A Higher Categories Approach for Stochastic Modelling of Biological Processes

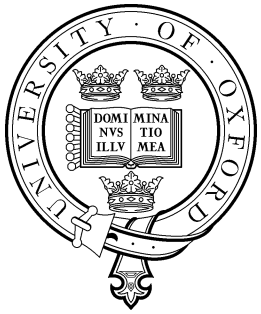

Katherine A. Casey
St. Catherine's College
University of Oxford

A dissertation submitted for the degree of
Master of Science in
Mathematics and the Foundations of Computer Science

Trinity 2015

Dominus illuminatio mea

Acknowledgements

I wish to thank my supervisor Jamie Vicary for his time, patience, and insight. Credit is also due to John Baez for his helpful input. Much of this work would not have been possible without the dedicated `globular` team, particularly the work of Krzysztof Bar and Caspar Wylie (and of Jamie, of course). I am grateful to Michael Casey for the many useful discussions and for his support, and to Christoph Dorn for his encouragement.

Abstract

We present a new way in which to model biological processes using graphical inputs and outputs. This method is grounded in higher category theory and can be viewed as an extension of the popular Petri net approach for biological modelling. We implement a stochastic process algorithm in `globular`, a rewriting program for higher categories. We use this to model a basic viral invasion process, following the membrane operations of the brane calculi of Luca Cardelli. Comparisons with the standard non-spatial modelling framework motivate a new chemical kinetics model for biochemical processes. Two potential kinetics models are proposed and discussed in detail.

Contents

Chapter 1

Introduction

The discipline of biology tends to rely heavily on different types of diagrams, accompanied by detailed qualitative descriptions, to describe processes. The motivation for the work put forth in this dissertation is as follows: Connecting mathematics to biological diagrams in an intuitive way is one of the most promising ways to establish a firm mathematical foundation for the field of biology. Such a foundation would secure a rigorous backbone for a subject which generally entertains mathematics on a case by case basis.

In creating this new foundation, it is important to conserve as much as possible from what is already accepted and useful in biology. To this end, our research is an extension of Petri nets, a tool for describing processes which already has widespread use in systems biology and which is compatible with quantitative data. By acknowledging that Petri nets are presentations of strict symmetric monoidal categories, we are able to harness the power of the categorical graphical calculus developed by the Quantum Group at Oxford. As this calculus has already been used to gain insight into quantum, thermodynamic, and classical processes [7, 14, 36], and as biological processes will necessarily fall into at least one of these areas [32], it is reasonable to expect this approach to be fruitful.

As the graphical calculus has been extended to higher categories, we can take advantage of the fact that process inputs and outputs can be two-dimensional structures, with the ultimate goal of having three-dimensional inputs and outputs. This allows spatial data to be included in our modelling approach. Additionally, through the use of the graphical reasoning engine `globular` [13], we can apply various dynamics to these processes. While deterministic or stochastic dynamics could be applied, in this work, we focus on the latter. Because our representation of processes is spatially rich, this motivates a new model for stochastic chemical kinetics,

one which can take advantage of a higher spatial resolution than standard methods. Two possible models are proposed and their physical implications explored.

1.1 Current Biological Modelling Methods

Before laying the groundwork for our modelling approach, we give a brief overview of current biological modelling methods. We introduce some of the most closely related modelling approaches as well as provide a more general scope of types of biological modelling. We delay discussion of Petri nets and graphical calculi until the next sections.

1.1.1 Differential Equations Modelling

Ordinary and partial differential equation modelling techniques are good for describing average behaviour, where the average may be over different time scales or different space scales. In the case of ordinary differential equations, the systems modelled are assumed to be well-mixed. Differential equations models also assume that species are continuous, rather than discrete, quantities [27, page 55].

1.1.2 Molecular Dynamics (MD) Modelling

Molecular dynamics modelling is concerned with the spatial trajectories of the molecular or atomic components of biological systems. It requires the calculation of the forces between atoms/molecules caused by chemical bonds, electrostatics and van der Waals interactions. Such detailed modelling is understandably computationally demanding, and much effort is spent on writing efficient algorithms and developing specialised hardware [39]. The aim is to see whether certain spatial arrangements of atoms/molecules occur more often on average than others; to predict such arrangements; and to see how these arrangements are affected under various conditions, such as the presence of different molecules [31]. Molecular dynamics modelling can be deterministic or stochastic, and while it was originally carried out with only classical laws, quantum approaches are now being undertaken [31].

1.1.3 Deterministic Modelling

Deterministic modelling is still appropriate for certain biological systems: there is enough variety in gene regulatory networks that some have stochastic behaviour while others have deterministic behaviour [5].

1.1.4 Stochastic Modelling

Stochastic systems behave differently than deterministic systems do in that given an initial state, one cannot predict the ending state of a stochastic simulation. Wilkinson notes: "The deterministic approach to kinetics fails to capture the discrete and stochastic nature of chemical kinetics at low concentrations. As many intra-cellular processes involve reactions at extremely low concentrations, such discrete stochastic effects are often relevant for systems biology models" [40, page 178]. Examples of biological processes exhibiting stochastic behaviour include gene regulation, human brain dynamics [29, page 2], cell polarization [29], cell differentiation [33], and circadian rhythms [38].

Stochastic modelling is often used to bridge the gap between the averaged differential equations picture and the highly detailed molecular dynamics models. This use of stochastic modelling is regarded as mesoscale modelling.

1.1.4.1 Gillespie's Algorithm

We briefly introduce the Gillespie algorithm approach; more discussion on this method can be found in section 3.1. In 1992, Gillespie provided a rigorous derivation for the chemical master equation [19]. His derivation is based on the assumption that the reactant molecules of any chemical reaction in a well-mixed environment at thermal equilibrium behave in the same way that particles of an ideal gas behave. These particles are assumed to be spherical in shape. It is assumed that when the distance between the centres of these two particles is less than or equal to the sum of their radii, the molecules will "collide;" whether the collision results in a reaction is determined by the rate for that reaction. Gillespie developed a stochastic modelling procedure based on this chemical master equation (CME) model, which has been subsequently extended in various ways [20]. Because it is based on standard chemical kinetics assumptions and is relatively computationally efficient when compared to other algorithms, "Gillespie's method is the method of choice in many situations" [26, page 108]. There are many different modifications of Gillespie's "Discrete Stochastic Simulation Algorithm," including an inhomogeneous (spatial) version; see [23, 37] for examples of applications of this method. This version involves breaking the system into smaller three-dimensional components called voxels, where, while collectively the voxels do not represent a well-mixed situation, each individual voxel is assumed to be well-mixed. The drawback with this algorithm is that there is a limit to how small the voxels can be: two reactants

for a chemical process need to be present in the same voxel for the reaction to occur. This means the voxel has to be big enough to include both molecules. Our proposed model does not enforce such limitations of spatial resolution.

1.1.5 Membrane-Based Modelling

1.1.5.1 Brane Calculi

Luca Cardelli's brane calculi [16] describes basic membrane operations, such as mating and fusion, using a process calculi. These operations can be used, for example, to describe the life cycle of a virus. We adapt these membrane operations into our modelling framework in 2.2.

1.1.5.2 Membrane Computing

Membrane computing, whose models are referred to as "P Systems", is an abstraction of the compartmentalization via membranes of biological systems. It thus falls under the discipline of Natural Computing. The guiding principle is that only certain types of computations or "rules" can be carried out in certain compartments. This approach is used to model not only biological systems but any other system capable of being expressed in this format [3].

While work has also been done with Petri nets in membrane computing, the inputs for the processes are still limited to being one-dimensional [6, Chapter 15]. Stochastic P systems have also been developed [6, Chapter 18].

Also of note is the chemtainer calculus, which is similar to many membrane computing approaches in its compartmentalisation approach, and is an extension of the brane calculi [17].

1.2 Foundations for Our Modelling Approach

The target application of the work done thus far is biological systems at the molecular level. Such systems are inhomogeneous: cells are not well-mixed bags of chemicals, but rather spatially arranged sets of components. Including this spatial characteristic allows for more realistic behaviour, particularly in situations where certain chemicals are present in small quantities. Additionally, biological systems are almost never in chemical equilibrium, so we will exclude this case as well. The proposed modelling scheme is stochastic, continuous in time, and discrete in species numbers.

It is important to keep in mind that Petri nets and category theory diagrams are mathematical structures and are therefore not restricted to biological applications. Our work relates features of biological systems to features of mathematical structures.

1.2.1 A Spatial and Local Approach

Particularly due to advances in microfluidics methods such as organs-on-a-chip technology, and more recently, high-speed multiple-mode mass-sensing [35], spatial data is now available. Part E of Figure 1.1 shows an example of a lung-on-a-chip, a microfluidics device which models basic lung processes and which can be imaged in vivo [15, 4].

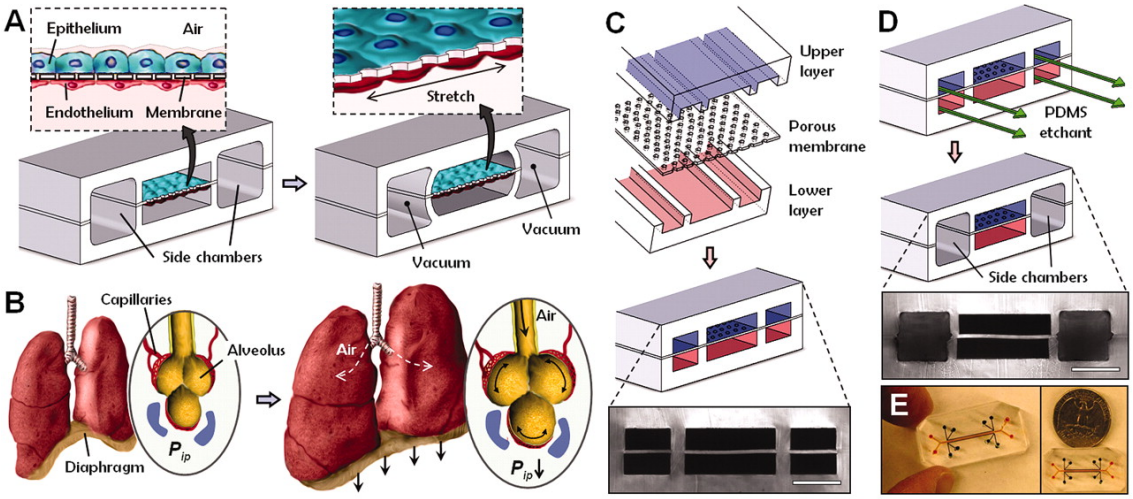

Figure 1.1: Lung-on-a-Chip Design and Device, produced by the Wyss Institute [25]

To emphasize what we mean by local, consider the following example: suppose one wanted to model the initial step of viral invasion of a cell, that is, a virus outside of a cell acts as input, and the output of the process is the virus being within the cell. If you include the whole cell in your input, as in Figure 1.2, then the operation of virus invasion is no longer local. This is because in a sense the process has to "see" the entire cell for the reaction to occur. A local, and thus more accurate, version of viral invasion would take the virus and only a *part* of the cell membrane as input (Figure 1.3).

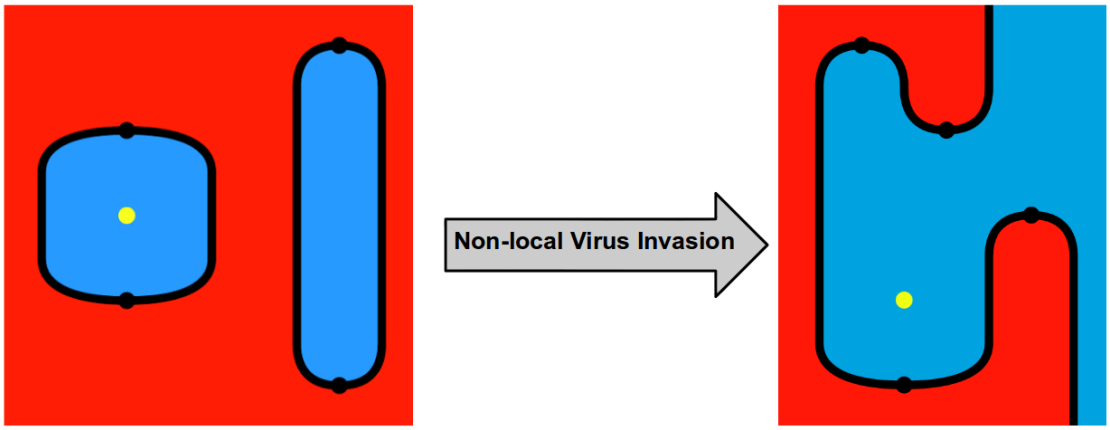

Figure 1.2: A Non-Local Version of Virus Invasion: note that the entire cell must be present for the virus to enter

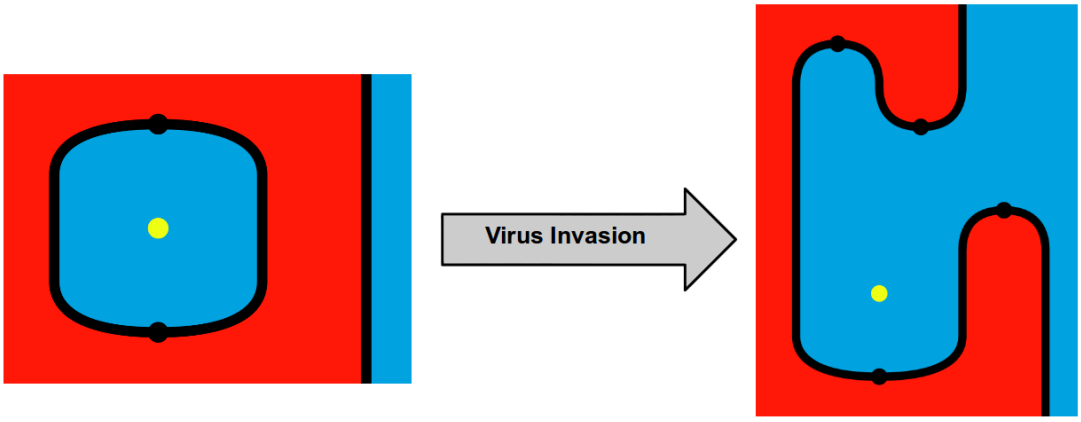

Figure 1.3: A Local Version of Virus Invasion

1.2.2 Petri Nets

Petri nets are a useful way to visualise a system's dynamics, provided that all aspects of the system behaviour can be specified as processes with particular inputs and outputs. The formal definition of a Petri net is:

Definition 1.2.1. *A **Petri net** consists of a set S of **species** and a set T of **transitions** (or processes), together with a function*

$$i : S \times T \to \mathbb{N}$$

*describing how many copies of each species are in the **input** for each transition, and a function*

$$o : S \times T \to \mathbb{N}$$

*describing how many copies of each species are in the transition's **output**.*

(Definitions adapted from [11, page 11].)

