

Mathematical Modelling of a Cytoskeletal System: Optimisation of a Plasmid Partitioning Mechanism Using Artificial Evolution

By

Minghua Yin

Supervised by Dr François Nédélec

A Project Report
in Partial Fulfillment of the Requirements for
the Degree of Master of Natural Sciences
at the University of Cambridge

Word Count: 5671

DECLARATION OF ORIGINALITY

'Mathematical Modelling of a Cytoskeletal System: Optimisation of a Plasmid Partitioning Mechanism Using Artificial Evolution'

I understand the University's definition of plagiarism. I declare that, in accordance with Discipline regulation 6, this dissertation is entirely my own work except where otherwise stated, either in the form of citation of published work, or acknowledgement of the source of any unpublished material.



Minghua Yin, Trinity College

04/05/2022

SUMMARY

Recently, computer simulations have been found to be very useful in cell research, with many different modelling tools now in development that aim to shed light on the mechanisms that allow cells to perform complex tasks. Cell division in particular is of great interest as an essential process that sustains life, but the large number of steps and cellular components involved means that it is difficult to gain a detailed understanding of how it works through experiments alone, highlighting the need for computational modelling. This project investigates how a particular artificial cell, modelled using the cytoskeletal simulation tool Cytosim, behaves after its DNA is replicated and the two copies need to be separated so that the cell can successfully divide. The cell was configured to have a relatively simple cytoskeleton consisting of microtubules assisted by plus-end directed motor proteins of a single type. A small number of cell parameters was selected and an optimisation method known as a genetic algorithm was used to look for values for these parameters that would maximise the cell's chances of successful partitioning. The cell's partitioning mechanism was found to be generally effective, often with a success rate of at least 0.9 in the fittest cases, although some possible limitations were identified that could warrant further investigation. The performance of the genetic algorithm was generally good and highly optimal solutions were often found, but certain avenues for further research, regarding both the way the algorithm is coded and the way it is configured, were identified that could help to better explain how the algorithm works and to improve some of the algorithm's less successful aspects.

ACKNOWLEDGEMENTS

I'm extremely grateful to Dr François Nédélec for his constant advice, support, patience and ambition for the whole duration of this project, and to the rest of the Nédélec group for the many thought-provoking chats and meetings.

I'm indebted to Yi Ting for teaching me about reverse-i-search in Bash, without which this project would have been excruciating.

Lastly, I'd like to thank Yuriy Tumarkin and Alejandro Epelde for enthusiastically letting me talk to them about, and show them pictures of, 'my big green balls'.

TABLE OF CONTENTS

DECLARATION OF ORIGINALITY	III
SUMMARY	IV
ACKNOWLEDGEMENTS	V
LIST OF TABLES	VII
LIST OF FIGURES	VIII
LIST OF ABBREVIATIONS	IX
1 Introduction	1
1.1 Computational modelling of cells	1
1.2 Cell division and the cytoskeleton	2
2 Results	7
2.1 Infinite energy case	9
2.2 Additional energy constraints	10
2.3 Partitioning checkpoint	16
3 Discussion	19
3.1 Conclusions	19
3.2 Other further work	21
4 Methods	22
4.1 Cytosim	22
4.1.1 Background	22
4.1.2 Main assumptions and techniques	23
4.1.3 Modifications	23
4.2 Genetic algorithm	25
4.2.1 Background	25
4.2.2 Implementation	25
REFERENCES	27
A Code Snippets	31
A.1 Modifications to Cytosim	31
A.2 Fitness functions	35
A.3 Heatmaps and fitness plots	37
B Configuration Files	43

LIST OF TABLES

2.0.1	The specifications for the artificial, rod-shaped cell used for the simulations. A few modifications were made for some of the simulations — these are explicitly stated in the relevant sections.	8
2.1.1	The parameters that the GA varied for the simulations in Subsection 2.1.	9
2.1.2	Two of the perfect scorers for the infinite energy simulation, demonstrating different optimal solutions.	10
2.2.1	The fittest individual when the GA was run with a population size of 50 and 40 simulations per individual.	12
2.2.2	The perfect scorer when the GA was run with a population size of 200 and 10 simulations per individual.	14
2.3.1	One of the fittest individuals from each of two runs of the GA. Ind 5 is from a run with the first evolutionary configuration, and Ind 6 is from a run with the second evolutionary configuration.	18
4.1.1	The three energy-costing actions that are accounted for in the simulations in Subsection 2.2 onwards. As with all properties defined in Cytosim, the values of the new properties are specified in the .cym configuration file, which makes the simulations highly customisable.	24

