tends to be very qualitative, and is frequently of the form of "Gene A expressed in cell type B activates pathway C." Nevertheless, such data are critical for establishing the mechanisms of the pathway, and the causal sequence of signaling events. It is instructive to follow the early literature on the MAP Kinase cascade (Reviewed in Cobb et. al. 1991). Initial studies indicated co-activation of various stages of the cascade. Intermediates were suspected, but not identified.



It is a measure of the rapid pace of the field that by 1992, some clear experiments were done which established the current version of our understanding of the cascade (reviewed by Ahn et. al 1992).



In the laboratory as well as in model development, such studies are usually followed by quantitative analyses which put some numbers onto the mechanistic arrows. A typical paper in this category would seek to purify a newly identified signaling molecule, verify its mechanism of activation in the test tube, and obtain rate constants. For enzymes this purification usually involves a sequence of chromatographic separations leading to ever higher specific activities and lower yields. Table 1 illustrates such a purification sequence for MAPK (Adapted from Sanghera et al 1990).

Purification stage	Volume (ml)	Total Protein (mg)	Total activity nmol.min ⁻¹	Specific activity nmol.min ⁻¹ .mg ⁻¹	% Recovery	Purification
--------------------	-------------	--------------------	---------------------------------------	--	------------	--------------

Cytosol	366	11488	1241	0.108	100	1
DEAE-cellulose	432	5490	898	0.16	72	1.5
Phenyl-Sepharose	70	28	378	13.5	30	125
Mono-Q	10	0.05	46.9	938	3.8	8685

Table 1: Purification series for MAPK.

Such papers provide not only activity parameters, but also an estimate of concentration of the signaling molecule in the preparation. These numbers are, of course, interdependent. The good news is that the degree of purification affects the activity and concentration in inverse proportion, so one does have a good estimate of total cellular enzyme activity even if concentration values are inaccurate. The bad news is that most regulatory interactions require a good estimate of concentration. Continuing with our MAPK example, it turns out that the concentrations for MAPK in the hippocampus are some ten times that in most other neural tissues (Ortiz et. al 1995). Such system-specific parameters greatly complicate quantitative modeling. As discussed below, there are still more fundamental difficulties in extracting system-specific parameters using most current experimental approaches.

The third category of paper, in which detailed reaction mechanisms and kinetics are analyzed, is much to be treasured. Such papers typically contain numerous concentration-effect curves for enzyme activation and binding of molecules, often with estimates of time-course (Døskeland and Øgreid 1984, Posner et. al. 1992). These curves tightly constrain parameters and reaction mechanisms.

Iterative parameterization

An overview of the process of parameterization is provided in the flowchart in Figure 2. The key aspect of this flowchart is that it is iterative. We start with a simple, plausible mechanism for a signaling molecule and results for a specific interaction involving it. Once this is parameterized to satisfaction, another interaction may be considered. This elaboration of the reaction mechanism will usually mess up the first stage of parameterization, which must be repeated, until now both interactions behave well. In systematic experimental papers, interactions are often examined in a progressive sequence which fits very well with such an iterative approach. The process of parameterization is illustrated in a worked example for developing a model of PKC (Box 1, example 4 on the CD). There is a great deal of similarity between this process and the Towers of Hanoi: each additional interaction may mean that the modeler must go through the whole process again. It is well to apply Occam's razor ruthlessly to prune elaborate mechanisms: do not assume any reactions unless the data demand it. The reader is invited to work out the number of iterations required to match, say, five graphs of experimental results, assuming that each new graph introduces reaction changes that require adjustment to all previously fixed parameters. This iterative process of model development typically goes through 50 or more versions for each new pathway modeled. A similar process may be needed to refine models in the light of new data or to incorporate additional mechanistic details such as cytoskeletal interactions (Alberts et al 1983).

Parameter searching

Several of the stages in the flowchart require parameter estimation based on fits between simulated and experimental data. This is an aspect of modeling that is especially suited to judicious automation. Parameter searching can be done in three main ways: user-guided, brute force, and automated. The user-guided approach takes advantage of human insight into the system to select the best of a small number of versions of the model. Typically rather complex judgments, for example, about the shape of a stimulus response curve, enter into such searches. It is obviously a labor intensive and somewhat subjective method. Nevertheless it is often the best approach especially in the initial stages of refining a simulation. The brute force and automated searching methods both require a measure of the 'goodness' of a given set of parameters. This can be difficult to define. Obvious methods include mean-squared error of the simulation results as compared to experimental data. The key requirement for such a measure of 'goodness' of fit is that it should improve monotonically and smoothly as the simulation improves. Often the mean-squared term behaves well near the best fit, but does not help much if the parameters are significantly off. More sophisticated measures may behave better but are correspondingly harder to set up. Brute force methods are effective if there are only a few parameters to examine, and if each simulation runs rather quickly. In this approach each parameter is scaled up and down in small increments which are expected to span the range of

interest. If there are N parameters each run for M scales, then the number of required simulations = M^N . A practical upper limit for N is typically four, using an M of 3 or 4. A variety of automated methods exist, such as simplex, gradient descent, and conjugate gradient (Vanier and Bower in press). 'Genetic algorithms' have also recently been used with some success. It is difficult to provide a general prescription here; each specific problem has its own unique features. The only common rule is not to let the searches proceed blindly. The modeler usually has a pretty good feel for 'reasonable' parameters, and can catch the search algorithm when it threatens to launch off into the realm of the absurd.