As an example of a Petri net, we follow John Baez's example [11] of modelling the population dynamics of rabbits and wolves. Suppose we have two species, rabbits (white) and wolves (black) and three processes with associated colors:

(i) asexual rabbit reproduction, that is, one rabbit can reproduce on its own (Blue)

(ii) rabbit killing by a wolf on its left (Red), and

(iii) rabbit killing by a wolf on its right (Yellow)

For each process we have:

(i) i(rabbits, rabbit reproduction) $= 1$; o(rabbits, rabbit reproduction) $= 2$
i(wolves, rabbit reproduction) $= 0$; o(wolves, rabbit reproduction) $= 0$

(ii) i(rabbits, left killing) $= 1$; o(rabbits, left killing) $= 0$
i(wolves, left killing) $= 1$; o(wolves, left killing) $= 1$

(iii) i(rabbits, right killing) $= 1$; o(rabbits, right killing) $= 0$
i(wolves, right killing) $= 1$; o(wolves, right killing) $= 1$

Figure 1.4 shows an example of a Petri net with these species and processes.

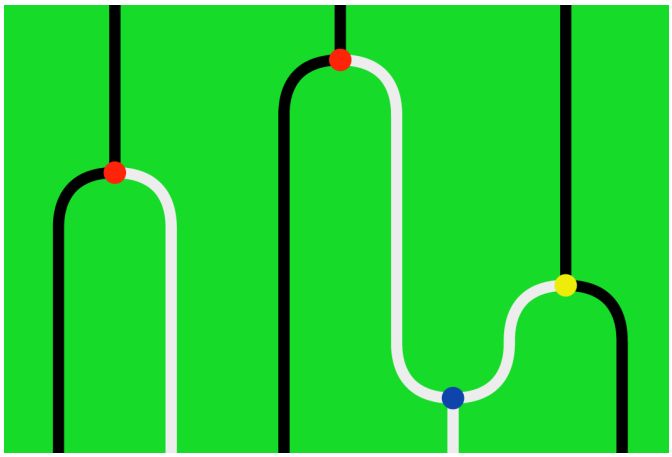

Figure 1.4: A Petri Net for Rabbit and Wolf Population Dynamics: We begin with a row of species in the order: wolf, rabbit, wolf, rabbit, wolf. We may interpret the green background loosely as grass. Taking time to be in the vertical direction, the first process event to occur is rabbit reproduction. Following this, a rabbit is killed by a wolf on its right, then another rabbit is killed by a wolf on its left, and finally, the last rabbit is also killed by a wolf on its left, leaving the system with three wolves and no possible processes which may be carried out.

1.2.2.1 Stochastic Petri Nets

Petri nets have been used to model stochastic biological processes at least since the work of Goss and Peccoud in 1998 [21]. For a detailed reference on modelling stochastic processes with Petri nets, see [40]; in this reference it is noted that a stochastic Petri net is "simply a convenient mathematical and graphical representation of a stochastic kinetic process" [40]. Formally, we have:

Definition 1.2.2. *A **stochastic Petri net** is a Petri net together with a function*

$$r : T \to (0, \infty),$$

*which assigns a **rate constant** to each transition.*

Much work with stochastic Petri nets and biological systems has been done already. The focus, however, has been on problems in chemical kinetics, on fitting experimental data, and on processes with inputs/outputs of only one-dimension [27, page 149]. While Petri nets do allow for a spatial separation of processes and species, most stochastic Petri net models keep the processes and species type separate but do not spatially separate the individual species members. In this approach, the species members are all referred to as "tokens." Because the tokens reflect only the

number of species and not the spatial arrangement of the species, any model using tokens is operating under the assumption that the species members are well-mixed. We discuss more about the relationship between our modelling approach and the well-mixed approach in Section 3.1.3.

Petri nets are often used to model concurrent systems, but we are not concerned with concurrency[1] in our model. Because we will be working with semi-strict categories, only one event can occur at any given time. It would be very unlikely for two events to occur at the same time in a continuous-time stochastic process anyway. However, this could be modelled by varying the time values by a precision level at least one order of magnitude finer than the dominant time scale of the model. Thus, up to whatever approximation we choose, we can distinguish the time of two events such that they are separate.

1.2.3 Petri Nets are Presentations of Categories

As stated in [11], "a Petri net is a presentation of a strict symmetric monoidal category that is freely generated." We briefly review the meaning of these terms.

Recall that a category **C** consists of a collection of objects, a collection of morphisms between every two objects (note that the collection may be empty), a composite for every pair of morphisms which share a common intermediate object, and an identity morphism for every object, where associativity of morphisms and composition with the identity hold.

In addition to the above, a monoidal category **C** has a tensor product functor $\otimes : \boldsymbol{C} \times \boldsymbol{C} \to \boldsymbol{C}$, a unit object I, an associator natural isomorphism $\alpha_{A,B,C}$ for any three objects A, B, C; a left unitor natural isomorphism λ_A for each object A; and a right unitor natural isomorphism ρ_A for each object A, where the data of the category satisfy the triangle and pentagon equations for all objects.

A monoidal category is *strict* if all isomorphisms $\alpha_{A,B,C}, \lambda_A$ and ρ_A are identities. A monoidal category is symmetric when it has a braiding, which is a natural isomorphism, $A \otimes B \xrightarrow{\sigma_{A,B}} B \otimes A$ satisfying the hexagon equations, and additionally where this braiding satisfies $\sigma_{B,A} \circ \sigma_{A,B} = id_{A \otimes B}$ for all objects A and B.

Finally, a category is free if it has a finite number of generators: a finite number of generating objects (species) and a finite number of generating morphisms (processes). Taking a tensor product of generating objects gives rise to a new object, and

[1] Here, as in [27], the definition of concurrency has been restricted to parallel executions of processes. This notion of concurrency appears to be the standard one when discussing Petri nets in biological modelling.

taking a composition of morphisms gives rise to a new morphism. Fitting in with our Petri net definition, set S will be the set of these species, and set T will be the set of morphisms, where elements of T take in tensor products of objects as input.

In the next section we discuss how to extend the species and processes (Petri net) connection with categories to higher dimensional categories in the context of our higher-dimensional rewriting tool, `globular`.

1.3 `globular`

As mentioned earlier, our simulations are carried out in `globular`, a reasoning engine which employs a graphical logic grounded in higher category theory. In this section we introduce the basic notions from higher category theory needed to get a sense of how `globular` and our modelling approach rely on higher categories. At some point in the near future, `globular` will be accessible on the internet. We emphasize that higher categories are being used because they allow us to have higher-dimensional input structures. 3-categories in particular allow us to model processes as sequences of two-dimensional pictures.

1.3.1 A Graphical Logic

By a graphical logic we mean a system of composing pictures using certain rules. Processes take pictures as inputs and produce pictures as outputs, and one cannot compose pictures unless the adjacent boundaries match, that is, strong typing is enforced. Interactions between pictures only occur for adjacent pictures, thus this logic captures the local nature of biochemical reactions.

1.3.2 Higher Categories

A higher category is a higher-dimensional abstraction of a category. In this language, a 1-category is a strict category with objects and one kind of morphism, as discussed above. A 2-category has in addition, a second kind of morphism, so its data consists of objects, 1-morphisms and 2-morphisms. Analogously, a 3-category will have objects and three types of morphisms. In general, an n-category will have objects and n kinds of morphisms. Higher categories allow for more ways to compose morphisms, for example, in 2-categories, there is both horizontal and vertical composition of morphisms.

For the 2-category case, the formal details are as follows:

Definition 1.3.1. A bicategory **C** (as defined in [24]) consists of a collection of objects, Ob(**C**); for every two objects A, B, a category $\boldsymbol{C}(A, B)$, with objects called 1-morphisms f drawn as $A \xrightarrow{f} B$, and morphisms μ called 2-morphisms drawn as $f \overset{\mu}{\Rightarrow} g$; for any two composable 2-morphisms in $\boldsymbol{C}(A, B)$ of type $f \overset{\mu}{\Rightarrow} g$ and $g \overset{\nu}{\Rightarrow} h$, an operation called vertical composition given by their composite as morphisms in $\boldsymbol{C}(A, B)$, denoted $f \overset{\nu \cdot \mu}{\Longrightarrow} h$; for any triple of objects A, B, C, a functor $\circ : \boldsymbol{C}(A, B) \times \boldsymbol{C}(B, C) \to \boldsymbol{C}(A, C)$ called horizontal composition; for any object A, a 1-morphism $A \xrightarrow{id_A} A$ called the identity 1-morphism; for any 1-morphism $A \xrightarrow{f} B$, invertible 2-morphisms $f \circ id_A \overset{\rho_f}{\Rightarrow} f$ and $id_B \circ f \overset{\lambda_f}{\Rightarrow} f$ called the left and right unitors satisfying a set of naturality conditions for all $f \overset{\mu}{\Rightarrow} g$; for any triple of 1-morphisms $A \xrightarrow{f} B$, $B \xrightarrow{g} C$, and $C \xrightarrow{h} D$, an invertible 2-morphism $(h \circ g) \circ f \overset{\alpha_{h,g,f}}{\Longrightarrow} h \circ (g \circ f)$ called the associator, such that for all $f \overset{\mu}{\Rightarrow} f'$, $g \overset{\nu}{\Rightarrow} g'$ and $h \overset{\sigma}{\Rightarrow} h'$ we have $(\sigma \circ (\nu \circ \mu)) \cdot \alpha_{h,g,f} = \alpha_{h',g',f'} \cdot ((\sigma \circ \nu) \circ \mu)$. This data must satisfy the triangle and pentagon equations.

By Theorem 8.2 from [24], we have: A monoidal category is the same as a bicategory with one object.

Additionally we have that:

Definition 1.3.2. A bicategory is strict, or a 2-category, when the 2-morphisms α, λ and ρ are the identity at every stage.

Thus a 2-category with one object is also a monoidal category. Since the object is unique, in this situation, it is standard to refer to the 1-morphisms as objects and the 2-morphisms as morphisms. In this sense our "objects" or "species" are now of a higher dimension than before.

Higher categories can be freely generated in a similar way as before, by a finite number of generating objects and a finite number of generating 0-, 1-, ..., n-morphisms.

We now have all the features of the categories of which Petri nets are a presentation except for the symmetric braiding.

Having a braiding is important for us when we want to model swapping of positions, and ultimately, the well-mixing of our biochemical systems, so that we can compare our methods to those of well-mixed models. The braiding allows us to interchange points, however it is not essential that this braiding be symmetric. Including a braiding requires us to have a 3-process and to thus be in a 3-category or an n-category such that $n \geq 3$.

With higher categories we can describe a more realistic picture since we are not restricted to modelling everything as point masses. With 2-categories for example, we can model two-dimensional regions as 2-cells.

For our purposes, the basic intuition for how 2-categories work can be used for n-categories in `globular`; we omit the formal technicalities.

1.3.3 Graphical Calculus for Higher Categories in `globular`

In `globular`, we have several types of "cells." The relationship between these cells and the objects and morphisms described above is as follows: 0-cells correspond to objects, and all higher dimensional cells correspond to the morphisms with the same number prefix; for example, 2-cells correspond to 2-morphisms. However, as we mentioned earlier, when a 2-category has only one object or 0-cell, then the 1-morphisms or 1-cells will be referred to as objects.

Figure 1.5 shows examples of 0-cells later used in our rabbit/wolf and membrane models. In the higher-dimensional cells, these 0-cells will become "background regions."

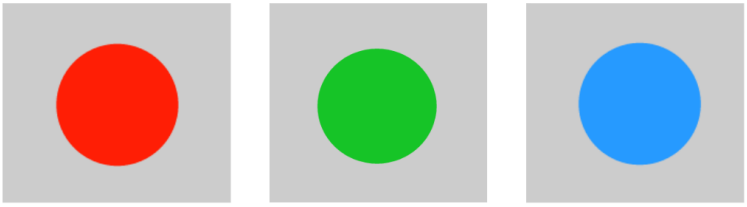

Figure 1.5: 0-cells from `globular`

Figure 1.6 shows 1-cells from these models. For each n-dimensional cell, the appropriate (n-1)-dimensional sources and targets need to be made. Since 1-cells are 1-morphisms, they are created by selecting an input or "source" and an output or "target," both of which must be 0-cells. Note that 1-morphisms are identified by the single dot in the middle of the 1-cell picture.

Figure 1.6: 1-cells from `globular`

Figure 1.7 shows 2-cells. Since these are 2-morphisms, the sources and targets will be 1-morphisms. Note that here the dots of the 1-cells have been raised to a higher dimension and have thus become lines. The dots now identify the 2-morphisms.

Figure 1.7: 2-cells from `globular`

Figure 1.8 shows an example of one of our 3-cells; we refer to this as a 3-process or 3-morphism. It takes an input of a 2-cell and outputs a 2-cell, that is, it maps one two-dimensional picture to another. Due to strong typing, the source of the source picture must match the source of the target picture, and the target of the source picture must match the target of the target picture. If you placed these two two-dimensional pictures with one in front of the other, the lines from the 2-cells would connect to form two-dimensional regions between the pictures.

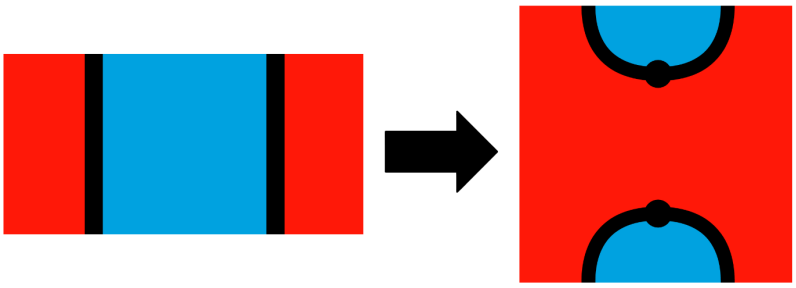

Figure 1.8: A 3-cell (3-process) from `globular`

It is possible to create cells of higher dimensions, but for our modelling needs 3-cells and below will suffice.

We implemented a stochastic process, the algorithm for which will be discussed in Chapter 2, in `globular` such that for each process of a given model, a rate may be assigned and used according to our stochastic process algorithm procedure. To make a simulation in `globular`, one must create the desired processes and the

desired initial state, which consists of the various species objects and their spatial arrangement. After this, one must select the stochastic process option and specify which n-processes are to be used. For each process, a numerical rate can be specified. We note that the input of a process needs to be compatible with the processes selected if a history is to be made. For example, for a situation in which the system state is a horizontal composition of 1-cells, the identity must be taken before a 2-process can be applied because the 2-process needs input species to be represented as lines rather than dots. See Figure 1.9 for an example of how to provide the proper setup for stochastic modelling with 2-processes. Note that application of 2-processes is the only case in which we need to perform a special modification step on the input.