LIST OF FIGURES

1.2.1	The cell cycle for a typical eukaryotic cell (Alberts, 2015, Chapter 17).	2
1.2.2	The structure of a typical MT (Alberts, 2015, Chapter 16).	3
1.2.3	The mitotic spindle of a typical animal cell during metaphase, just before the sister chromatids are pulled apart (Alberts, 2015, Chapter 17).	4
1.2.4	The walking mechanism by which kinesin moves along an MT (Korosec and Forde, 2017).	5
2.2.1	Frames of a successful simulation for Ind 3 from Table 2.2.1. The two plasmids are represented by the green balls, which are purely visual as the plasmids are actually modelled as single-point beads. The red dots are complexes, some of which are free-floating in the cytoplasm and some of which are bound to one or two MTs.	12
2.2.2	Heatmap and fitness plots for runs of the GA with a population size of 50 and 40 simulations per individual. The red line is the maximum fitness plot and the white line is the average (mean) fitness plot.	13
2.2.3	Heatmaps and fitness plots for runs of the GA with the cell configured to have a very high viscosity of 100 Pa s. In each Subfigure, the red line is the maximum fitness plot and the white line is the average (mean) fitness plot.	15
2.3.1	Heatmaps and fitness plots for runs of the GA with a partitioning checkpoint added to the cell. In each Subfigure, the red line is the maximum fitness plot and the white line is the average (mean) fitness plot.	17
2.3.2	Frames of a successful simulation for Ind 5 from Table 2.3.1. Interestingly, no complexes attached to both MTs! The single MT attached to each plasmid seemed to be enough to push the plasmids towards the ends of the cell, fortunately to opposite ends.	18
4.2.1	The steps of the GA used in this project. The first stopping condition in Step 3 was in fact disabled in all the simulations by setting the threshold to 0.	26

LIST OF ABBREVIATIONS

GA	Genetic Algorithm
MAP	Microtubule-Associated Protein
MT	Microtubule
MTOC	Microtubule-Organising Centre

1 Introduction

1.1 Computational modelling of cells

Cells are incredibly complex building blocks of life and there remains much to be discovered about their inner workings. In recent decades, computer simulations have been increasingly employed to help better understand the details of cellular processes. For example, successful cell division requires numerous structural proteins, regulatory proteins and enzymes to correctly interact in several steps and checkpoints, and it is difficult to investigate specific mechanisms at high resolution by purely empirical means. Previous work that has benefitted from the use of simulations has included looking at the effect of chromosomal movements and rotations on the acceleration of mitotic spindle assembly (Paul et al., 2009), and showing that dynamic instability, a property of microtubules (see Subsection 1.2), greatly improves the efficiency of locating chromosomes (Holy and Leibler, 1994).

Ultimately, a central aim in computational biology is to continually develop better and more accurate models of cellular processes. Not only is this useful for understanding life, but it can also shed light on *unnatural* but successful new mechanisms not found in any known organisms. The potential implications this has in synthetic biology are difficult to predict, but they certainly exist.

This project focuses on cell division, particularly the mechanisms by which genetic material is transported in cells. The main aims of the project were as follows:

- to show that Cytosim, a cytoskeletal simulation tool (see Subsection 4.1 for more details), which uses a number of simplifying assumptions, can reproduce successful mitotic processes with different starting configurations and parameters;
- to improve the simulation tool by coding in simple energy considerations and checkpoints;
- to use a genetic algorithm (see Subsection 4.2 for more details) to optimise sets of parameters for various simulation configurations.




Fig. 1.2.1: The cell cycle for a typical eukaryotic cell (Alberts, 2015, Chapter 17).

1.2 Cell division and the cytoskeleton

Cell division is an essential mechanism by which all organisms reproduce, and thus is key to sustaining life. Although the precise steps vary between different species, all dividing cells must replicate their DNA and then accurately separate the two copies in space before splitting in two. In the eukaryotic cell cycle, this process of DNA partitioning takes place in *M phase* (see Figure 1.2.1), consisting of *mitosis*, during which the DNA is partitioned and nuclear division occurs, and *cytokinesis*, during which the cytoplasm is divided in two and the daughter cells are formed (Morgan, 2007). Similar steps occur in prokaryotic cells such as bacteria. The cell cycle is complex and involves precise activatory and inhibitory processes at various checkpoints by enzymes and regulatory proteins.