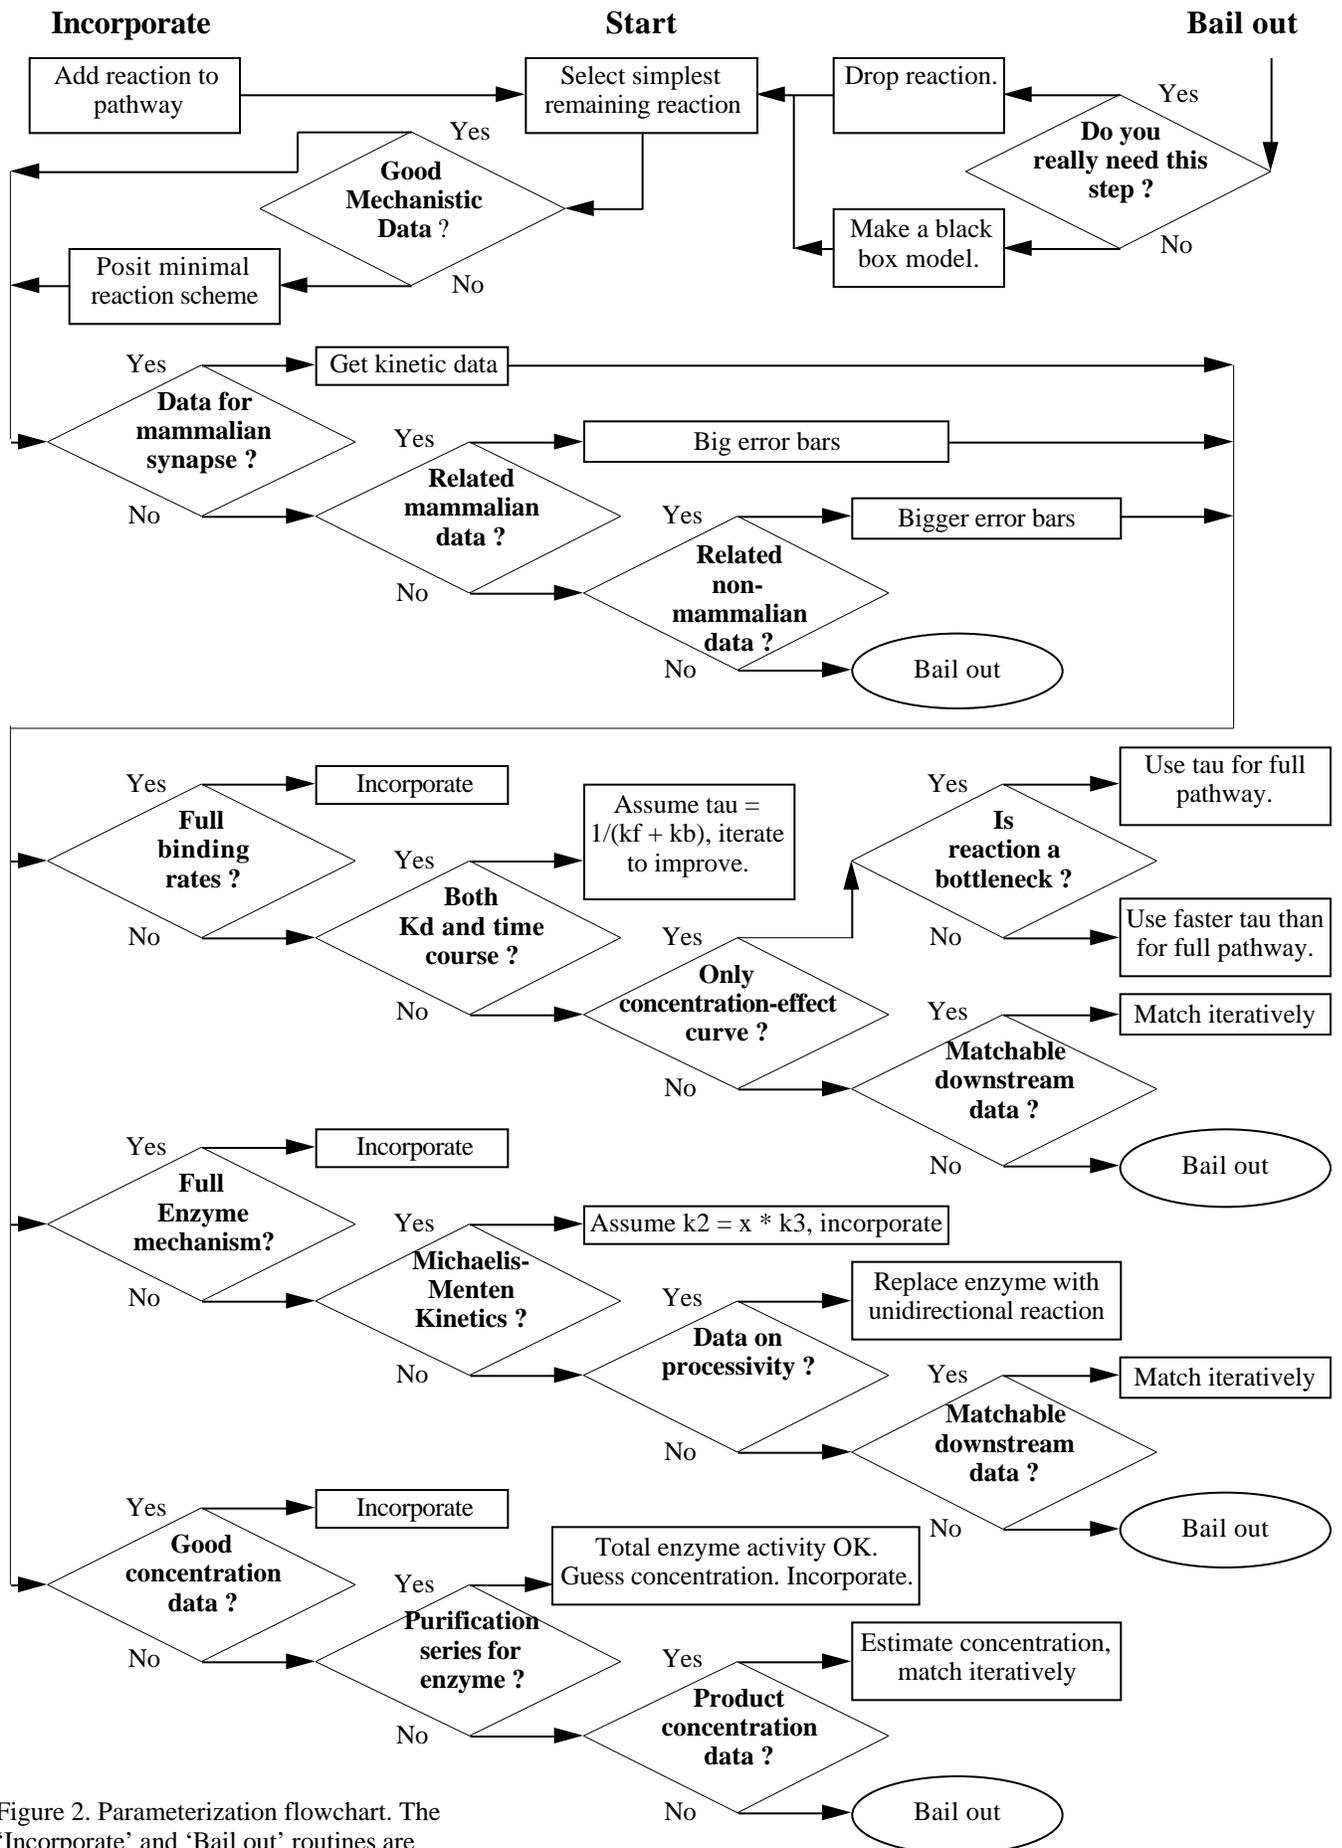


Figure 2. Parameterization flowchart. The 'Incorporate' and 'Bail out' routines are indicated on top.