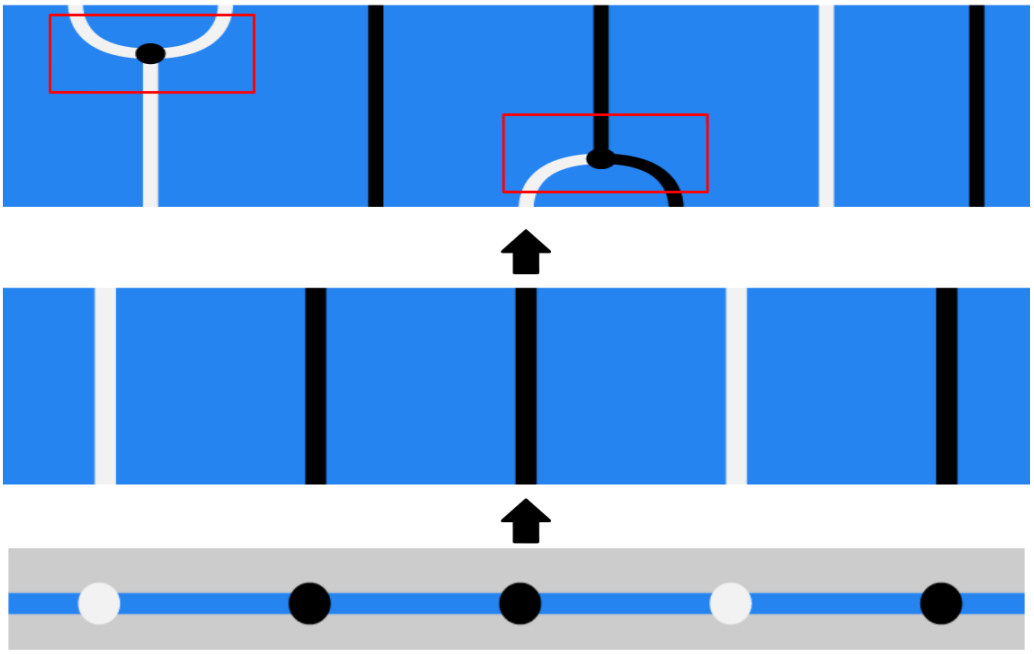

Figure 1.9: Setting Up for a Stochastic Process with 1-cell species: After the 1-cells have been horizontally composed in a line, the identity operation must be taken so that a 2-cell is generated. This then allows 2-processes, such as the ones identified by the red rectangles, to be attached to the input diagram, forming a history.

1.3.4 Modelling Technicalities in `globular`

It should be noted that strong typing occasionally means that a bit of thought needs to be put into making the before and after pictures of a process since the boundaries must match up properly. While initially this may make including certain

processes in a model a bit difficult, with practice the matching constraint ceases to be an issue. As a result of the category-based approach, only one "dot" can occur at a single height. This can make pictures expand vertically much more rapidly than horizontally when running stochastic processes. Including interchanges in the stochastic processes can help widen the picture, but not in a deterministic way. Currently, `globular` is overly-sensitive to certain setup details, and this makes matching less than straightforward. For example, Figure 1.10 shows a situation where one would think that the membrane straightening process from Figure 1.11 could be applied. However, because of the dot on the right-hand side in Figure 1.10, `globular` sees the height difference between the two dots on the curved pieces of membrane as being 2 units tall rather than 1 unit tall. Since the height between said dots is 1 unit in the process definition in Figure 1.11, the membrane straightening process cannot occur here. Work is currently being done to relax the notion of matching in a reasonable way.

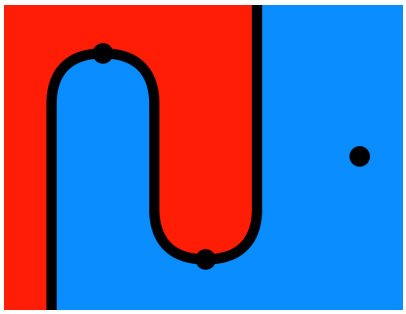

Figure 1.10: Matching Prevention due to Height: The dot on the right-hand side makes the distance between the two curved parts of the membrane equal to 2, rather than 1, as in the input for the membrane straightening process (Figure 1.11).

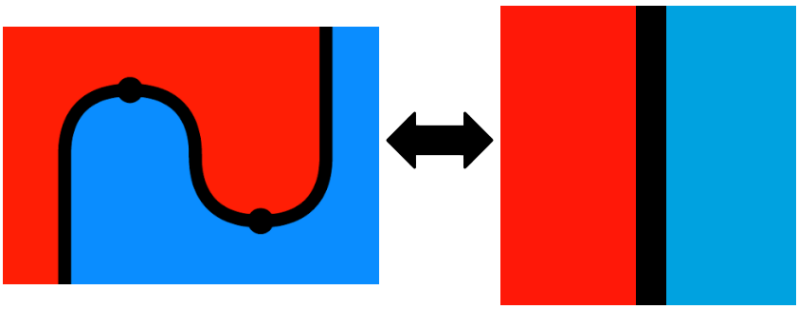

Figure 1.11: Membrane Straightening (Forward Arrow) and Buckling (Backward Arrow)

Another consequence of `globular`'s setup is that when given an input as in Figure 1.12, while only one cell is present and therefore there should only be one way in which the cell could divide, because of the dot on the right-hand side, `globular` will see two different opportunities (as marked) for the mitosis process (Figure 1.13) to occur. This will result in cell division having a higher rate than intended. A resolution to this issue is also in progress.

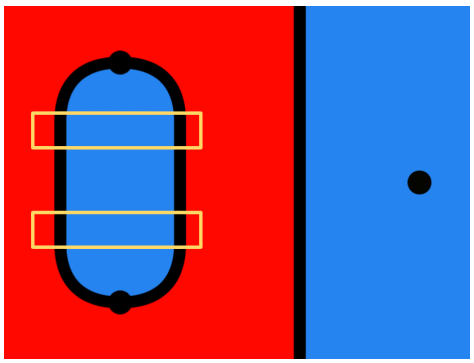

Figure 1.12: A situation in which `globular` recognizes two ways (marked by the yellow rectangles) for the mitosis process to occur

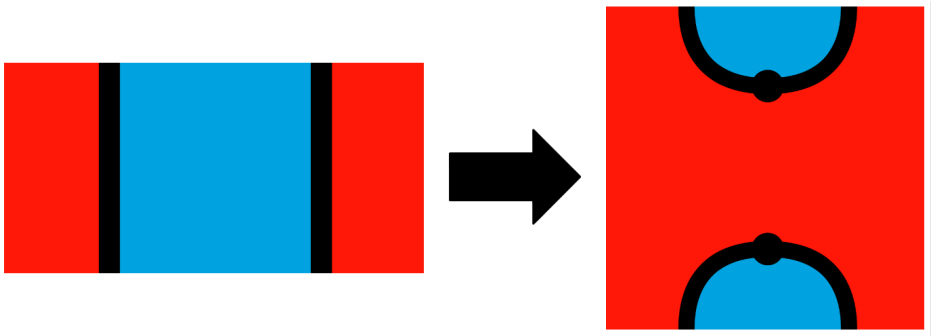

Figure 1.13: Representation of Mitosis as a 3-process in `globular`

Chapter 2

Stochastic Modelling

2.1 Stochastic Process Algorithm

Our input will be a set of objects with a spatial configuration. In the one-dimensional case, this will effectively be a row or string of objects which we represent by letters. In the two-dimensional case, the input will be a two-dimensional picture, a composition of 2D objects.

2.1.1 Markov Processes

We are simulating the stochasticity by using a continuous time Markov process, which is a memoryless process, that is, only present events, not past events, will affect future events. Stated another way, the likelihood of an event occurring now is the same as it was at an earlier time regardless of what has happened in between these times. This seems reasonable since physical systems are local in time and space: since processes only take in a certain input, how that input came to be is irrelevant to the execution. Additionally, if the past were relevant, the system would need to store this information; if this does occur in biological systems, then we will at least begin with the simplifying approximation that the history is not important. Modelling time as continuous is the most realistic way to include time, and given that we do not expect major computational power to be needed, the main advantage of using discrete time is not applicable.

As stated in [40], if X is a random variable, then it is memoryless if:

$$P(X(t+dt) = x | \{X(t) = x(t) | t \in [0,t]\})$$
$$= P(X(t+dt) = x | X(t) = x(t)), \forall\ t \in [0,\infty), x \in S$$

More explicitly, if X is a nonnegative random variable, then it will be memoryless if:

$$P[X > s+t | X > t] = P[X > s]\ \forall\ s,t \geq 0$$

For each process rate, there is an associated probability distribution, from which we will sample time values. This probability distribution is an exponential distribution because that is the only possibility for a continuous probability distribution which is memoryless [22, page 140]. More specifically, if r is the rate of the process, then $F(x) = re^{-rx}$ is the probability distribution for that process. The expectation value for this distribution is $\frac{1}{r}$, which is also the standard deviation. A higher rate will result in an increased likelihood of a lower time being sampled, that is, events of a process with a higher rate are more likely to occur sooner than events of a process with a lower rate. Figure 2.1 illustrates this difference.

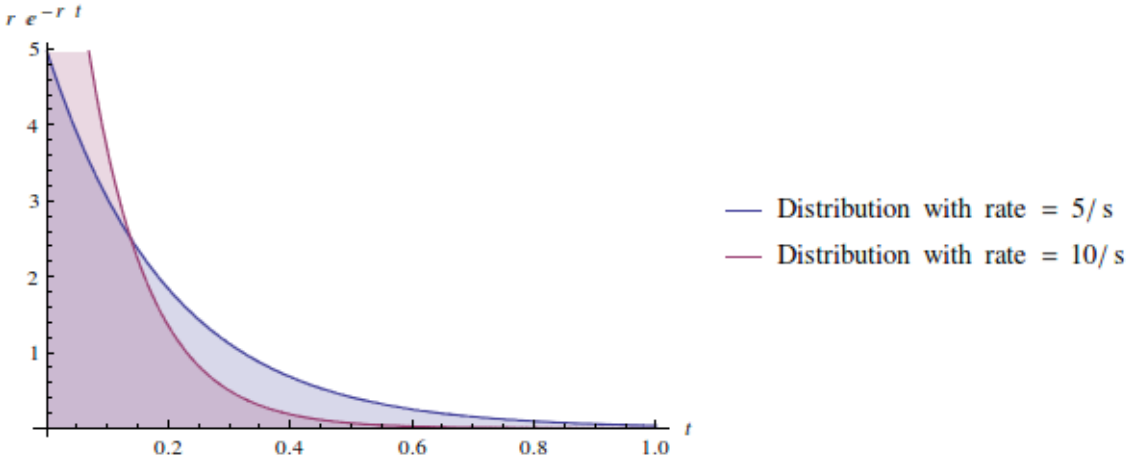

Figure 2.1: Comparison of Two Negative Exponential Distributions. Process A has rate 10/s, while process B has rate 5/s. Note that since process A has a higher rate, a greater amount of the distribution occurs at low time values.

2.1.2 One-dimensional Model

The simplest possible model will have a one-dimensional input. The model is not expected to be realistic since data from the other two dimensions is missing, but it can still provide some useful intuition. Recall the rabbit/wolf system from Chapter 1 in which we have two species, rabbits and wolves (Figure 2.2), and three processes: rabbit reproduction, rabbit killing by a wolf on the left and rabbit killing by a wolf on the right (Figure 2.3). Our rabbit and wolf model will be a monoidal 2-category with one 0-cell, which we will call “grass”, with two 1-cells, the rabbits and wolves, and three 2-cells, our processes. We will not include a braiding for this model, so no interchange moves will be allowed by default.

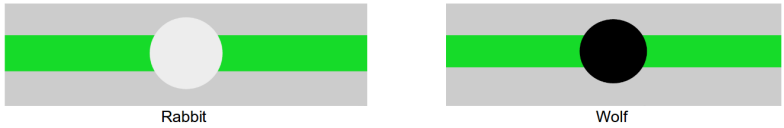

Figure 2.2: A Rabbit (white) and a Wolf (black) in Our One-Dimensional Model

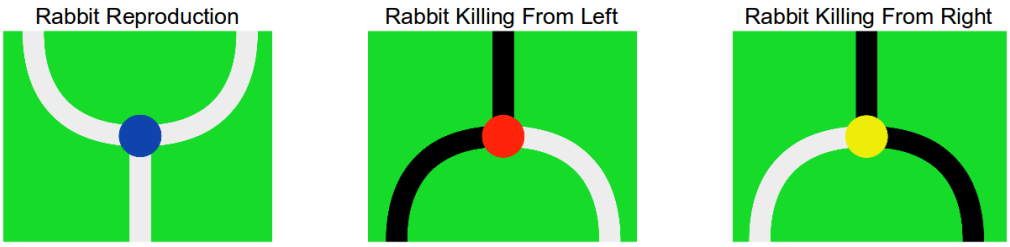

Figure 2.3: 1D Model Processes: Rabbit Reproduction (blue), Rabbit Killing from Left (red) and Rabbit Killing from Right (yellow)

This is a "one-dimensional" model because we are treating our species as point masses, or letters: R for rabbit; W for wolf. Figure 2.4 gives an example of a possible input for this model.

Figure 2.4: Sample Input for 1D Model: Wolf, Rabbit, Wolf, Rabbit, Wolf

After receiving the input, the next step in the algorithm will be to identify which of our processes can occur given this initial state, and for each process, where that process can occur. Note that only adjacent R,W pairs count as valid inputs for rabbit killing; this is an example of how spatial considerations are taken into account by the model. Figure 2.5 shows which processes would be identified given the input in Figure 2.4.

Each process is assigned a rate, and this rate is then fed into the rate parameter in the negative exponential distribution. For each process instance identified, a time is sampled from that process's associated probability distribution. Figure 2.6 shows an example of how times could be assigned for our rabbit/wolf model.

The event with the lowest time will be the event executed next. Given the times in Figure 2.6, the least event time is $t = 0.15$ seconds, corresponding to

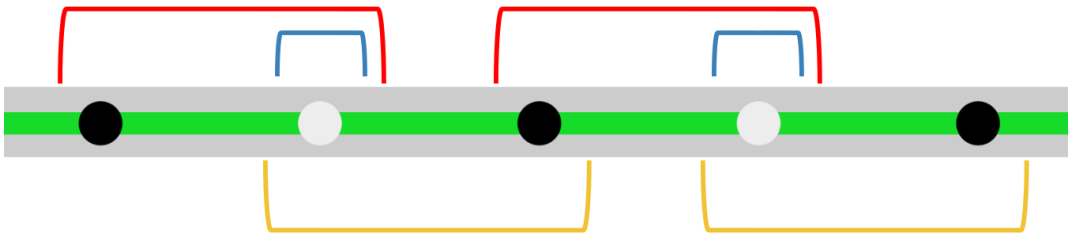

Figure 2.5: Possible Processes for Sample Input, Identified by Process Color

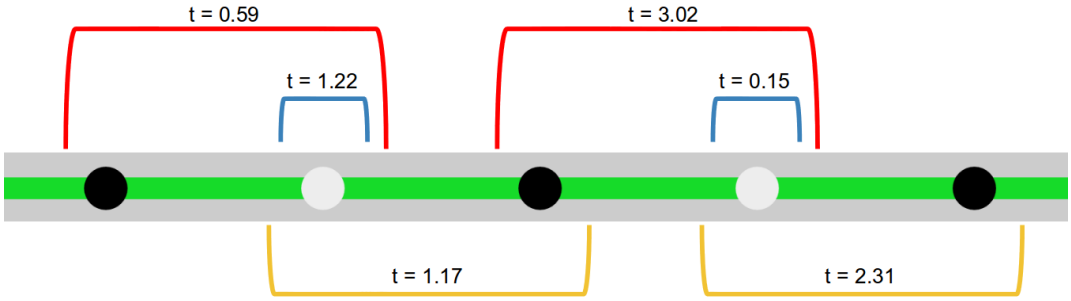

Figure 2.6: Possible Event Times (in Seconds) for Rabbit/Wolf Processes

the reproduction of the right-hand rabbit. This event will therefore be executed (provided that the input is correctly modified as discussed earlier), yielding the state in Figure 2.7.