The movement of chromosomes and other genetic material in a cell is controlled by the cell's *cytoskeleton*, in particular by filaments such as *microtubules* (MTs) in eukaryotes (Aldaz et al., 2005) and *ParM*, an actin homologue, in bacteria (Salje and Löwe, 2008). MTs are hollow tubes consisting of 13 parallel protofilaments, each of which is made up of 8 nm-long subunits of $\alpha\beta$ -*tubulin* (see Figure 1.2.2) (Kerssemakers et al., 2006). The numerous longitudinal and lateral contacts between these tubulin heterodimers make MTs very stiff, and previous experiments have shown that an MT has similar mechanical properties to an elastic rod (Dogterom and Yurke, 1997). MTs are polar, with the *plus-end* polymerising much faster than the *minus-end*. Polymerisation of MTs is greatly assisted by nucleators, typically γ -*tubulin ring complexes*, which act as a blueprint to

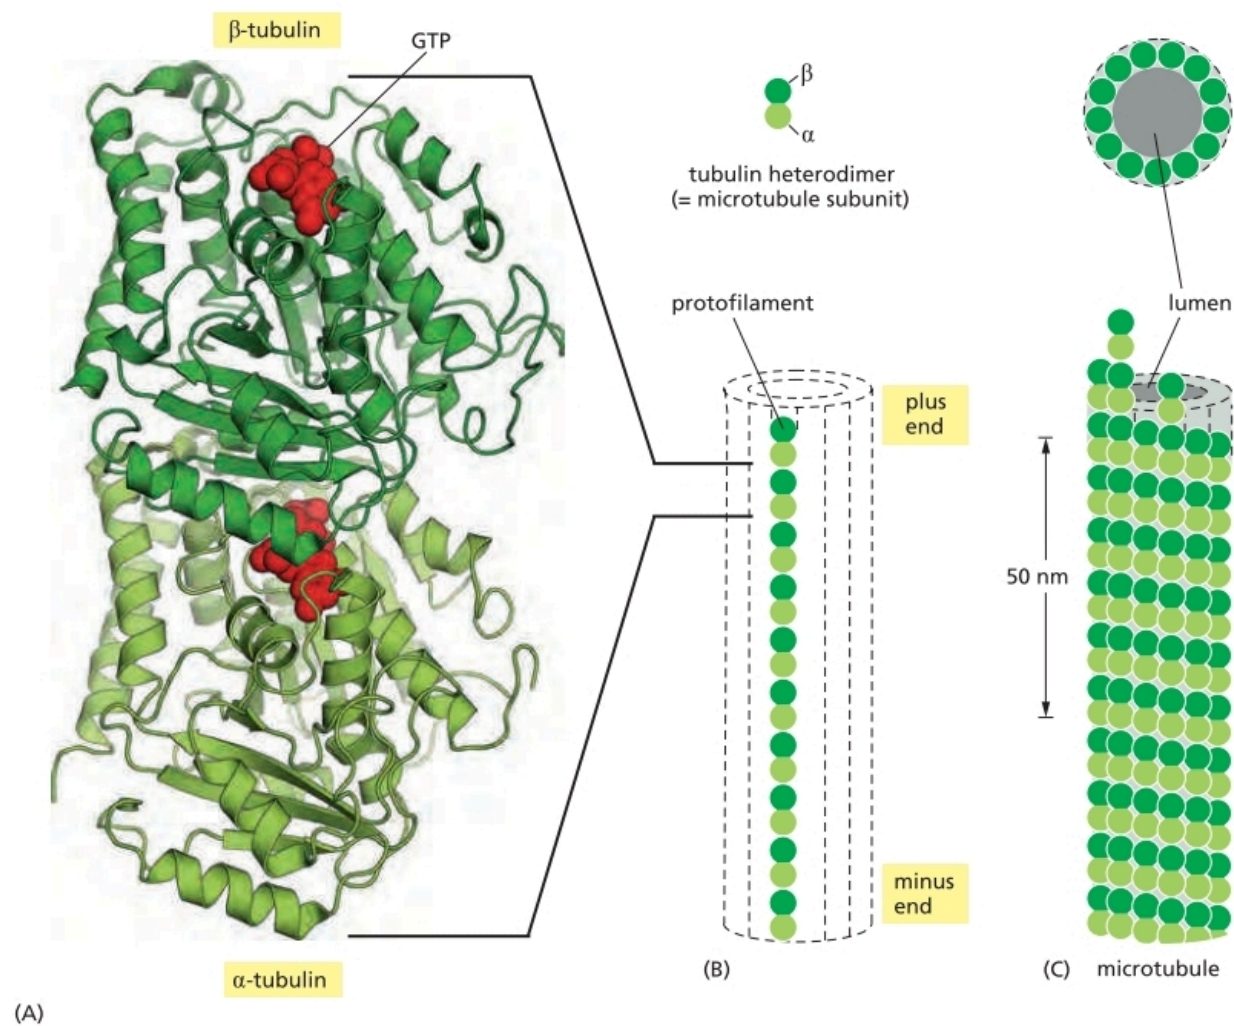


Fig. 1.2.2: The structure of a typical MT (Alberts, 2015, Chapter 16).