Model integration.

A well designed 'library' of models of pathways would ideally be plug-and play: one should be able to take individual pathway modules and connect them together as if one were drawing a block diagram. For this to work, the interactions between pathways should be incorporated into the pathway models themselves, rather than pasted on after the composite simulation has been assembled. This principle makes the parameterization process relatively independent of subsequent use of the pathways. It is also cleaner in the sense that emergent properties of the system are not susceptible to being 'built in' by tweaking all parameters together. Three common motifs for inter-pathway interactions are second messengers, multimer formation, and covalent modification.

Interactions between pathways via second messengers are trivial to implement: cAMP produced by any of the 15-odd adenylyl cyclases (or applied through a pipette) will activate PKA identically.

Multimer formation (such as binding of Ca_v CaM to CaMKII) or dissociation of multimers (as in G-proteins) is somewhat more involved. One complication is that all components of the multimer may undergo modification of activity in a substrate-specific manner. For example, the catalytic activity of the G-protein alpha subunit is different in isolation, and when bound to target enzymes (Biddlecome et al. 1996). Furthermore, the GTPase activity depends on which target enzyme it is bound to. Another difficulty is that unlike second messengers, every participating protein in a multimer may be present in multiple isoforms. Such detailed data may be difficult to come by, and incorporation of all isoform permutations into a library is tedious at best.

Communication between pathways through covalent modification, such as by phosphorylation, is still more problematic. First, enzyme-substrate interactions are often highly specific. The isoform permutation problem mentioned above is, if anything, even worse. In addition to specificity of rates for enzyme action, there is also isoform specificity of signaling responses following covalent modification. Determination of the signaling responses due to a specific enzyme-substrate pair often boils down to educated guess work: Can one assume that the phosphorylation rates for PKC- α on ACII are the same as for PKC- α on ACIII? Second, covalent interactions are often modulatory rather than switch-like. Worse, the modulatory effects are typically applicable for a specific activity state of the enzyme. An example of this situation might be the experimental result that MAPK phosphorylates PLA_2 to increase its activity some 3-fold (Lin et al. 1993). Closer examination of the methods reveals that the activity was measured in 5 mM Calcium. It would now be useful to know what happens over a more physiological range of Ca^{2+} . Yet more detailed models would need to incorporate increase in activity as compared to other activators of PLA_2 as well. Practically speaking, pathway interactions may involve even more parameters than the pathways themselves.

Model integration, that is, converting a set of models of individual pathways into a network, is still a long way from the ideal plug-and-play situation. Current models may have to fall back on the following compromises between the modular 'library' approach and completely uninhibited parameter tweaking.

- Depending on data sources, it may be simpler to develop two or more tightly coupled pathways as a single unit. The MAPK pathway is in fact a sequence of 3 kinases, with interspersed inhibitory phosphatases. These reactions are so tightly coupled that in experiments as well as in modeling, it is simplest to treat them as a single entity.
- Data may be more readily available for a series of pathways than for the individual elements. The EGFR pathway and the MAPK cascade in Equation 10 illustrate a situation where the direct outputs of the EGFR are difficult to measure, but a good deal of data are available for MAPK activation by EGF (e.g., Wahl et al. 1992). Given the input (EGF) and the output (MAPK activity) and known parts of the pathway (Ras, MAPK), one can 'solve for' the unknown (EGF pathway). In other words, the MAPK pathway just acts as an output black box in our parameter selection process for the EGF pathway.
- In several cases (e.g., MAPK \rightarrow PLA_2), the individual models are well defined and the composite effect is known but the details of the interaction are not. Here the interaction strength is the unknown which can be parameterized while the structure of the individual models are kept fixed.

The final step of model integration, when several individual pathway models are merged into a grand simulation, is precisely the point at which models scale beyond the grasp of most unaided human minds. Well-designed interfaces play a key role at this stage. First, they can provide sanity checks for trivial problems like duplication of enzymes in different pathway models, or incompatible units. Second, they allow quick checks for convergence of molecule concentrations to reasonable steady-state values. Most

importantly, the interface should make it easy for the modeler to match simulations against experiments which provide data for an entire network of pathways. This is, after all, the purpose of the entire exercise.

Model evaluation

Model evaluation involves at least two questions: How good are the data; and, how good is the model? The 'garbage in, garbage out' rule means that the first sets fundamental limits to accuracy. This is reflected in the flowchart above. Error bars, while they inspire confidence in a given experiment, may not be the best way of evaluating data accuracy. It is common for different groups to have widely divergent results for the same enzyme, each internally accurate to around 10% (Table 2).

<i>Source</i>	<i>Km</i>	<i>Vmax (min⁻¹)</i>
Paudel et al, JBC 268 (9): 6207-6213 1993	16.6+-4.7	643+-85
Huang et al ABB 305(2):570-580 1993 (Using PKC)	26.4+-1.3	285+-11

Table 2: Phosphorylation of neurogranin by PKC. Both sets of experiments were run at 30 degrees C, however Huang et al separately measured rates for the α , β and γ isoforms of PKC (α isoform reported here). The methods used in each paper appear tightly controlled and reliable.