Figure 2.7: System State after First Event

There are two ways in which the algorithm could proceed. The current implementation in `globular` uses the straightforward method of repeating the above sequence of steps, replacing the initial input with the current state. An alternative version, dubbed the "time conservation" method keeps the time values of all the events which were unaffected by the previous event. We will discuss the relationship between these two approaches in section 2.4.3.

We reproduce Figure 1.4 as Figure 2.8 to illustrate a history of the rabbit/wolf population produced by these dynamics.

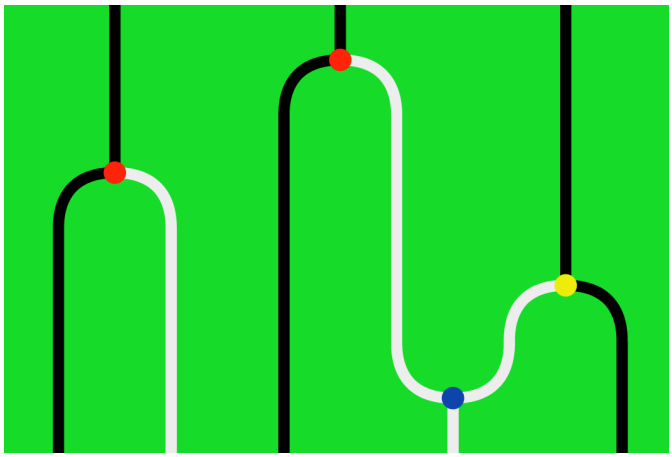

Figure 2.8: A History of Rabbit and Wolf Population Dynamics: We begin with a row of species in the order: wolf, rabbit, wolf, rabbit, wolf. We may interpret the green background loosely as grass. Taking time in to be in the vertical direction, the first process event to occur is rabbit reproduction. Following this, a rabbit is killed by a wolf on its right, then another rabbit is killed by a wolf on its left, and finally, the last rabbit is also killed by a wolf on its left, leaving the system with three wolves and no possible processes which may be carried out.

2.1.3 Steady-states

Some basic intuition about the system can be gleaned even in this simple model. For example, since wolves are not allowed to reproduce, the number of wolves in this system will always remain constant. This factor then influences the system to reach two different steady-states: one in which all the rabbits are eaten by the wolves, and one in which rabbit reproduction dominates rabbit death so much that there will be an explosive population of rabbits. In the first case, we would expect the rabbit killing rate to be higher than the rabbit reproduction rate, while in the second case, the opposite would probably be true. The more "interesting" behaviour will thus occur when the reproduction and death rates are similar. In the breeding situation, the simulation will never stop since rabbits can continue reproducing as long as they exist. However, in the killing dominated situation, the simulation will stop once no more processes can be performed on the system, that is, when all the rabbits are gone.

2.1.4 Stochastic Process Algorithm

For reference, here is a basic outline of the Stochastic Process Algorithm steps:

(i) Obtain initial state.

(ii) Based on the initial state, identify all possible processes whose inputs are present in the initial state.

(iii) Assign a time value for each possible process; this time value should be sampled from the negative exponential distribution with parameter equal to the rate constant of the process being considered.

(iv) Identify the lowest time value.

(v) Execute the event with the lowest time value. If there are multiple events with the same least time value, pick one at random.

(vi) Repeat (i)-(v) on the new state of the system

Time-Conservation Version:

(i) Execute steps (i)-(v) as above.

(ii) Save all time values for process events which are still valid in the new system state; this should be all processes whose input instances did not include any of the instances in the input just consumed in the executed process. Repeat steps (i)-(v) on the new state, only sampling for time values for process instances which do not yet have a time assignment.

2.2 Models of Membrane Processes

Higher categories allow you to describe different regions of space, so we can use them to describe simple processes involving membranes, like the processes in [16], which we model below. Some of these processes have also been implemented in the P-systems modelling framework [6, chapter 14]. Membranes play a vital functional role in cellular dynamics: the efficient production of ATP, the molecular form of energy, in the mitochondria would not occur without its double-membrane structure; quick transport of important molecules from one part of the cell to another would not be possible without transport by vesicles, which are membrane-bound compartments for holding small molecular cargo; and many proteins are modified through a membrane-mediated process in the Golgi apparatus outside the nucleus. In addition, by definition all organelles in the cell are membrane-bound: for a basic schematic refer to Figure 2.9.

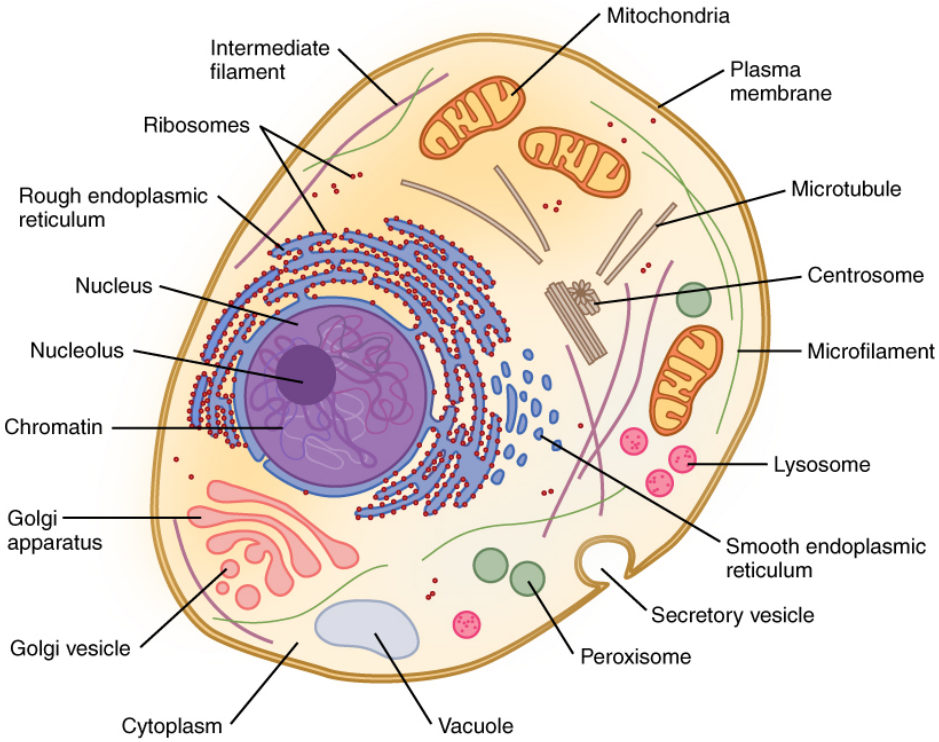

Figure 2.9: An Animal Cell [1]

As a brief terminology note: "cytosol" is the fluid within the cell, and the "extracellular fluid" is the fluid outside of a cell. We color the cytosol blue, and the extracellular fluid red in our cell process diagrams.

One simple cell process is mitosis, or cell division, referred to as "mito" in [16]. We will put all process names assigned in [16] in double quotes when first introducing them. See Figure 2.10 for confocal fluorescence microscopy images of input and output to a mitosis event. Figure 2.11 shows our representation of mitosis.

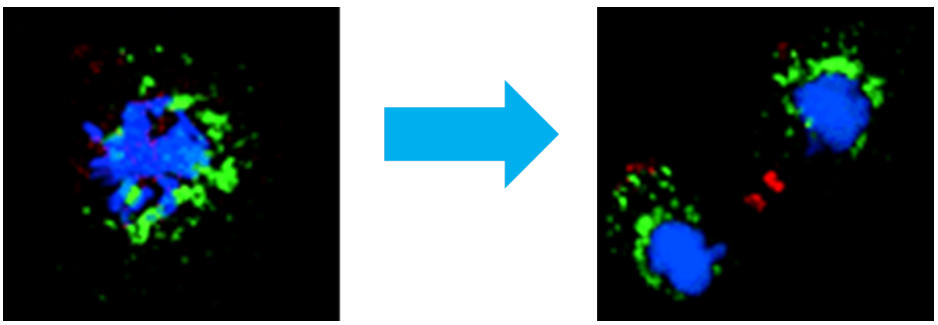

Figure 2.10: Mitosis, images from [2]

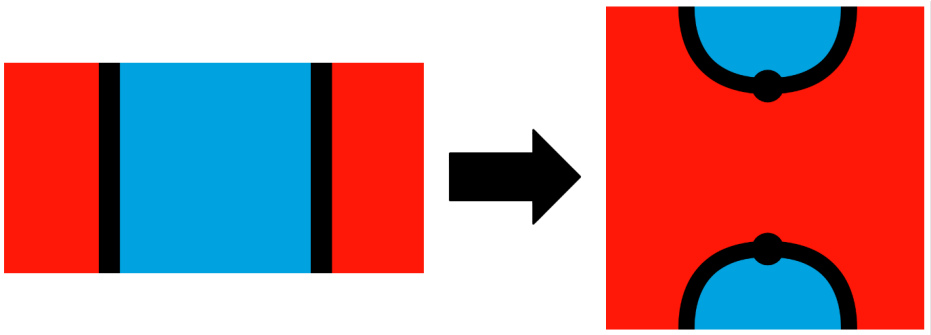

Figure 2.11: Representation of Mitosis as a 3-process in `globular`

Figure 2.12 shows our representation of the reverse process, mating, or "mate." Henceforth we shall use double-headed arrows to indicate a process and its reverse in the same image.

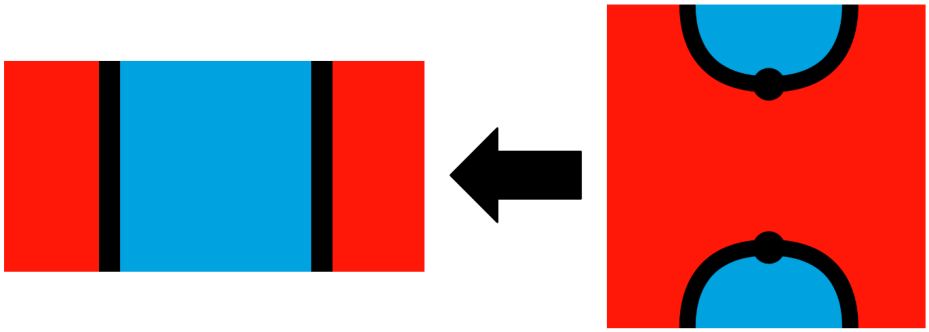

Figure 2.12: Representation of Mating as a 3-process in `globular`

Figure 2.13 shows our representation of cell generation ("froth") and its reverse process, apoptosis, that is, cell death ("fizz").

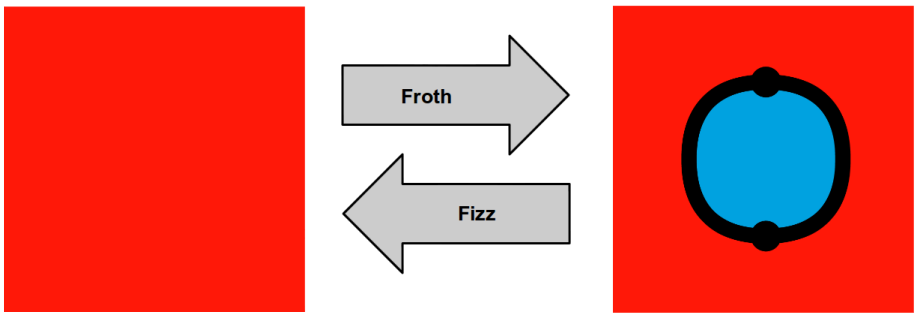

Figure 2.13: Representation of Cell Generation/Cell Death as a 3-process in `globular`

Figure 2.14 shows our representation of pinocytosis, often called 'cell drinking' ("pino"), and the reverse process, exocytosis of fluid.

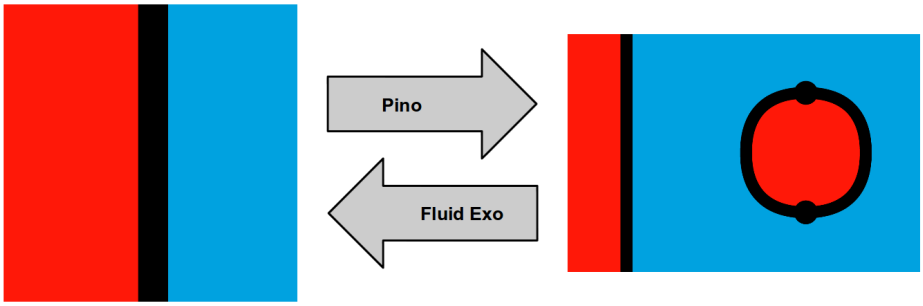

Figure 2.14: Representation of Pino/Fluid Exocytosis as a 3-process in `globular`

We note that as a consequence of the local nature of our graphical logic, we need a second pair of processes so that pino/fluid exo can occur on the right-hand edge of a cell (recall that a cell's interior is blue in our representation). Figure 2.15 shows this alternate pino/fluid exo pair of processes. We did not need to double up on processes before due to the symmetry of the input and output images. To be completely thorough in our process definitions for two-dimensional inputs/outputs, processes should consist of left-hand, right-hand, top-down and bottom-up versions. For some processes this will not require four separately defined 3-processes. In three dimensions, there would need to be six separate specifications, again not necessarily all unique.

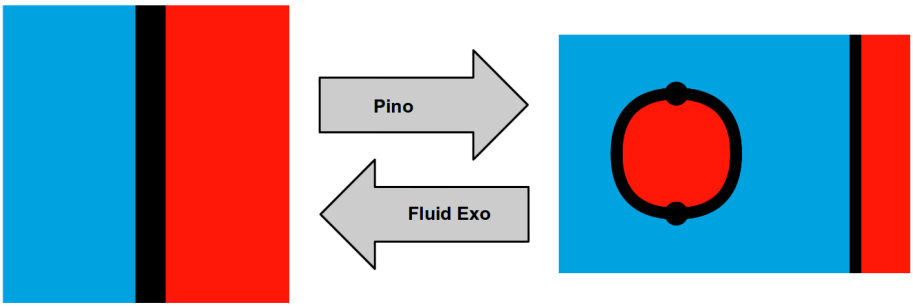

Figure 2.15: Right Side Pino/Fluid Exocytosis Representation as a 3-process in `globular`

Similar to pinocytosis is the process endocytosis ("endo"). Its reverse process is exocytosis ("exo"); in general the term exocytosis on its own refers to the reverse of endocytosis, rather than pinocytosis. While we only show one representation of

endo/exo (Figure 2.16), there are again two versions as with pino/fluid exo. Note that for simplicity we represent what could be many particles as a single particle.