which the tubulin heterodimers can attach. MTs are often nucleated from specific locations called *microtubule-organising centres* (MTOCs), resulting in striking, radiating formations called *asters* (De Simone et al., 2016). Animal cells usually have a single MTOC called the *centrosome*, which duplicates before mitosis into two parts that move to opposite sides of the nucleus, forming the poles of the *mitotic spindle*, the structure that separates sister chromatids during mitosis (see Figure 1.2.3). Other eukaryotic cells also have mitotic spindles, although their structures may be somewhat different — for example, the cells of higher plants do not have centrosomes (Stoppin et al., 1994). It is not too difficult to model MTs at the scale of monomers, and models with this resolution are capable of producing reasonably accurate and illuminating results.




Fig. 1.2.3: The mitotic spindle of a typical animal cell during metaphase, just before the sister chromatids are pulled apart (Alberts, 2015, Chapter 17).

MTs and other filaments exhibit *dynamic instability*, where they can rapidly switch between shrinkage and growth (Hotani and Horio, 1988). The β -tubulin within each subunit is associated with a GTP molecule which is hydrolysed to GDP soon after the subunit attaches to the MT. (In actin and related filaments, ATP is found instead of GTP (Pernier et al., 2016).) Hydrolysis of the GTP reduces the binding affinity of the subunit for neighbouring ones and hence makes it more likely to dissociate from the polymer. An MT will therefore consist mostly of GDP-tubulin monomers, but at each end it is possible for there to be a GTP cap if the rate of addition of subunits at that end is greater than the rate of GTP hydrolysis. A GTP cap favours growth, while an end containing GDP-tubulin favours depolymerisation (by a factor of about 100 compared to GTP-tubulin). MTs are commonly observed to alternate between a state of slow growth and one of rapid depolymerisation. These changes of state are random, with a probability that depends on the concentration of free GTP-tubulin subunits. A transition from growth to shrinkage is called a *catastrophe* and a transition from shrinkage to growth is called a *rescue*.

Eukaryotic cells contain many different types of *microtubule-associated proteins* (MAPs), which attach to MTs and alter their behaviour, such as by initiating rescue or catastrophe, or by helping to transport the MTs through the cytoplasm (Galjart, 2010). In particular, *kinesins* and *dyneins* are *motor proteins* that bind to and move along MTs. Most (but not all) kinesins move towards the plus-end of an MT while dyneins are minus-end directed (Kikkawa, 2013). Some kinesins can




Fig. 1.2.4: The walking mechanism by which kinesin moves along an MT (Korosec and Forde, 2017).

bind to multiple MTs, while others have sites to which intracellular cargo such as chromosomes can bind (Verhey et al., 1998). Kinesins typically have two heavy chains, each of which has a tubulin-binding domain (a ‘foot’) and a site that binds ATP (Howard et al., 1989). The two feet of a kinesin molecule attach to adjacent heterodimers on an MT. Kinesin moves via a ‘walk’ (see Figure 1.2.4): hydrolysis of one ATP molecule causes the back foot to detach from the MT, swing forwards and reattach in front of the other foot (Vale and Milligan, 2000). Hence the length of each step that kinesin makes is 8 nm. Motor proteins attach to MTs at a rate that depends on their concentration, while their detachment is affected by the force experienced by the molecular link in the motor protein, e.g. due to tension from cargo. The relationship between the detachment rate k_{off} and the magnitude f of the experienced force is exponential and obeys *Bell’s law*, $k_{\text{off}} = k_0 \exp(f/f_0)$, where k_0 is the load-free rate and f_0 is a characteristic force that depends on temperature and a molecular scale (Walcott, 2008). Intracellular events such as motor protein detachment that can be described using well-defined equations can be readily modelled.