Subtle differences in the methods may lead to several-fold differences in rates. One useful rule of thumb is that enzyme rates in-vivo are likely to be higher than test-tube rates: all else being equal, choose the highest reported rate. Multiple data sources for the same parameter are extremely valuable as they provide a feel for the range of values and often a range of conditions and supplementary results. The preferred technique for model evaluation is to test its predictive capacity. Within a paper, some plots may contain mutually redundant data, though it would be all but impossible to derive one from the other without explicitly making a model. As part of the iterative process described earlier, one might imagine approaching each new graph in a paper hoping that the existing interactions and model are already capable of predicting it. (Box 1-Figure 2). The disparity between results from different papers usually makes it difficult to accurately predict a figure in one paper completely from another. One should, however, be able to 'tune' a model to a specific paper with relatively little effort if the model is fundamentally sound. It is worth making a distinction here between the fundamental mechanisms and rates of a model (which should be the same for a given enzyme) and the actual concentrations in a given experiment. In test-tube experiments in particular, concentrations are entirely up to the experimenter. This lays a particular onus on the modeler to accurately model experimental conditions. This seemingly trite recommendation is non-trivial. For example, consider a two-state enzyme whose activity is measured using a test substrate and a concentration-effect curve for levels of product. An obvious shortcut in simulating such an experiment would be to simply measure (in the simulation) the levels of the active form of the enzyme. There are at least two pitfalls here. First, experiments to measure enzyme activity often utilize saturating substrate levels. If the enzyme acts on some other regulatory component the substrate will competitively inhibit this other enzyme activity (Figure 3A). Unless the substrate assay is modeled explicitly, this competitive interaction will be missed. Second, the mechanism for regulation of the enzyme usually assumes that the enzyme-substrate complex is a transient, and is present at very low levels. This assumption breaks down at saturating substrate levels. Here we would have to explicitly include the enzyme-substrate complex as a participant in other interactions. (Figure 3 B). These models are illustrated on the demo CD as example 3.

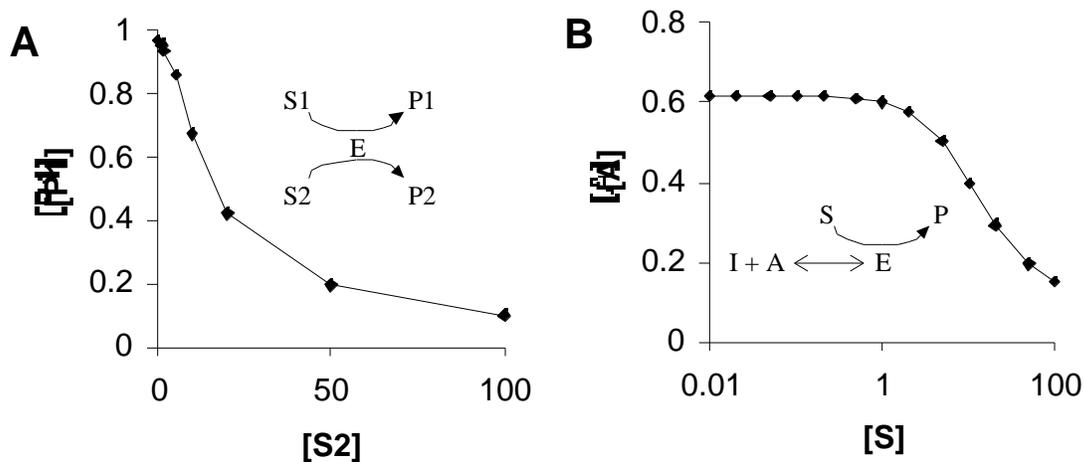


Figure 3. Effects of high substrate concentrations. A. Competitive inhibition of formation of P1 by high levels of S2. The reaction runs for 100 sec, with 1 μ of S1 and 1 μ M of enzyme. B. Shifted activation curve depending on substrate levels. The inactive enzyme I and activator A start out at 1 μ M. When substrate levels approach saturation the amount of free enzyme declines, leading to the shift in the curve. Any other reaction depending on the activator would be affected. In both cases all enzyme rates are set to 0.1, so $V_{max} = 0.1$ and $K_m = 2$.

Steady state data often account for most of the parameters in a simulation of signaling pathways. Such experiments can often be replicated by more than one possible reaction mechanism. Models based on such data are clearly somewhat 'loose'. Experimental designs which explicitly incorporate time-series measurements provide much tighter constraints on model parameters (e.g., De Koninck and Schulman 1998).

Complications and future directions.

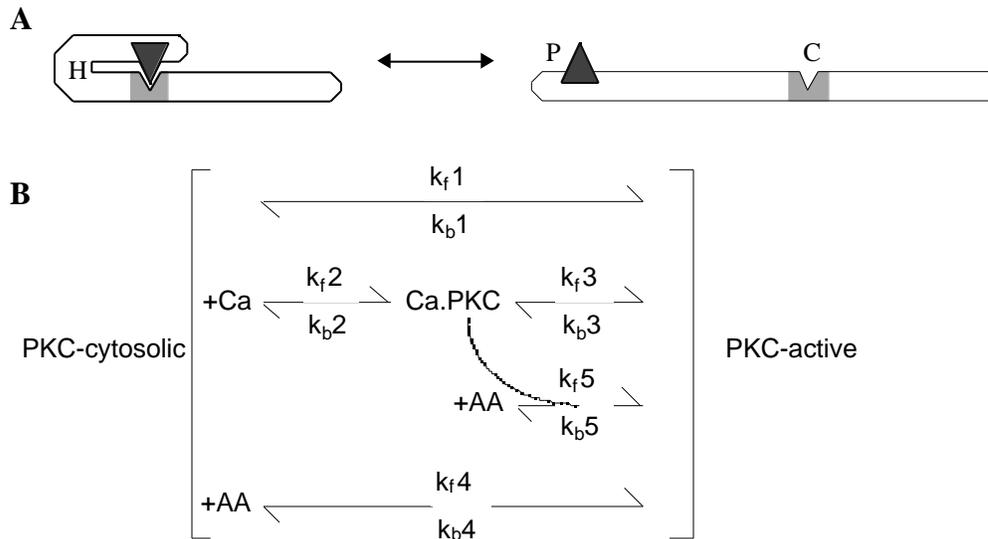
A cynic would identify two serious problems with current attempts to model signaling pathways. The first is that there is not enough detail in the models. The second is that there is too much. The first criticism refers to the fact that biological systems are far more intricate and computationally rich than crude test-tube chemistry allows. As described much more comprehensively in subsequent chapters, our simple differential equations are only a first approximation to the three-dimensional, compartmentalized, stochastic events in a real cell. Even at the very limited level of 'well-stirred' biochemistry, current models barely begin to describe the known menagerie of signaling molecules. The second problem relates to the limitations of experimental data. It is perhaps presumptuous to attempt to devise models involving dozens of signaling pathways given the current state of our knowledge. So far no more than two or three signaling molecules have ever been simultaneously monitored in real time in vivo (De Bernardi and Brooker 1996). There is also a growing recognition that macroscopic, test-tube experiments are seriously flawed as models of cellular signaling. Cellular chemistry is not a point process, or even simply diffusive. Complex sequences of reactions are known to occur in an assembly-line fashion on cytoskeletal scaffolds (Garrington and Johnson 1999). Such reactions are not necessarily well described by bulk concentrations, rates, or even mechanisms.