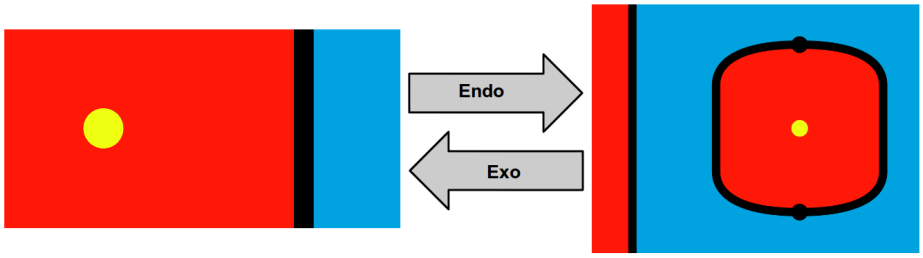

Figure 2.16: Representation of One Form of Endo/Exo as a 3-process in `globular`

One other variant for the input of materials into a cell is phagocytosis (Figure 2.17) or 'cell eating' ("phago"); it does not have an exact reverse like the other processes do.

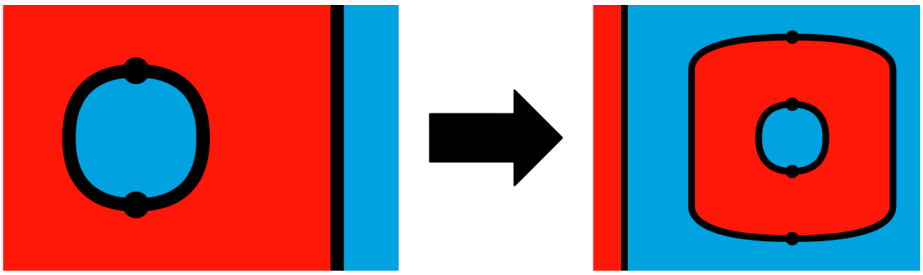

Figure 2.17: Representation of Phago as a 3-process in `globular`

The "drip" process (Figure 2.18) also has no reverse process.

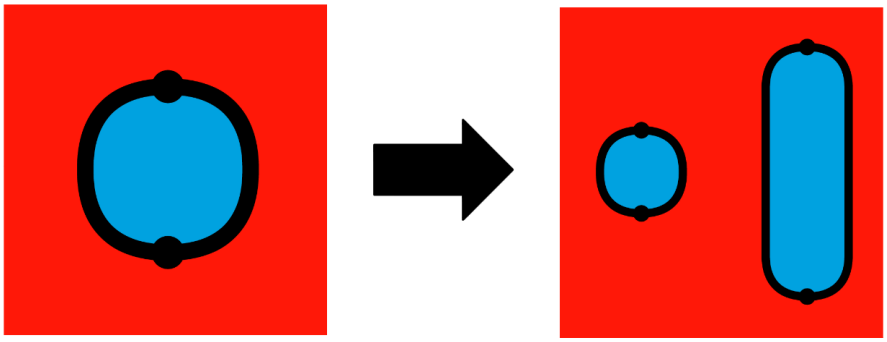

Figure 2.18: Representation of Drip as a 3-process in `globular`

Finally, we have the “bud” process (Figure 2.19), which has no direct reverse process.

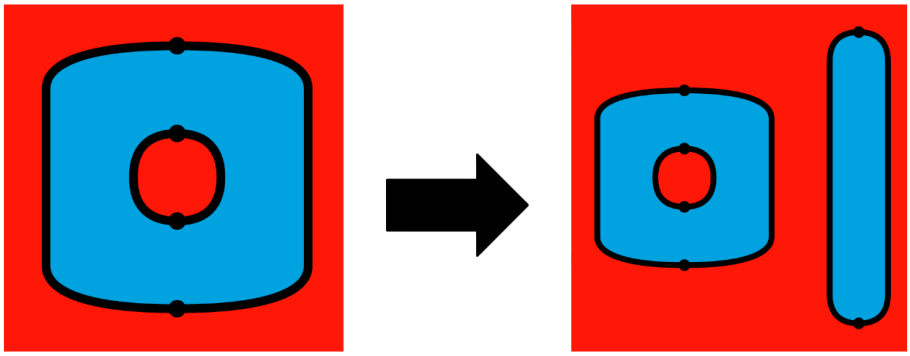

Figure 2.19: Representation of Bud as a 3-process in `globular`

2.2.1 Virus Model

Cardelli uses these basic membrane operations to model the invasion of a cell by the Semliki Forest virus. A diagram of the viral invasion and reproduction process can be seen in Figure 2.20.

In order to create a similar, but somewhat coarser, virus picture, we need to create a few more processes: fizz and froth of extracellular fluid (Figure 2.21), fusion and splitting of membrane-bound extracellular fluid (Figure 2.22), and membrane buckling and straightening from extracellular fluid to the cytosol (Figure 2.23) and vice versa (Figure 2.24).

Figure 2.25 shows the virus invasion process. The Semliki Forest virus first enters the cell through phagocytosis. The virus invasion process is an instance of membrane splitting (Figure 2.22): phagocytosis is technically the sequence of steps of membrane splitting followed by the process we call phago (Figure 2.17).

We note again that all processes should include variants to account for all the directions in which the process might occur. In Figure 2.26, we explicitly show how an interchange process could occur.

In Figure 2.27, we a version of membrane straightening adapted for the virus situation.

Three other processes specially made for this model are virus replication (Figure 2.28), virus migration (Figure 2.29), and virus exit (Figure 2.30).

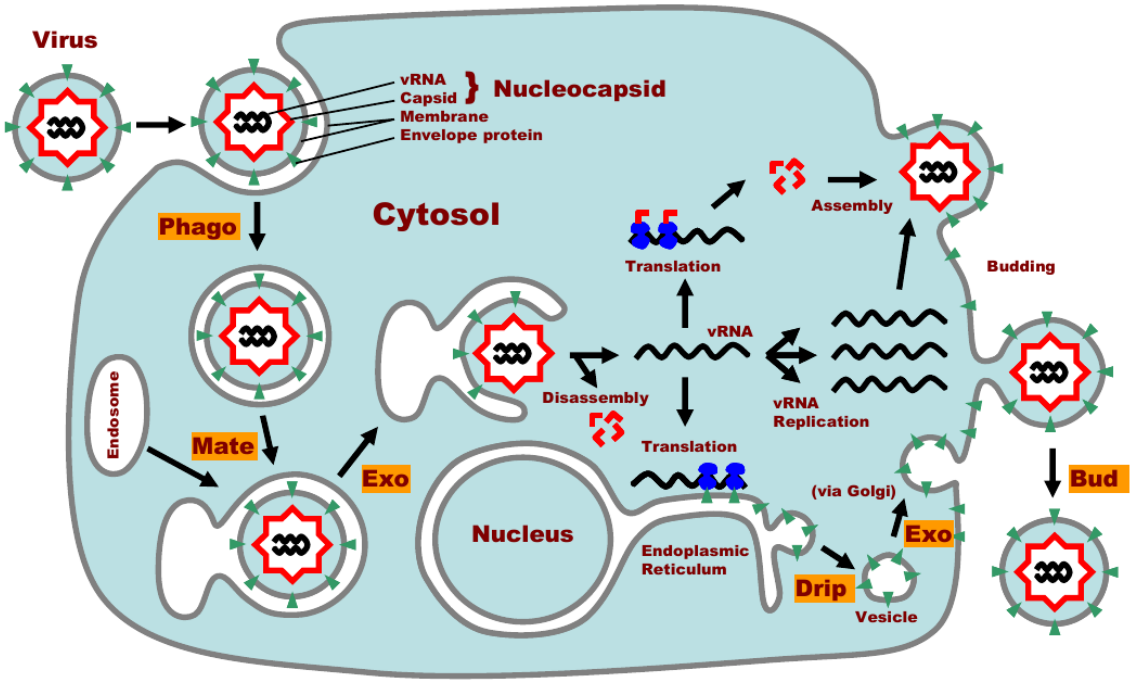

Figure 2.20: Viral Infection and Reproduction [8, page 279]

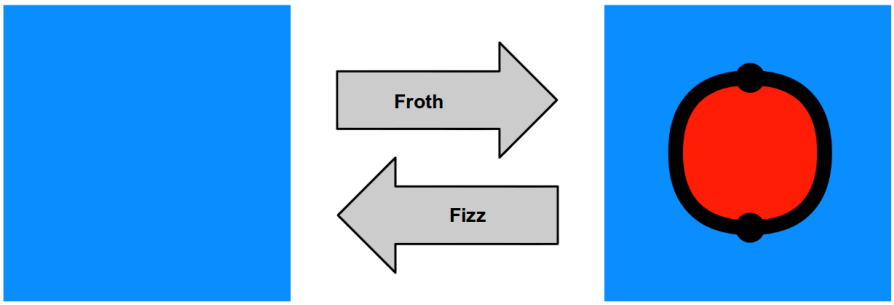

Figure 2.21: Fizz and Froth of Extracellular Fluid as a 3-process in `globular`

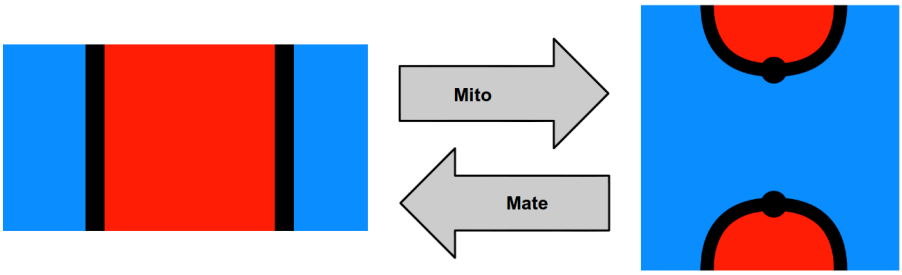

Figure 2.22: Fusion and Splitting of Membrane-Bound Extracellular Fluid as a 3-process in `globular`

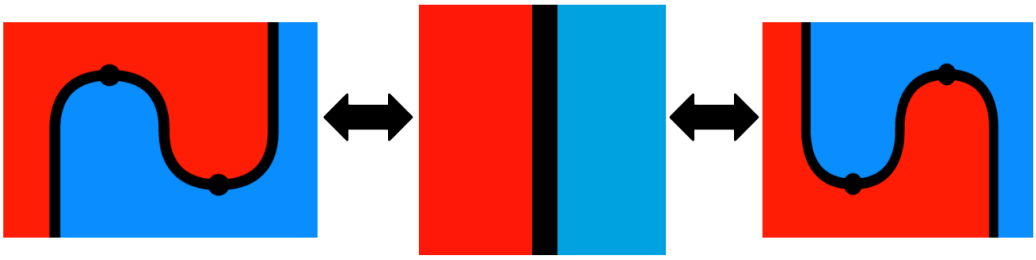

Figure 2.23: Membrane Buckling and Straightening from Extracellular Fluid to Cytosol as a 3-process in `globular`

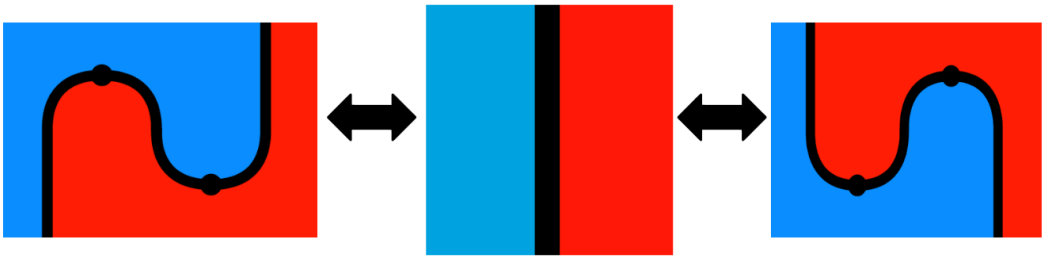

Figure 2.24: Membrane Buckling and Straightening from Cytosol to Extracellular Fluid as a 3-process in `globular`