From a practical viewpoint, both problems are addressed by the simple observation that many models do indeed replicate and predict experimental data. These models are therefore at least empirically useful. To the extent that such models remain faithful to biology, they promise to provide the theoretical counterpart to the biology of signaling. It is increasingly evident that such a theoretical understanding will be critical to overcoming two of the major difficulties in the field: complexity and patchy experimental data. In concert with web-based databases of signaling pathways (URLs below), the methods described in this book form an emerging framework for a quantitative approach to studying cellular signaling.

Box 1. PKC parameterization.

PKC is a 'mature' enzyme: it has been studied in sufficient depth that most of the relevant data are available. Here we use it as an example of the parameterization process. We start with the following mechanistic data (Nishizuka 1992):

- PKC is activated by Ca^{2+} , AA and DAG. We will examine only the first two for this exercise.
- PKC, like many serine/threonine kinases, is kept in the inactive form by an inactivation domain on the kinase itself which contains a pseudo- substrate. (Box 1-Figure1 A) Normally the protein folds over onto itself so that the pseudo-substrate domain covers the catalytic site and prevents access to other substrates. Activators unfold the protein and expose the catalytic site. This suggests that we could model it as a two-state enzyme, where the activity is always the same when open, and different activators open different fractions of the total enzyme (or open it for different times).
- Membrane translocation is known to be involved in PKC activation
- The time-course of activation of PKC is under 10 sec. Ca^{2+} activation is under 1 sec (Nishizuka 1992, Schaechter and Benowitz 1993).



Box 1 Figure 1. Activation of PKC. A: Activation of PKC. In the inactive form it is folded over at the hinge region H so that the catalytic site C is covered by the pseudosubstrate P. When activated, the catalytic site is exposed. B: Reaction scheme for PKC activation model. The cytosolic form can undergo any of reactions 1, 2 or 4. The active form of PKC is formed by each of the reactions 1, 3, 4 or 5.

In order to model the kinase, we go sequentially through the data plots in Box 1-Figure 2. The solid lines represent experimental data (Schaechter and Benowitz 1993).

1. Basal activation (Box 1 Figure 2 A). Inspecting the plot, we first note that there is a basal level of activation even at very low Ca^{2+} . This amounts to about 5% of maximal activity. Ignoring the 1 nM Ca^{2+} , we can set up this activity by postulating a basal reaction where 5% of the cytosolic enzyme goes to the active state. So,

$$k_{f1}/k_{b1} \approx 1/20 \quad (1)$$

Setting this to a 1 sec time course,

$$1 \text{ sec} \approx 1/(k_{b1} + k_{f1}) \quad (2)$$

Solving:

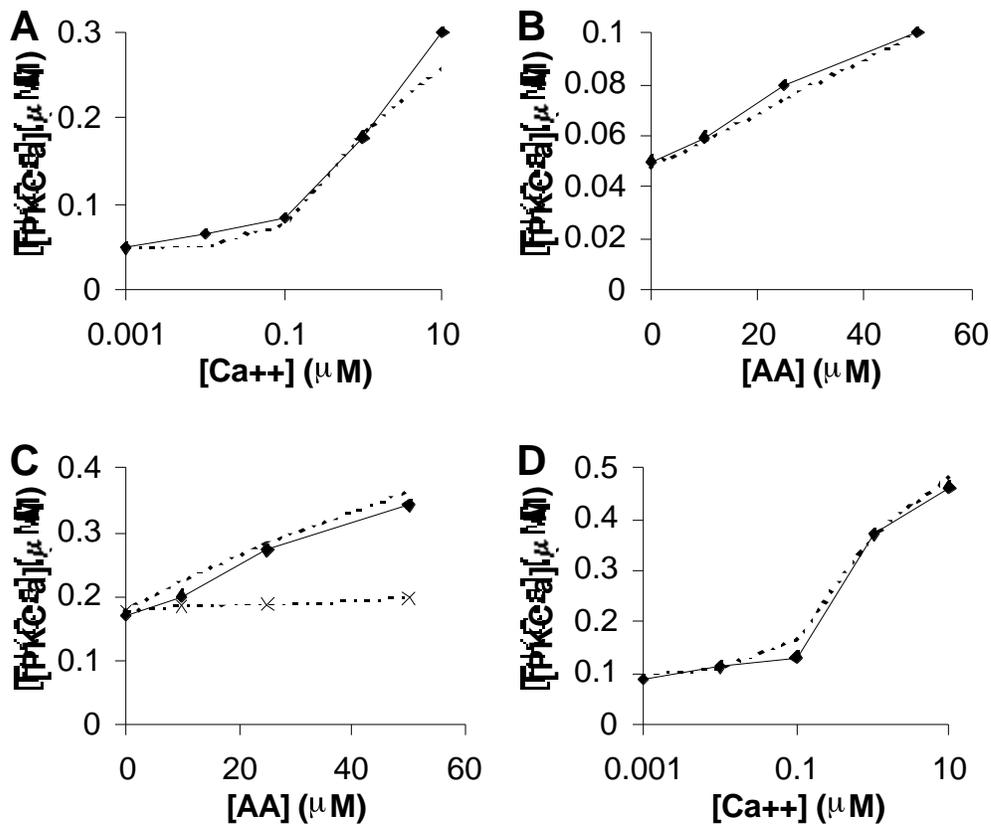
$$k_{b1} \sim 1, k_{f1} \sim 0.05 \quad (3)$$

2. Ca^{2+} activation. The curve in Box 1-Figure 2 A appears to have a half-max of about $1 \mu\text{M Ca}^{2+}$. Unfortunately this produces nothing at all like the desired curve. The maximal activity with Ca^{2+} stimulation alone is only about 30% of peak PKC, whereas a simple Ca-binding reaction will give us 100% activity. This is a situation where the known mechanistic information about membrane translocation gives us a useful hint. Let us keep the half-maximal binding where it was, at about $1 \mu\text{M Ca}^{2+}$ (Reaction 2), and assume that the membrane translocation step (Reaction 3) is what allows only 1/3 of the kinase to reach the membrane. This additional degree of freedom lets us match the curves nicely. The final parameters are pretty close to our initial guesses:

$$kf2 = 0.6, kb2 = 0.5 \quad (4)$$

$$kf3 = 1.27, kb3 = 3.5 \quad (5)$$

It is important to keep in mind that this is still just an empirical model for the true reaction mechanisms. Although we can justify the additional reaction step by invoking membrane translocation, strictly speaking we do not have any direct experimental basis for this. Indeed, a blind parameter search will quite effectively utilize the two additional parameters to give us the desired curve, without any insight into mechanistic details.



Box 1-Figure 2. PKC Regulation: experimental (solid lines) and simulated (dashed lines). A: Basal and Ca-stimulated activity. B: AA-stimulated activity without Ca. C: AA-stimulated activity in the presence of $1 \mu\text{M Ca}$. Crosses indicate curve obtained without using a separate reaction to represent synergy. D: Ca-stimulated activity in the presence of $50 \mu\text{M AA}$. The simulated curve was predicted from the model based on panels A-C.

3. Matching the AA activation in the absence of Ca^{2+} . Examination of the curves (Box 1-Figure 2 B) indicates that the activation is almost linear in this concentration range. Therefore, almost any large Kd should do as a starting point. After a little exploration, values of

$$kf4 = 1.2e-4, kb4 = 0.1 \quad (6)$$

seem reasonable. We use a slower 10-sec time-course for the AA binding steps.

4. Matching AA activation with $1 \mu \text{Ca}^{2+}$. A quick run with the previous version of the model shows that it is not able to account for the synergistic activation of PKC by AA and Ca^{2+} combined (Crosses in Box 1-Figure 2 C). We therefore must introduce this synergy in the form of combined binding of Ca^{2+} and AA to the kinase (Reaction 5). We already have Ca^{2+} - and AA-bound forms of the kinase. Which should we use as the starting point? Given that Ca^{2+} binding is faster, it is reasonable to use the Ca^{2+} .PKC form as our first intermediate, and subsequently bind AA to this to give our synergistically active form. Furthermore, this combined form should account for the bulk of the kinase when both Ca^{2+} and AA are present, so its reaction should have a much tighter affinity for AA than does the PKC-AA reaction. Working through the matches, we find that a 10-fold increase in affinity produces a good match:

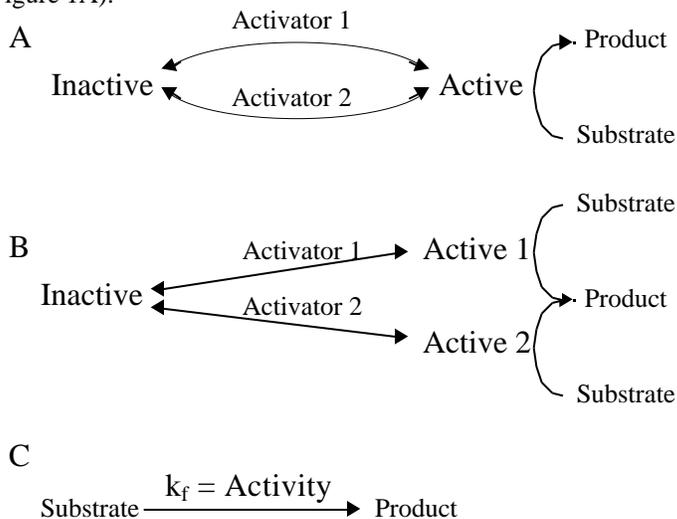
$$kf5 = 1.2e-3, kb5 = 0.1 \quad (7)$$

5. Matching Ca^{2+} activation with fixed $50 \mu \text{AA}$. This is the test of the completeness of our model. In principle we have already accounted for all combinations of Ca^{2+} and AA activation. As it turns out, our model does indeed do a good job of replicating this curve, without adding any further reactions (Box 1-Figure 2 D). Not only is our model empirically able to match the previous figures, it has been able to generalize to a new situation. This improves our confidence in the model being applicable to other signaling situations as well. Our 'final' model is shown in Box 1-Figure 1 B. The model is also set up as example 4 on the CD. As an exercise, the reader is invited to try to match further interactions involving DAG.

Box 2. Recurring biochemical motifs in signaling.

There are a few specific motifs that recur in models of biological signaling, and are worth examining more closely.

1. The two-state system. Several enzymes, especially the protein kinases, can be reasonably well modeled as two-state systems: either completely on or completely off. The level of activity of the system is determined simply by the proportion of the enzyme in the 'on' state. For example, protein kinases typically have an autoinhibitory domain, which binds to and blocks the active site. Activators of the enzyme, to first approximation, simply release this block (Bhalla and Iyengar 1999). As modeled above for PKC, one plausible mechanism for such an enzyme is for the activators to bind to the inactive enzyme, and convert it to the active state. Similar 'switches' seem to take place in many enzymes by phosphorylation. (Box 2- Figure 1A).



Box 2-Figure 1: Models of enzymatic regulation. A: Two-state enzyme where either of two activators form the same active state. B: Empirical enzyme. Each activator produces a distinct activity level. C: Unregulated enzyme. A simple unidirectional reaction suffices.

2. The 'empirical' enzyme. A remarkably good empirical description for an enzyme about which little is known is simply that each activator 'turns on' a distinct level of activity (Box 2-Figure 1B). This mechanism is relatively easy to relate to experimental data, but it does lead to a proliferation of activity states and enzyme sites with different rates. PLA2 is an enzyme which has been modeled in this manner (Bhalla and Iyengar 1999).

3. The unregulated enzyme. If a pathway has an enzyme whose activity is unlikely to be saturated, and which does not undergo any known regulation, it is safe to treat it as a simple unidirectional reaction with a fixed rate constant (Box 2-Figure 1C). This situation frequently arises in degradative pathways of second messengers (e.g., DAG and AA).

4. The reaction loop. This is a common trap for the unwary. Consider the quite plausible scheme for G-protein coupled receptor activation as illustrated in Figure 1 C. There is a strong constraint on the rates: the product of the Kds around the loop must always be one. The proof of this is left as an exercise for the reader. (Hint: Consider free energy.) There is no problem with reaction loops where there is an energy-releasing step, such as phosphorylation or ATP hydrolysis.