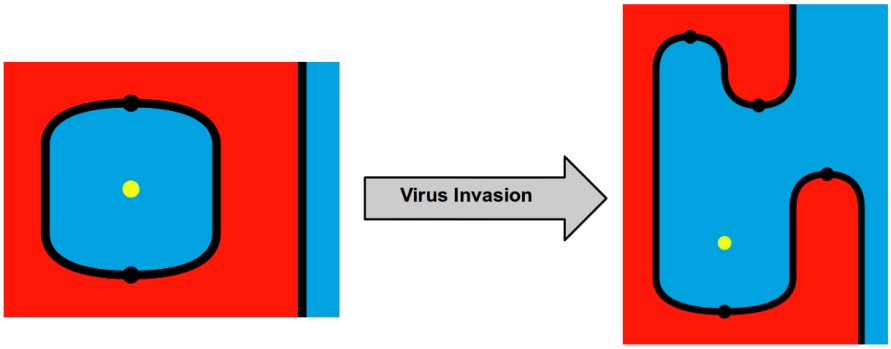

Figure 2.25: Virus Invasion as a 3-process in `globular`

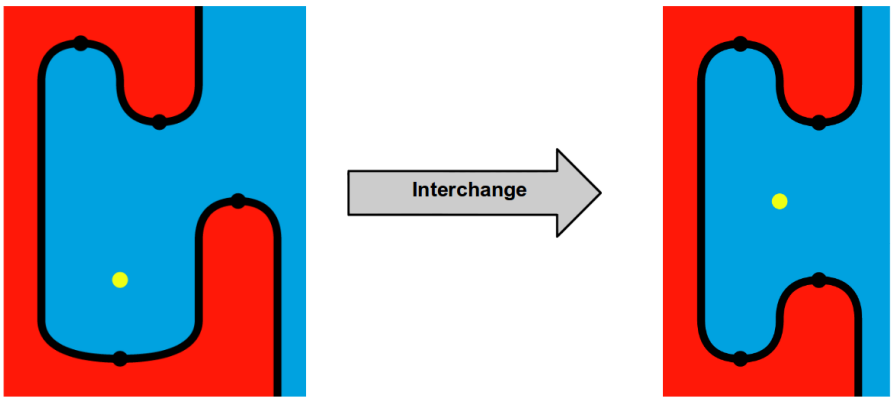

Figure 2.26: Interchange as a 3-process in `globular`

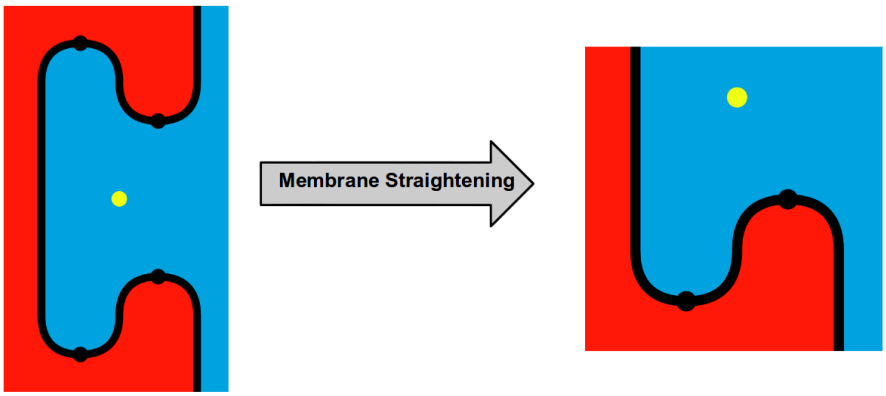

Figure 2.27: Membrane Straightening as a 3-process in `globular`

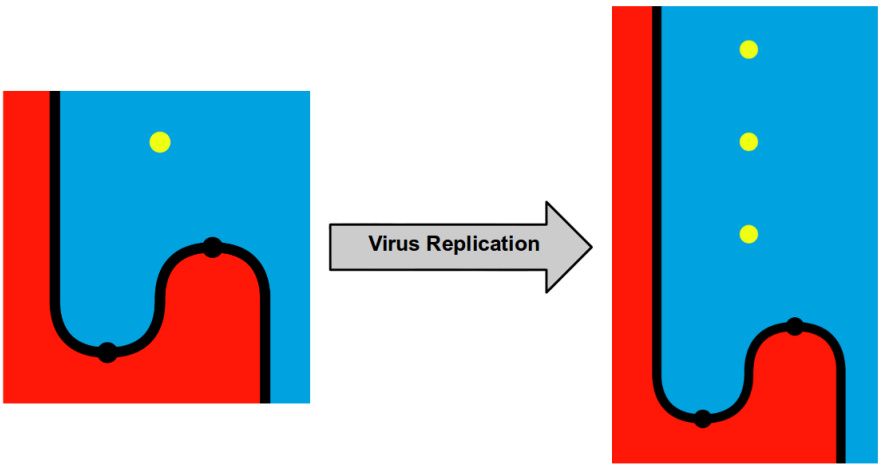

Figure 2.28: Virus Replication as a 3-process in `globular`

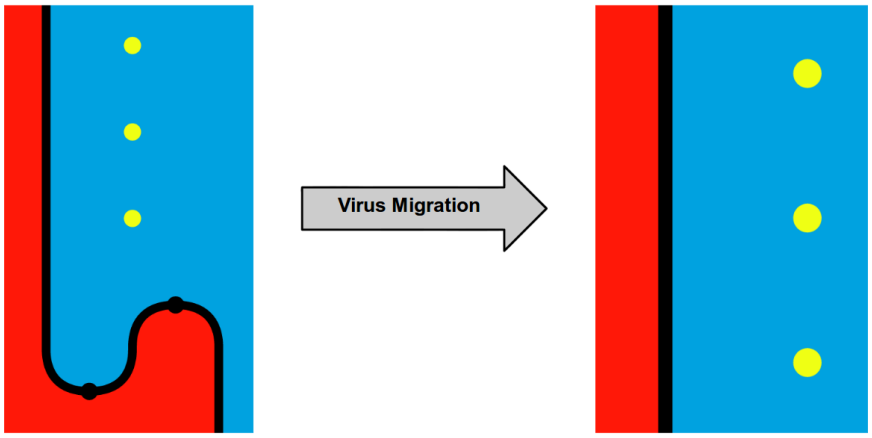

Figure 2.29: Virus Migration as a 3-process in `globular`

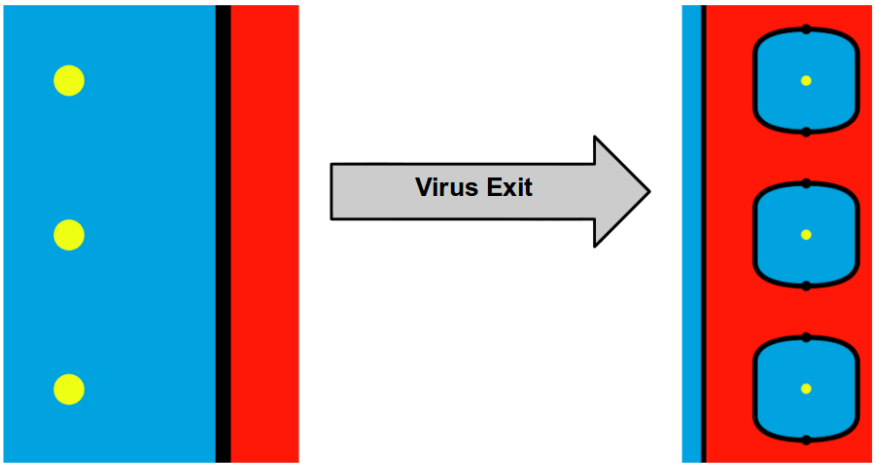

Figure 2.30: Virus Exit as a 3-process in `globular`

Figure 2.31 gives an example of the output from a stochastic simulation with the above 3-processes.

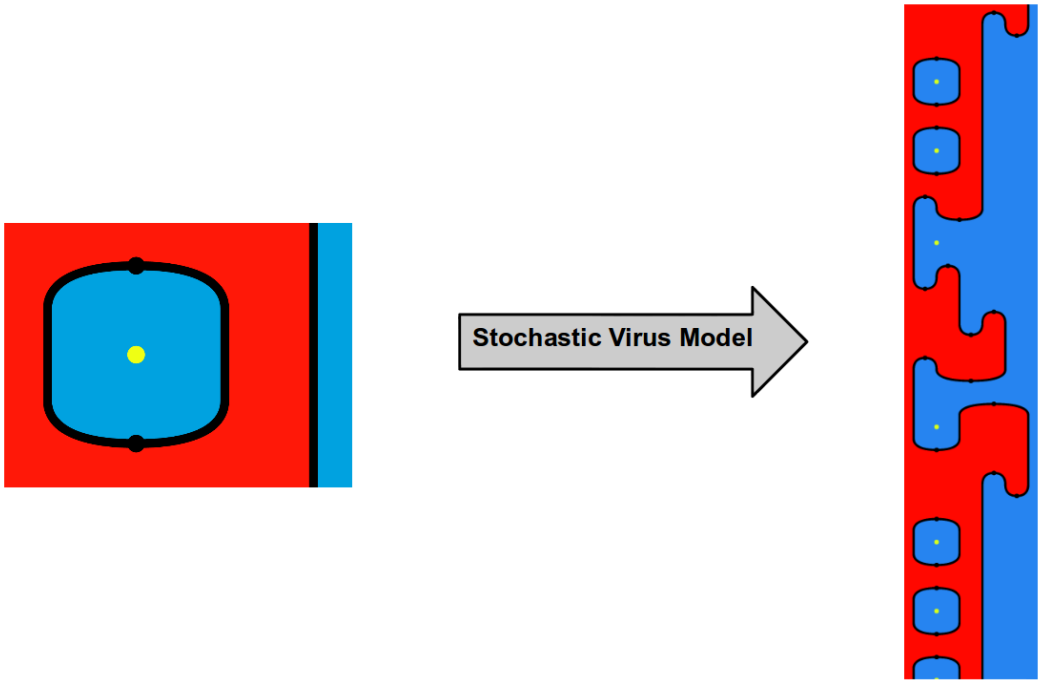

Figure 2.31: Input and Output of a Stochastic Virus Model in `globular`

2.3 Rate Conversion for Stochastic Modelling

It is important to note that the rate constants used in stochastic simulations should be for numbers of molecules, rather than concentration rates [40]. For zeroth-order reactions, that is reactions with no input, the stochastic rate constant c is defined as:

$$c = n_A V k$$

where n_A is Avogadro's number: $n_A \approx 6.023 \times 10^{23}$, V is the volume of the system (for example, the volume of a cell), and k is the rate constant obtained from species concentration data. For first-order reactions (1 reactant), $c = k$. For second-order reactions (2 reactants) in which the reactants are of different species type,

$$c = \frac{k}{n_A V}$$

While for second-order reactions with reactants of the same type,

$$c = \frac{2k}{n_A V}$$

These conversions are based on deterministic rate laws rather than stochastic chemical kinetics, so they will still be valid even if we use a different underlying chemical kinetics framework, as we do in the next chapter.

2.4 Algorithm Comparisons

2.4.1 Steps of the Gillespie Algorithm

For the sake of a full comparison, here we provide the Gillespie Algorithm written in terms of the hazard functions (reproduced from [40, page 183]), where a hazard function is the overall rate at which a particular process occurs:

(i) Initialise the system at $t = 0$ with rate constants $c_1, c_2, \ldots, c_v$ and initial numbers of molecules for each species, $x_1, x_2, \ldots, x_u$.

(ii) For each $i = 1, 2, \ldots, v$, calculate $h_i(x, c_i)$ based on the current state, x.

(iii) Calculate $h_0(x, c) \equiv \sum_{i=1}^{v} h_i(x, c_i)$, the combined reaction hazard.

(iv) Simulate time to next event, t', as an $Exp(h_0(x, c))$ random quantity.

(v) Put $t := t + t'$.

(vi) Simulate the reaction index, j, as a discrete random quantity with probabilities $h_i(x, c_i)/h_0(x, c), i = 1, 2, \ldots, v$.

(vii) Update x according to reaction j. That is, put $x := x + S^{(j)}$, where $S^{(j)}$ denotes the jth column of the stoichiometry matrix S. [The stoichiometry matrix consists of the net change in species numbers after the reactions have occurred.]

(viii) Output x and t.

(ix) If $t < T_{\text{max}}$, return to step 2.

The well-mixed assumption was originally made when spatial data was not readily available and for the sake of keeping the computations less demanding. We know that this is not realistic, but models can provide insight without being fully realistic in their assumptions, and given the widespread use of Gillespie's algorithm, clearly this is such a case.

2.4.2 Gillespie and Stochastic Process Algorithms Comparisons

There are two key differences between our stochastic process algorithm and Gillespie's algorithm (and its variants). The first is that our algorithm recognizes and keeps track of spatial data. The second is that while both algorithms employ exponential distributions to supply their stochastic values, Gillespie's algorithm simulates all reactions in the system until the random time value is reached, while our algorithm assigns time values to all reactions and chooses to execute a single event at each step of the algorithm. While some variants of Gillespie's algorithm, such as the τ leaping method decrease the number of reaction events simulated by "leaping" from one time value to the next, this is still computationally demanding [12].

A more in depth look at the physical implications for this difference in whether time or reaction events act as the stochastic events in a simulation should be explored.

While for Gillespie's algorithm the simulation of reaction events is the demanding step, for the stochastic process algorithm, the most demanding step is the finding of potential matches for the various processes. This step scales according to the dimension of the diagram, with one-dimensional diagrams having a linear run-time; two-dimensional diagrams, quadratic, and n-dimensional diagrams having an n^n running time. All of the other basic methods in `globular` are linear in diagram size. For example, finding the potential interchanger processes in a 2-diagram consists

of going through all the height values in the diagram and finding the unique dot at each height.

2.4.3 Stochastic Process Algorithm with Time Conservation

To be completely rigorous, we will need to find an equivalence between the entire stochastic process algorithm and the entire stochastic with time conservation algorithm. Here we sketch the underlying hypothesis for why these two algorithms would be equivalent.

The way we are assigning times and conserving times which are not used to execute the next event can be shown to be equivalent, within a certain approximation, to assigning an entirely new set of times at each step.

For each iteration of the algorithm, we assign a set of times, which is then ordered into a sequence in which the lowest element is picked as the event to execute. In the case of two events having the same time, only one of these events will be executed: the one whose time corresponds to the first element in the position-event-time list.

Suppose that we have m events at a certain iteration n of the algorithm. Let L be the list of tuples $(position, event, time)$ for this iteration. For the time conservation approach, we note that after the earliest time event occurs, we will be left with $m-1$ tuples, but not all of these tuples will necessarily be relevant for the new state at the $n+1$th iteration. Let L_{-1} denote the set of tuples excluding the tuple of the event which will be executed at the nth step. Let L_r denote the set of tuples of events which are still relevant at the $n+1$th iteration. Let α denote the triple which contains the lowest time value in L_r (which may even be equal to the executed event's time value)–we choose this event for no specific reason other than that it can always be identified, so long as L_r is not an empty list, in which case the issue under consideration would not occur in the first place.

For the new time assignments approach, let L' be the set of tuples for the $n+1$th iteration, and let α' be the triple whose event and position values are equal to α's.

Now, because α and α' have the same event, both of their time values will be chosen from the same probability distribution, so the mean value for both α and α' is $\frac{1}{r}$, with r being the rate for the process of which the event is an instance. That is, on average the two values are the same: we are sampling from the same distribution both times, thus we are going to get similar values, so whether we use a value picked earlier rather than later is of no consequence.

Given that the efficiency of not sampling distributions a second time comes at a memory usage cost, it is not yet clear whether the time conservation algorithm is more efficient overall.

Chapter 3

Chemical Kinetics and Other Model Comparisons

We first present the chemical kinetics most widely used to model biochemical processes. Through a series of illustrative examples, we show how this approach varies in predictions from what we call "the spatial lattice model." We define this model explicitly in terms of appropriate "hazard functions." After presenting this approach, we present an interpretation of how in the limit of fast interchanging, that is, the fast permuting of positions, the well-mixed approach is equivalent to what we call "the well-mixed-compatible model." While our two spatial models do not agree with each other, because they both rest on valid assumptions, they each promise to offer different useful insights.

3.1 The Well-Mixed Chemical Kinetics Model

Gillespie's stochastic kinetics model has become the "theoretical ground for most of the stochastic modelling research conducted today in the systems biology domain" [27, page 123]. The Gillespie algorithm is the standard stochastic simulation procedure used to model the evolution of a biochemical system. It is theoretically grounded by the Chemical Master Equation (CME), a chemical kinetics model with rates, usually referred to as "hazard functions," whose definitions (from [40, pages 180-181] and included below) are based on combinatorics arguments. Most other simulation methods are based on the same stochastic kinetics theory on which Gillespie's method is based [40, page 221].

3.1.1 Assumptions of the CME approach

We state the assumptions of the standard chemical kinetics model, as relayed in [27]:

Assumption 3.1.1. *The chemical system is under thermal equilibrium conditions.*

Assumption 3.1.2. *The chemical system is such that, at any time t, the concentration of each species is homogeneous in the reaction vessel (i.e., does not depend on space).*

Assumption 3.1.3. *In a bimolecular reaction, the time to the occurrence of the reaction is largely determined by the time to the reactive collision, whereas the time necessary for the chemical transformation of the colliding species into the reaction products is negligible.*

3.1.2 Hazard Functions for the CME Approach

An n^{th}-order reaction is a reaction whose input consists of n reactants. There are three types of reactions which are relevant for us to consider: zeroth-, first- and second-order reactions. We will make the standard assumption that a reaction with three input species will only occur as a sequence of first-order and second-order reactions [40, page 182],[27, page 125]. Zeroth-order reactions do not occur in nature, but they can be useful for models in which there is an influx of some species. We will often refer to zeroth-order reactions as "creation/generation processes," as in Figure 2.13.

The Chemical Master Equation (CME) describes how many ways a particular process can happen, given the process rate and the numbers of input species. The basic idea is that if given a process and m ways to start it, then there are m ways for the process to happen. If the rate constant of the process is r, then since there are m ways for the process to happen, it is m times as likely to happen as it was with one way, so the total reaction rate for the process is $m \cdot r$. This CME approach is not intended to be spatial in nature, and so it does not take dimensional data into account.