5. Buffers. Almost all test-tube experiments rely on buffers. For example, the standard way to obtain a concentration-effect curve, either experimentally or in a simulation, is to buffer the control molecule at a series of concentrations and monitor the resulting effect on the target. A good chemical buffer is actually quite unpleasant to model. By definition, a good buffer must a) be able to compensate for the largest chemical concentration changes that may happen in the system, and b) do so faster than other reactions. Ca^{2+} buffering systems, for example, have large amounts of avid chelators like EGTA with balancing

amounts of Ca^{2+} . Cooperative binding makes such buffers even more effective. Accurate models of such buffered systems therefore require very small timesteps. If one is willing to settle for a 'perfect' rather than a 'realistic' buffer, it is almost embarrassingly easy to do so numerically: Simply fix the concentration of the molecule at the target level. One drawback of such perfection is that any one-way reaction fed by such buffers will quite happily build up infinite levels of product.

Box 3. Useful equations and conversions.

1. Michaelis-Menten equation and relationship to rate constants.

Standard formulation:



Where S is substrate, E is enzyme, E.S is enzyme-substrate complex, and P is product.

$$V_{\max} = \text{maximum velocity of enzyme} = k_3. \quad (2)$$

Derivation: Substrate is saturating, so all of E is in E.S form. So

$$V_{\max} \cdot [E_{\text{tot}}] = [SE] \cdot k_3 = [E_{\text{tot}}] \cdot k_3 \quad (3)$$

The units of V_{\max} come in a variety of forms, due care is necessary in conversion as described below.

$$K_m = (k_3 + k_2)/k_1 \text{ (by definition)} . \quad (4)$$

In most kinetic experiments, only V_{\max} and K_m are obtained for a given enzyme. This means that we are short one parameter for constraining k_1 , k_2 and k_3 . An assumption which seems to work well is

$$k_2 = k_3 * 4 \quad (5)$$

This is not entirely arbitrary. In many enzymes $k_2 \gg k_3$. If k_2 is small compared to k_3 then a large proportion of the enzyme will be in the complex form, assuming K_m is fixed. In such cases it may be better to explicitly model the enzyme reaction as a series of conventional reaction steps, rather than hide the enzyme complex. The factor of 4 keeps the amount of enzyme complex fairly small while avoiding extreme rate constants which might cause numerical difficulties. In previous studies, this factor has been varied over a range from 0.4 to 40, with very little effect on the simulation results (Bhalla and Iyengar 1999).

2. Standard bimolecular reaction



Where A and B are reactants and C is the product. Note that this is completely reversible. The equilibrium dissociation constant is:

$$K_d = k_b/k_f \quad (\text{by definition}) \quad (7)$$

So if B is limiting, and half of B is bound, then at equilibrium:

$$[A][B_{\text{half}}] \cdot k_f = [C_{\text{half}}] \cdot k_b = [B_{\text{half}}] \cdot k_b \quad (8)$$

Since the other half of B has been converted to C. So we get the standard relationship for K_d :

$$[A_{B_{\text{half}}}] = k_b/k_f = K_d \quad (9)$$

or, K_d is that conc of A at which half of B is in the bound form. And, obviously, the association or binding constant is:

$$K_a = k_f/k_b = 1/K_d \quad (10)$$

3. Biological ratios (for mammalian cells)

Proportion of protein to cell weight: 18% (Alberts et. al. 1983 pp. 92) (11)

Proportion of lipid to cell wt: 3% (Alberts et. al. 1983 pp. 92). (12)

Approx. 50% of membrane wt is protein
(up to 75% in some organelles) (Alberts et al. 1983 pp 264). (13)

So maybe 3% of cell mass is membrane protein. (14)

4. Concentration and other units

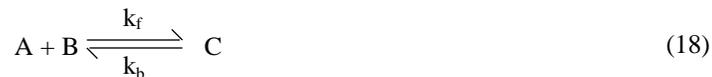
When working within a single compartment, it is simplest to use standard concentration units such as μM . However, transfers between compartments or to and from the membrane are best handled using the number of molecules. Indeed, the 'concentration' of a molecule in the membrane is a somewhat tricky quantity. To illustrate, we will consider conversion of standard enzymatic units to number units in a cubical model cell of 10μ sides.

Vol of cell = $1\text{e-}6 \mu\text{l}$ (15)

$1 \mu\text{M}$ in cell will have $1\text{e-}18$ moles $\sim 6\text{e}5$ molecules/cell (16)

V_{max} of $1 \mu\text{mol}/\text{min}/\text{mg}$ will convert to $\text{Mwt}/6\text{e}4 \text{ \#}/\text{sec}/\text{\#}$ (17)

where Mwt is the molecular weight of the enzyme and numbers/cell are represented by the # symbol. The units can be interpreted as "# of molecules of product formed per second per # of enzyme molecules". Similarly, rates for binding reactions will need to be scaled to the appropriate units. For reaction 6 above:



Take $k_b \sim 1/\tau \text{ sec}^{-1}$ (19)

Then

$k_f = k_b/K_d \sim 1/(\tau * K_d * 6\text{e}5) \text{ sec}^{-1}\text{\#}^{-1}$ (20)

So the units of K_d would be #.

These equations give us the rates k_f and k_b in terms of K_d and τ , and would be a good place to start refining a model to fit an experimental concentration-effect curve.

GLOSSARY OF ABBREVIATIONS

PKC: Protein Kinase C

PKA: Protein Kinase A

CaMKII: Calcium Calmodulin Activated Protein Kinase Type II

MAPK: Mitogen-Activated Protein Kinase

AA: Arachidonic Acid

DAG: Diacylglycerol

cAMP: cyclic Adenosine MonoPhosphate

CaM: Calmodulin

GTP: Guanosine TriPhosphate

ACII, ACIII: Adenylyl cyclases type II, III

PLA₂: Phospholipase A₂

EGF: Epidermal Growth Factor

EGFR: Epidermal Growth Factor Receptor

FURTHER READING:

Modeling signaling pathways:

Lauffenberger, D.A., and Linderman, J.J. (1993) *Receptors: Models for Binding, Trafficking and Signaling*, Oxford University Press, New York.

Thomas, R., and Thieffry, D. (1994) *Developing a logical tool to analyse biological regulatory networks*, in: *Computing with Biological Metaphors*, ed. Paton, R. Chapman and Hall, London.

Bhalla, U.S. (1998) *The network within: Signaling Pathways*, in: *The Book of GENESIS*, 2nd edition, ed. Bower, J.M. and Beeman, D. Springer Verlag, New York.