For zeroth-order reactions

$$\mathcal{R}_i : \emptyset \xrightarrow{c_i} X$$

where c_i is the reaction rate, the hazard function is defined to be

$$h_i(x, c_i) = c_i$$

This is because there are no input species to increase the likelihood of the process occurring. Later, it will be seen in the `globular` simulations that a "0-cell" corresponds to the $\emptyset$ of a zeroth-order reaction. Let x_i denote the number of instances of species X_i. Then for first-order reactions

$$\mathcal{R}_i : X_j \xrightarrow{c_i} ?$$

the hazard function is:

$$h_i(x, c_i) = c_i x_j$$

For second-order reactions

$$\mathcal{R}_i : X_j + X_k \xrightarrow{c_i} ?$$

the hazard function is:

$$h_i(x, c_i) = c_i x_j x_k$$

For the special case of dimerisation, that is when two molecules of the same type combine to form one molecule, called a dimer, we have:

$$\mathcal{R}_i : 2X_j \xrightarrow{c_i} ?$$

the hazard function is:

$$h_i(x, c_i) = c_i \frac{x_j(x_j - 1)}{2}$$

3.1.3 Spatial Limitations of the CME Approach

For zeroth-order and first-order reactions (also referred to as "equations"), the hazard functions are also valid in a spatial model. However, second-order equations necessarily rely on the distance between reactants, and this fact causes a discrepancy between the CME approach and our spatial approach. In subsection 3.1.4, we define a new hazard function for second-order reactions in our spatial model. In this section however, we motivate the creation of this new hazard function by looking at several cases and discussing the associated physical interpretations. We will take all rate constants to have units of s^{-1}.

Example 3.1.1. *Consider a system with the following two reactions:*

$$\alpha : A \to Z \text{ and } \beta : AB \to W$$

Reactions α and β are assigned the same rate constant value, which for simplicity's sake, we will take to be 1. In our system, substance A occurs in abundance compared to B:

for the sake of illustration, let us assume that there are 1000 times as many molecules of A as there are of molecules of B, and let us take the number of molecules of B to be 1. From subsection 3.1.2, we know that the overall rate for reaction α is:

$$h_\alpha(x, c_\alpha) = c_\alpha x_A,$$

where c_α is the rate constant and x_A is the number of molecules of A. And similarly, the overall rate for reaction β is:

$$h_\beta(x, c_\beta) = c_\beta x_A x_B,$$

where c_β is the rate constant, x_A is as before, and x_B is the number of molecules of B. For reaction α, $h_\alpha(x, 1) = 1 * 1000$ or 1000/s, but note that reaction β's hazard function, $h_\beta(x, 1) = 1 * 1000 * 1$ or 1000/s has the same value of 1000/s.

Now consider the case where spatial data does matter. Following our one-dimensional procedure, let us take the input for our system to be a string of letters, where each instance of A is called A and similarly for B. We expect a string of length 1001, where a certain position b in the string is the location of B. For example "$A - \ldots - A - A - A - B - A - A - A - A - A - A - A - \ldots - A$" is such an input; we place dashes between the letters to emphasise that being adjacent to a letter does not imply that a reaction will necessarily happen. For this input, there are two ways for reaction β to occur since 'B' has an 'A' on either side of it ('AB' or 'BA' can act as input for β), and we assume that the reaction is indifferent to the order of the letters. Thus the overall rate with which reaction β will occur is $1 * 2 = 2$ or 1 in the case that B is on either end of the string, and the overall rate with which reaction α will occur is $1 * 1000$ or 1000/s. As we claimed, the overall likelihood with which the first-order equation (α) occurs is equal in both the spatial and well-mixed models. For the second-order reaction (β), however, we see that the well-mixed case predicts a likelihood which is at least 500 times greater than the spatial case's.

Even if we tried to induce well-mixed conditions in the spatial model by swapping positions of adjacent letters with some separate rate, the one-dimensional constraint, coupled with the adjacent positions requirement, forces the spatial model form of reaction β to always have maximum overall rate of 2/s.

The one-dimensional case is an oversimplification of a biological system. We can see that in this case, the concept of being well-mixed cannot be properly portrayed because the one-dimensional constraint restricts all possible interactions to only be adjacent interactions. Therefore in this case, the well-mixed model does not make physical sense, but this is in part expected because a string of letters can never have the properties expected of a well-mixed higher dimensional system.

Example 3.1.2. *Consider the same reactions and initial conditions as above where we are now in a two-dimensional setting. We ignore the first-order reaction since we know it will always agree with the spatial model. Since the well-mixed model does not take dimensions into account, its overall rates are the same as above. For the spatial model, now our input is an array, or lattice, of letters. We shall consider valid adjacencies to be those directly above, below, to the right or to the left of B. There are three different configurations of A's which could be around B:*

(i) If B is in a corner of the lattice, then it will have two A's which are close enough to it to count as neighbours.

(ii) If B is on an edge of the lattice but not a corner, then it will have three adjacent A's.

(iii) If B is in the interior of the lattice, it will have four adjacent A's, the maximum number of neighbours.

Despite the different configurations, clearly in the majority of cases, B will have four neighbours. In this case, the likelihood of reaction β is $1 * 4$ or 4/s. This is still significantly smaller than the well-mixed rate of 1000/s and cannot be resolved by interchanging letters quickly.

Example 3.1.3. *Now take the same system in three dimensions. Again, the well-mixed likelihoods are the same. For the spatial model input we will have a three-dimensional lattice, so in the majority of cases, B will be surrounded by 6 neighbours, resulting in a rate of 6/s, again still significantly lower than the well-mixed rate.*

Considering the physical implications of these two models, the well-mixed model is stating that the rarity of a reactant does not affect the overall rate of a reaction, whereas the spatial model states that it does.

Given the many transport mechanisms a cell has in place for delivering special molecules to the right parts of the cell, the spatial model intuition appears to correlate with biological systems better than the well-mixed model does.

Example 3.1.4. *We provide one final example of a more complex system. Here there are two reactions $\gamma : AA \to Y$, and $\beta : AB \to W$, as before. Reactions γ and β have the same rate constant, which we will again set to be 1. There are 1000 molecules of A and 2 molecules of B. Note that γ is a dimerisation type reaction. Since the three-dimensional setting produced the closest comparison with the last system, let us stay in three dimensions. The well-mixed model will predict a rate of $1 * 1000 * (1000 - 1)/2$ or*

roughly 500,000 molecules of Y per second for reaction γ, and for reaction β, it will predict a rate of 2000 molecules of W per second. In the spatial model, the rate of β will be the sum of each rate for the B molecules. Since it is most likely that both B molecules will have 6 neighbours, let us take each of their rates to be 6/s, so overall, the rate of β is 12/s. The reaction of γ will be equal to the sum of all the possible adjacent A, A pairs; this will be roughly equal to the difference between the number of possible adjacent pairs and the number of A, B pairs. The three-dimensional lattice we are looking at is roughly a cube with side length 10. The four corners of this lattice will each have 3 adjacent points, the other edge points will have 4 neighbours, and all interior points will have 6 neighbours. The total number of distinct adjacent pairs in a three-dimensional square lattice of side length 10 is:

(number of adjacent pairs per horizontal row) · (number of horizontal rows)

+ (number of adjacent pairs per vertical row) · (number of vertical rows)

+ (number of points in a layer of the lattice cube)· (number of layers in the lattice cube)

$$= 9 \cdot 10 + 9 \cdot 10 + 100 \cdot 9 = 1080$$

If we assume that the two instances of B will be far apart but both in the interior of the cube, then we will need to subtract $6 * 2 = 12$ pairs from the 1080 total adjacent pairs. Therefore the overall reaction rate for γ is roughly 1070 molecules of Y/s, which is still several orders of magnitude less than the rate for the well-mixed model. It does not matter how many interchanges we perform, because when an event happens you will necessarily be in a single configuration of the system, which means that the spatial model will hold. The well-mixed model will always over-count the number of possible reactions.

Unsurprisingly, if the well-mixed model over-estimates the reaction likelihoods, then a model like this spatial lattice model on the opposite extreme would be likely to under-estimate the reaction likelihoods, as seen above.

3.1.4 The Spatial Lattice Model for Stochastic Kinetics

For a reaction with a single input, the position of a_i, a molecule of type A, was irrelevant to the hazard function for the associated reaction. However, when two reactants A and B are involved in a reaction β, the spatial relation between instances a_i and b_i affects the likelihood of the reaction. For example, if a_i and b_i are close

to each other, then the likelihood of a collision occurring in the future seems much higher than the likelihood if they were far apart. In the first case, this pair would contribute positively to β's hazard function, while in the second case, the pair should not be counted with the same weight as the close-to-each-other pair, and if the distance is large enough, the pair should perhaps not be counted at all, as the odds of reaction β happening as a result of this pair colliding are so low that it can be assumed that the reaction never happens.

While one could: calculate the distance between every pair of molecules, establish a threshold value for when pairs could interact, and then for each state derive a new set of overall reaction rates for each reaction, it is simplest to first follow a model at the other extreme from the well-mixed model, that is to count the contributions of only the adjacent molecules to a particular likelihood. This spatial model will not be able to account for molecular density when such density is greater than that of a square lattice model.

Zero-order and first-order reactions will have the same form as in the well-mixed model since spatial and dimensional considerations are not relevant in these cases.

The general form for the hazard function for second-order reactions with different reactants is then:

$$h_i(x, c_i) = \sum_{i=1}^{x_r} c_i \cdot 1 \cdot x_{adj}$$

where x_r is the number of molecules of the rarer substance, and x_{adj} is the number of adjacent molecules of the second reactant type, which will vary with each instance of the rarer substance. This adjacency parameter will also be influenced by the dimension of the model, with the maximum adjacency level for one-dimensional cases being 2, that for two-dimensional cases being 4, and that for three-dimensional cases being 6.

The general form for the hazard function for second-order dimerisation reactions is:

$$h_i(x, c_i) = \sum_{i=1}^{x_r} c_i \cdot 1 \cdot x_{adj}$$

where x_r is the number of molecules of substance X_R, and x_{adj} is the number of molecules of type r which are above or to the right of this lattice point. The directionality with which we check for adjacent pairs will prevent over-counting of dimerisation pairs.

3.2 The Well-Mixed-Compatible Model

In this approach, rather than considering a lattice structure to account for spatial data, we instead focus on the following observation: in the well-mixed model, reactants have to first find each other, and then the reaction has a certain likelihood of occurring, whereas in the spatial setup, the reactants have already found each other: whether the reaction occurs or not is entirely dependent on the relevant rate constant for the reaction.

We need to be careful when discussing the various rate terms: a rate constant is an experimentally derived quantity which gives the rate at which product molecules will be produced when the proper reactants are available. The overall rate or likelihood or "hazard" of a reaction is the total rate at which a certain type of reaction will occur; this quantity depends on the number of possible instances of the reaction which can occur and also on the rate constant for the reaction.

Since the rate constant value is likely to come from a well-mixed type experimental setup, the rate constant is probably not accurate for a spatial model. We would expect the rate constant to be lower for the well-mixed model due to the delays caused by one reactant having to first encounter the other reactant before the reaction can potentially take place. The guiding principle is thus to scale the hazard function appropriately in the fast swapping situation in order to obtain the same values that the CME approach gives.

Again because spatial considerations do not affect zero-order and first-order reactions, we will only compare second-order reaction predictions. We take the same examples used before (reproduced below) and note the differences with this model.

Example 3.2.1. *(Compare with example 3.1.1.) Consider a system with the following two reactions:*

$$\alpha : A \to Z \text{ and } \beta : AB \to W$$

Reactions α and β are assigned the same rate constant value, which for simplicity's sake, we will take to be 1. In our system, substance A occurs in abundance compared to B: for the sake of illustration, let us assume that there are 1000 times as many molecules of A as there are of molecules of B, and let us take the number of molecules of B to be 1. From subsection 3.1.2, we know that the overall rate for reaction α is:

$$h_\alpha(x, c_\alpha) = c_\alpha x_A,$$

where c_α is the rate constant and x_A is the number of molecules of A. And similarly, the overall rate for reaction β is:

$$h_\beta(x, c_\beta) = c_\beta x_A x_B,$$

where c_β is the rate constant, x_A is as before, and x_B is the number of molecules of B. For reaction α, $h_\alpha(x, 1) = 1 * 1000$ or 1000/s, but note that reaction β's hazard, $h_\beta(x, 1) = 1 * 1000 * 1$ or 1000/s has the same value of 1000/s. However, consider the case where spatial data does matter. Following our one-dimensional procedure, let us take the input for our system to be a string of letters, where each instance of A is called A and similarly for B. We expect a string of length 1001, where a certain position b in the string is the location of B. For example "$A-\ldots-A-A-A-B-A-A-A-A-A-A-A-\ldots-A$" is such an input; we place dashes between the letters to emphasise that being adjacent to a letter does not imply that a reaction will necessarily happen.

*The second-order reaction likelihood will be 1*2 or 2/s since there are only at maximum two correct input instances for β. Having a much lower value for the spatial model agrees with our intuition that the rate constant for the spatial model should be higher than the rate constant for the well-mixed model: if the rate constant for the spatial case could be calculated, it would be a larger number than 1, which could help to decrease the discrepancy between the well-mixed and spatial likelihoods.*

*If however, fast swapping was to occur in the spatial situation, effectively simulating the well-mixed conditions, then the rate constant would be appropriate. If we interpret the fast swapping to allow for an adjacent pair of A, B at each location in the string as a possible reaction event, then since there are 1000 such placements, the overall reaction rate for β is 1*1000, matching the well-mixed prediction.*

Example 3.2.2. *(Compare with example 3.1.2.) Consider the same reactions and initial conditions as above where we are now in a two-dimensional setting. Since the well-mixed model does not take dimensions into account, its overall rates are the same as above. For the spatial model, now our input is an array or lattice of letters. We focus on the well-mixed case, since the non-well-mixed case will be as with the spatial lattice model. Since we have roughly 1000 molecules, let us think of the lattice as being a 10 x 100 grid. If we count up all the possible arrangements in which A, B could be an adjacent pair, for the first row, our answer is 99. Since there are 10 rows, we know that at least $99 \times 10 = 990$ possible arrangements are available. These calculations have excluded the vertical pairings, so we multiple $9 \times 100 = 900$. Overall, we thus have $990 + 900$, or roughly 2000 possible pairings. Now the overall rate for the spatial model*

*will be 1*2000, which is double that of the well-mixed model. However, if we add to the definition of the overall rate a constant factor of 1 over the dimension of the system, then we recover the well-mixed rate of 1000/s.*

Example 3.2.3. *(Compare with example 3.1.3.) Now take the same system in three dimensions. Again, the well-mixed likelihoods are the same. For the spatial model input in the well-mixed case, we will again determine the rate by counting the possible pairings. This time we visualize a lattice cube of side length 10. By the same method as used in the two dimensional case, we know that each layer of the lattice will have* $90 + 90 = 180$ *pairings, and that since we have 10 such layers, there will be at least 1800 pairings. We must consider both the horizontal and vertical layers however, so this leads to* $2 \times 1800 = 3600$ *pairings. Again, if we merely plug this number into the rate constant times number of pairings formula, we will end up with a much larger rate than predicted by the well-mixed model. However, if we again try dividing by the dimension of the system, we obtain 1200/s as our rate, which is fairly close to recovering the 1000/s value.*

Example 3.2.4. *(Compare with example 3.1.4.) We provide one final example of a more complex system. Here there are two reactions* $\gamma : AA \to Y$ *and* $\beta : AB \to W$ *from above,* γ *and* β *have the same rate constant, which we will again set to be 1. There are 1000 molecules of* A *and 2 molecules of* B*. Note that* γ *is a dimerisation type reaction. Let us stay in three dimensions. The well-mixed model will predict a rate of roughly 500,000 molecules of* Y *per second for reaction* γ*, and for reaction* β*, it will predict a rate of 2000 molecules of* W *per second. In the spatial model, we still expect to obtain low values in the non-well-mixed case due to the rate constant being lower problem. In the fast swapping limit, for* β *we have roughly the same 1000 molecule lattice cube as in the previous example, so there are roughly 3600 pairings, per* B *molecule, and thus 7200 pairings, which when divided by the dimension, yields a rate of roughly 2000* W *molecules per second as desired. For* γ*, since we are counting the number of pairings of 1000 molecules of* A*, we will end up with exactly the same result as the well-mixed model due to the* $n(n-1)/2$ *hazard function formula.*

Based on the conclusions above, since the dimerisation case was equivalent, we define one new hazard function for this model when compared with the CME model. For second-order reactions with different reactants, we have:

$$h_i(x, c_i) = \frac{c_i}{D} \cdot n_{adj}$$

where D is the dimension of the system, and n_{adj} is the number of possible adjacent pairings.