Computation by signaling pathways:

Bray, D. (1995). Protein Molecules as Computational Elements in Living Cells. *Nature* **376**: 307-312.

Enzyme basics:

Chapter 8, Stryer, L. (1995) *Biochemistry 4th edition*, W. H. Freeman and Company, New York.

Cell biology basics:

Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J.D. (1983). *Molecular biology of the cell*. Garland Publishing, Inc. New York.

URLs:

Kinetic simulation software:

GENESIS/kinetikit: <http://www.bbb.caltech.edu/GENESIS>

XPP: <http://www.pitt.edu/~phase/>

Vcell: <http://www.nrcam.uchc.edu>

Models and databases of pathways:

Signaling pathways, concentrations and rates from Bhalla and Iyengar, 1999:

<http://www.ncbs.res.in/~bhalla/ltploop/> or <http://piris.pharm.mssm.edu/urilab/>

Genomics database of pathways: <http://www.genome.ad.jp/kegg/>

Cross-reference records of signalling molecules: <http://bioinfo.mshri.on.ca>

References

- Ahn, N.G., Seger, R. and Krebs, E.G. (1992) The mitogen-activated protein kinase activator. *Curr. Op. Cell. Biol.* **4**: 992-999.
- Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J.D. (1983). *Molecular biology of the cell*. Garland Publishing, Inc. New York.
- Bhalla, U.S. and Iyengar, R.I. (1999) Emergent Properties of Networks of Biological Signaling Pathways. *Science*. **283**: 381-387.
- Biddlecome, G.H., Berstein, G., and Ross, E.M. (1996) Regulation of phospholipase C- 1 by Gq and m1 muscarinic cholinergic receptor. Steady-state balance of receptor-mediated activation and GTPase-activating protein-promoted deactivation. *J Biol Chem* **271(14)**: 7999-8007.
- Bray, D. (1995). Protein Molecules as Computational Elements in Living Cells. *Nature* **376**: 307-312.
- Cobb, M.H, Boulton T.G. and Robbins D.J. (1991) Extracellular signal-regulated kinases: ERKs in progress. *Cell Regulation* **2**: 965-978.
- De Bernardi, M.A. and Brooker, G. (1996). Single cell Ca/cAMP cross-talk monitored by simultaneous Ca/cAMP fluorescence ratio imaging. *Proc. Natl. Acad. Sci. USA*. **93(10)**: 4577-4582.
- De Koninck, P., and Schulman, H. (1998). Sensitivity of CaM Kinase II to the Frequency of Ca²⁺ Oscillations. *Science* **279**: 227-230.
- Døskeland, S.O. and Øgreid, D. (1984). Characterization of the interchain and intrachain interactions between the binding sites of the free regulatory moiety of protein kinase I, *J. Biol. Chem.* **259**: 2291-2301.
- Garrington T. P. and Johnson G. L. (1999) Organization and regulation of mitogen-activated protein kinase signaling pathways. *Curr Opin Cell Biol.* **11(2)**: 211-8
- Hanson, P.I. and Schulman, H. (1992). Neuronal Ca²⁺/Calmodulin-dependent protein kinases. *Annu. Rev. Biochem.* **61**: 559-601.
- Hines, M. (1984) Efficient computation of branched nerve equations. *Int. J. Bio-Med. Comput.* **15**: 69-79.
- Huang, C-Y. F. and Ferrel Jr., J.E. (1996). Ultrasensitivity in the mitogen-activated protein kinase cascade. *Proc. Natl. Acad. Sci. USA* **93**:10078-10083.
- Lin, L.-L., Wartmann, M., Lin, A.Y., Knopf, J.L., Seth, A., and Davis, R.J. (1993). cPLA₂ is Phosphorylated and Activated by MAP Kinase. *Cell* **72**:269-278.
- MacGregor, R.J. (1987) *Neural and Brain Modeling*, Academic Press, San Diego.
- Mascagni, M.V. and Sherman, A.S. (1998). Numerical methods for neuronal modeling, in C. Koch and I. Segev (eds.), *Methods in Neuronal Modeling, 2nd edition*, MIT Press, Cambridge, MA, chapter 14.
- Michaelis, L., and Menten, M.L. (1913). *Biochem Z.* **49**: 333.
- Nishizuka, Y. (1992) Intracellular signaling by hydrolysis of phospholipids and activation of protein kinase C. *Science* **258**: 607-614.
- Ortiz, J., Harris, H.W., Guitart, X., Terwillinger, R.Z., Haycock, J.W. and Nestler, E.J. (1995). Extracellular Signal-Regulated Protein Kinases (ERKs) and ERK Kinase (MEK) in Brain : Regional Distribution and Regulation by Chronic Morphine. *J. Neurosci.* **15(2)**: 1285-1297
- Posner, I., Engel, M., and Levitzki A. (1992). Kinetic model of the epidermal growth factor (EGF) receptor tyrosine kinase and a possible mechanism of its activation by EGF. *J. Biol. Chem.* **267(29)**: 20638-20647
- Press, W.H., Flannery, B.P., Teukolsky, S.A., and Vetterling, W.T. (1988). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge.
- Sanghera, J.S., Paddon, H.B., Bader, S.A. and Pelech, S.L. (1990). Purification and Characterization of a Maturation-activated Myelin Basic Protein Kinase from Sea Star Oocytes. *J. Biol. Chem.* **265(1)**:52-57.
- Schaechter, J.D., and Benowitz, L.I. (1993). Activation of protein kinase C by arachidonic acid selectively enhances the phosphorylation of GAP-43 in nerve terminal membranes. *J. Neurosci.* **13(10)**: 4361-4371.
- Vanier, M.C., Bower, J.M. (in press) A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models. *J. Comput. Neurosci.*
- Wahl, M.I., Jones, G.A., Nishibe, S., Rhee, S.G., and Carpenter, G. (1992) Growth factor stimulation of phospholipase C- 1 activity. Comparative properties of control and activated enzymes. *J. Biol. Chem.* **267(15)**: 10447-10456.

List of Figures.

1. Reaction loops
2. Parameterization flowchart.
3. Effects of high substrate concentrations.

Box 1 Figure 1. Activation of PKC.

Box 1 Figure 2. PKC regulation: experimental and simulated. (Box 1)

Box 2 Figure 1. Models of enzymatic regulation

List of JAVA/Catacomb examples.

1. Forward reaction kinetics
2. Enzyme kinetics
3. Saturating substrate levels
4. PKC parameterization