3.2.1 Flexibility of the Well-Mixed-Compatible Model

Since diffusion dominates in the well-mixed case, if a reaction event happens, then it is unlikely that the same event can occur in the near future because diffusion will occur in between events. This provides a reasonable model of, for example, an enzyme reaction in which the enzyme requires ATP. After the reaction has taken place, the enzyme will be left with ADP because the phosphate group was released as energy in the reaction. This means that the enzyme will be inactive until it is replenished with the missing phosphate group which makes it active. Because the enzyme is not immediately ready for a second reaction, there is a necessary delay time, just as diffusion produces a necessary delay time between reaction events.

Not every biochemical reaction has the ATP setup, however. Our modelling setup (the non-interchanging version) allows for reaction types which do not require a delay time in between. Diffusion also seems less likely to dominate in the way it does with gas particles in a vacuum (Gillespie model's physical basis [19, page 407]) when you have such particles, or even more complex molecules, in a soup of other molecules (cytosol).

If delays between reactions were desirable, then the rate at which interchanges occur could be increased so that more mixing/diffusion events occur before reaction events. Thus our model has greater flexibility in modelling various biochemical situations than the CME approach.

3.2.2 Well-Mixing in Globular

We simulate a well-mixed situation by applying a series of interchange moves or swaps. These are localized permutations.

The default setup in `globular` is that once interchanges are possible, that is when 3-processes are available, a set of the various dots which can be interchanged is identified. The process list is updated to include each possible interchange event, and the time distribution for each interchange event is based on the change in height, interpreted here as "tension." Specifically, the rate for an interchange event is $e^{-\Delta\tau/T}$, where $\Delta\tau$ is the height difference between the two dots, and T is the temperature of the system. After the rate has been calculated, it is passed into the negative exponential distribution, from which, as with all processes, times are sampled. Figure 3.1 shows an example of an interchange event.

Because it is often the case that 3-process inputs include at least one dot, the number of possible interchanges in a given diagram composed of 2-cells is likely to

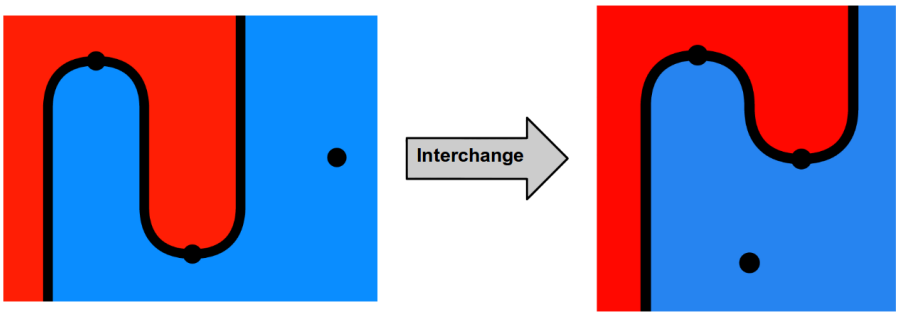

Figure 3.1: An Example of an Interchange Operation: here the height of the cup-shaped part of the membrane was interchanged with the height of the particle in the cytosol so that at the end of the process the particle is lower than the cup-shaped portion of the membrane.

be higher than the number of possible occurrences of inputs for any other process in the diagram. To this end, some sort of normalization factor should be included in the calculation of the rate for interchanges. Another update to the stochastic process algorithm in `globular` for biochemical processes would be dividing the rate of processes whose inputs are 0-cells by the number of (corresponding) 0-cells in the given state. This is because reactions with 0-cells are equivalent to zeroth-order reactions in the sense that their input is just a segment of the "background" for the diagram: a 0-cell would certainly not be referred to as an object in a monoidal 2-category or in any category in which the effective objects are 1-morphisms. If 0-cells are thus not counted among the other species in the model, they can be considered to be equivalent to nothing.

3.2.3 Connection to Deterministic Modelling

As an aside, the CME is equivalent to the rate equation for systems in chemical equilibrium [9, 11]; this provides a link between our stochastic modelling approach and the deterministic approach of the rate equation. (The rate equation model treats species in terms of continuous concentrations rather than discrete species numbers.) Therefore, our well-mixed-compatible spatial kinetics model seems to be the most general kinetics model in the sense that it can describe a range of behaviours, including the CME kinetics in the fast-swapping limit, which in turn describes deterministic kinetics at chemical equilibrium.

Chapter 4

Conclusions

4.1 Summary

From the discussions in these chapters, it should be clear that using higher categories to stochastically model biological processes is still very much in its infancy, but at the same time shows great promise. The graphical logic of higher categories, as implemented in `globular`, can capture the local nature of biological processes. Additionally, the ability to model in different dimensions without changing the underlying framework provides a rare type of robust and desirable flexibility. While it remains to determine the exact chemical kinetics model underlying the stochastic process algorithm, this suggests that there may be a need for a new basic physical model to describe biological processes in spatial detail.

4.2 Proposed Extensions of the Model

Up to this point, we have discussed how our model differs in its underlying stochastic kinetics from the standard approach in terms of spatial considerations (Assumption 3.1.2). However, Assumption 3.1.1 is also not realistic for biological systems. A system is in thermal equilibrium if the temperature within the system is uniform in space and time. It has been noted in [30, page 23] and [34, page 341] that cell death occurs under thermal equilibrium conditions and moreover that "cellular processes in general are far from thermal equilibrium" [30, page 23].

Dropping the thermal equilibrium assumption means that each species in a process would need to be assigned a specific temperature. Given that experimental techniques can now account for spatial distinctions, it does not seem improbable that localized temperature data could be obtained in the near future, if such data has yet to be obtained. Since reaction rate constants are temperature dependent,

including thermal data will affect the hazard function, which will in turn influence the exponential distribution from which event times are sampled. Thermal data could be represented in `globular` as an overlay of temperatures on the input species. We also note that the interchange rate would need to be adjusted to include a change in temperature between the two nodes, rather than a single equilibrium temperature.

4.3 Future Work

4.3.1 Theory

It would be nice to derive a single general hazard function which could account for all four cases for our spatial kinetics. The question of density of molecules and whether the most accurate spatial model would account for more adjacent molecule possibilities should be addressed.

4.3.2 Applications

Other biological systems that seem particularly accessible to modelling given the work that we have already done in `globular` include cell polarization in yeast cell mating [29], since the mechanisms involved are primarily membrane deformations, and finding a steady-state in cell division process dynamics to better understand cancer growth.

Once three-dimensional inputs are implementable in `globular`, torsion reducing operations of topoisomerases on DNA [18] and other topological surgery processes [10], [28] could be modelled.

It would be nice to add a feature to `globular` that immediately generated all the forms a single process, for example, if a vesicle entered a cell from the left, then `globular` would have a one-click feature for adding the cases in which a vesicle entered from the right, and above and below the cell (for the two-dimensional case).

References

[1] 0312_animal_cell_and_components.jpg (JPEG Image, 813 645 pixels) - Scaled (80%). https://upload.wikimedia.org/wikipedia/commons/0/0d/0312_Animal_Cell_and_Components.jpg. 23

[2] BioMed Search - Biomedical Image Search Engine. http://www.biomed-search.com/search?q=mitotically&s=54&r=9. 23

[3] Home - The P Systems Webpage. http://ppage.psystems.eu/. 4

[4] Lung-on-a-chip : Wyss Institute at Harvard. http://wyss.harvard.edu/viewpage/240/. 5

[5] Mathematical Biosciences Institute :: Workshop 2: The Interplay of Stochastic and Deterministic Dynamics in Networks. http://mbi.osu.edu/event/?id=897#description. 2

[6] *The Oxford Handbook of Membrane Computing*. Oxford University Press, Oxford; New York, 1 edition edition, February 2010. 4, 22

[7] Samson Abramsky and Bob Coecke. Physics from computer science: a position statement. *International Journal of Unconventional Computing*, 3(3):179, 2007. 1

[8] Bruce Alberts, Dennis Bray, Julian Lewis, Martin Raff, Keith Roberts, and James D. Watson. *Molecular Biology of the Cell 3E*. Garland Science, New York, 3 edition edition, March 1994. 28

[9] David F. Anderson, Gheorghe Cracium, and Thomas G. Kurtz. [0803.3042] Product-form stationary distributions for deficiency zero chemical reaction networks. 48

[10] S. Antoniou and S. Lambropoulou. Dynamical systems and topological surgery. *arXiv preprint arXiv:0812.2367*, 2008. 50

[11] John C. Baez and Jacob D. Biamonte. Quantum Techniques for Stochastic Mechanics. 7, 9, 48

[12] Harvey Thomas Banks, Anna Broido, Brandi Canter, Kaitlyn Gayvert, Shuhua Hu, Michele Joyner, and Kathryn Link. Simulation Algorithms for Continuous Time Markov Chain Models. *Applied Eletromagnetics and Mechanics*, 37, 2007. 34

[13] Krzysztof Bar, Aleks Kissinger, and Jamie Vicary. Globular. 2015. Presented at Higher Dimensional Rewriting and Applications 2015. 1

[14] Krzysztof Bar and Jamie Vicary. Groupoid Semantics for Thermal Computing. *arXiv preprint arXiv:1401.3280*, 2014. 1

[15] Sangeeta N. Bhatia and Donald E. Ingber. Microfluidic organs-on-chips. *Nature Biotechnology*, 32(8):760–772, August 2014. 5

[16] Luca Cardelli. Brane calculi. In *Computational methods in systems biology*, pages 257–278. Springer, 2005. 4, 22, 23

[17] Harold Fellermann and Luca Cardelli. Programming chemistry in DNA-addressable bioreactors. *Journal of The Royal Society Interface*, 11(99):20130987, October 2014. 4

[18] Maxim D. Frank-Kamenetskii. DNA topology. *Journal of Molecular Structure: THEOCHEM*, 336(23):235–243, June 1995. 50

[19] Daniel T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1):404–425, 1992. 3, 47

[20] Daniel T. Gillespie. Stochastic Simulation of Chemical Kinetics. *Annual Review of Physical Chemistry*, 58(1):35–55, May 2007. 3

[21] Peter JE Goss and Jean Peccoud. Quantitative modeling of stochastic systems in molecular biology by using stochastic Petri nets. *Proceedings of the National Academy of Sciences*, 95(12):6750–6755, 1998. 8

[22] Geoffrey Grimmett and David Stirzaker. *Probability and Random Processes*. Oxford University Press, Oxford; New York, 3 edition edition, 2001. 18

[23] Stefan Hellander, Andreas Hellander, and Linda Petzold. Reaction rates for mesoscopic reaction-diffusion kinetics. *Phys. Rev. E*, 91(2):023312, February 2015. 3

[24] Chris Heunen and Jamie Vicary. Categorical quantum mechanics: An introduction. 2015. 11

[25] Dongeun Huh, Benjamin D. Matthews, Akiko Mammoto, Martn Montoya-Zavala, Hong Yuan Hsin, and Donald E. Ingber. Reconstituting organ-level lung functions on a chip. *Science*, 328(5986):1662–1668, 2010. 5

[26] James Keener and James Sneyd. *Mathematical Physiology: I: Cellular Physiology*. Springer, New York, NY, 2nd edition edition, October 2008. 3

[27] Ina Koch, Wolfgang Reisig, and Falk Schreiber, editors. *Modeling in Systems Biology: The Petri Net Approach*. Springer, London; New York, 2011 edition edition, October 2010. 2, 8, 9, 37, 38

[28] Sofia Lambropoulou, Stathis Antoniou, and Nikola Samardzija. Topological Surgery and its Dynamics. *arXiv preprint arXiv:1406.1106*, 2014. 50

[29] Michael J. Lawson, Brian Drawert, Mustafa Khammash, Linda Petzold, and Tau-Mu Yi. Spatial Stochastic Dynamics Enable Robust Cell Polarization. *PLoS Computational Biology*, 9(7):e1003139, July 2013. 3, 50

[30] Mark C. Leake. *Single-Molecule Cellular Biophysics*. Cambridge University Press, Cambridge; New York, 1 edition edition, March 2013. 49

[31] Ben Leimkuhler and Charles Matthews. *Molecular Dynamics: With Deterministic and Stochastic Numerical Methods*. Springer, 2015 edition edition, June 2015. 2

[32] Seth Lloyd. Quantum coherence in biological systems. *Journal of Physics: Conference Series*, 302:012037, July 2011. 1

[33] Richard Losick and Claude Desplan. Stochasticity and Cell Fate. *Science*, 320(5872):65–68, April 2008. 3

[34] Jay Newman. *Physics of the Life Sciences*. Springer, New York, 2008 edition edition, October 2008. 49

[35] Selim Olcum, Nathan Cermak, Steven C. Wasserman, and Scott R. Manalis. High-speed multiple-mode mass-sensing resolves dynamic nanoscale mass distributions. *Nature Communications*, 6:7070, May 2015. 5

[36] Mike Stay and Jamie Vicary. Bicategorical Semantics for Nondeterministic Computation. *Electronic Notes in Theoretical Computer Science*, 298:367–382, November 2013. 1

[37] M. Sturrock, A. Hellander, A. Matzavinos, and M. A. J. Chaplain. Spatial stochastic modelling of the Hes1 gene regulatory network: intrinsic noise can explain heterogeneity in embryonic stem cell differentiation. *Journal of The Royal Society Interface*, 10(80):20120988–20120988, January 2013. 3

[38] Jos MG Vilar, Hao Yuan Kueh, Naama Barkai, and Stanislas Leibler. Mechanisms of noise-resistance in genetic oscillators. *Proceedings of the National Academy of Sciences*, 99(9):5988–5992, 2002. 3

[39] Harel Weinstein. *Comprehensive Biophysics: Simulation and Modeling*, volume 9. Academic Press, Amsterdam; New York, 1 edition edition, May 2012. 2

[40] Darren J. Wilkinson. *Stochastic Modelling for Systems Biology, Second Edition*. CRC Press, Boca Raton, 2 edition edition, November 2011. 3, 8, 17, 33, 37, 38