Motor proteins are essential for the action of the mitotic spindle. Dynein and different types of kinesin help to organise the spindle and transport cargo along MTs (Verhey et al., 2011). For example, in animal cells, kinesin-5 binds to two antiparallel MTs originating from opposite poles

and pushes the MTs part, since each pair of feet moves towards the plus-end of its respective MT. Kinesin-4 and kinesin-10 push attached chromosomes towards the plus-ends of MTs. Kinesin-14 is minus-end directed and tends to pull the poles together, and dynein links MTs with the actin cytoskeleton at the cell cortex, providing anchorage to the poles (Walczak et al., 2010). Despite each motor protein having a simple action, together they allow the spindle to perform important and complex tasks such as organising chromosomes and separating sister chromatids. Many other MAPs also play important roles — for example, proteins such as CLASP and CLIP-170 bind to MTs and promote rescue (Fees and Moore, 2019). The simplicity of each motor protein's action also means that it is relatively straightforward to make computational models of the cytoskeleton that can reproduce the complex behaviour observed in real cells. It is worth noting that, although the focus in this section has been on animal cells, the cytoskeletons of other eukaryotic cells and prokaryotic cells will have many similarities to the above, and the basic concepts described here are applicable to cells from a wide variety of species.

2 Results

A significant amount of time in this project was spent working on modifications to the Cytosim code that were required for the simulations (see Subsection 4.1.3) and completing the code for the genetic algorithm (GA) (see Subsection 4.2.2). The utility of the GA was then tested with a large number of simulations. Examples of the configuration files used in the project, for both the simulations and the GA, can be found in Appendix B.

The simulations are of an artificial, rod-shaped cell, similar in size to a typical bacterial cell. Inspiration was loosely taken from such single-celled organisms as the yeast *Schizosaccharomyces pombe* and the bacterium *Escherichia coli*, forming a sort of ‘hybrid’ cell, with its (already replicated) DNA found in two plasmids (modelled as single-point beads), but with MT fibers as in eukaryotes. The cell has two types of protein: complexes, which are couples consisting of plus-end directed kinesins (similar to kinesin-5), and nucleators, which are singles. Instead of the cell having a specialised centrosome, the nucleators are directly attached to the plasmids so that MTs are nucleated with their minus-ends bound to the plasmids. Complexes are synthesised free-floating in the cytoplasm and can bind to a pair of antiparallel MTs, one from each plasmid. Since the kinesins are all plus-end directed, a doubly attached complex will push the MTs away from each other and hence move the plasmids further apart — this is the primary mechanism by which the artificial cell aims to separate its DNA. The MTs exhibit dynamic instability with a hydrolysis rate of 1 s^{-1} , which corresponds to a catastrophe rate of approximately 0.12 s^{-1} (see Subsection 4.1.2).

Table 2.0.1 gives the exact specifications of the cell and its components. Most of the values were taken from empirical data from past research (Campbell and Mullins, 2007; Footer et al., 2007; Garner et al., 2004; Gittes et al., 1993).

This hybrid cell was designed to strike a good balance between simplicity, given the time constraints of the project and the limitations of Cytosim itself, and complexity, so that the cell’s mechanisms would not be too distant from reality and a significant step could be made towards the long-term aim of modelling and optimising more complex cells (see Subsection 3.2). A rod shape was chosen (rather than, say, a sphere) as it provided a clear, simple way of determining whether

Property	Value	Property	Value
Cell		Nucleator	
Radius	$0.4 \mu\text{m}$	Stiffness	$1000 \text{ pN } \mu\text{m}^{-1}$
Length	$2 \mu\text{m}$	Nucleation rate	1 s^{-1}
Viscosity	0.1 Pa s	Init length of created MT	$0.01 \mu\text{m}$
Plasmids		Kinesin	
Confinement stiffness within cell	$100 \text{ pN } \mu\text{m}^{-1}$	Binding rate	10 s^{-1}
MTs		Unbinding rate	0.3 s^{-1}
Rigidity	$20 \text{ pN } \mu\text{m}^2$	Unbinding force	2.5 pN
Confinement stiffness within cell	$100 \text{ pN } \mu\text{m}^{-1}$	Unloaded speed	$0.02 \mu\text{m s}^{-1}$
Unit length	$0.008 \mu\text{m}$	Stall force	6 pN
Shrinking speed	$0.5 \mu\text{m s}^{-1}$	Kinesin complex	
Hydrolysis rate	1 s^{-1}	Stiffness	$100 \text{ pN } \mu\text{m}^{-1}$
Growing force	1.7 pN	Diffusion coefficient	$0.05 \mu\text{m}^2 \text{ s}^{-1}$
Total # monomer units available	200	Natural length	$0.025 \mu\text{m}$

Table 2.0.1: The specifications for the artificial, rod-shaped cell used for the simulations. A few modifications were made for some of the simulations — these are explicitly stated in the relevant sections.

Property	Range
Growing speed of MTs ($\mu\text{m s}^{-1}$)	[0.05, 0.5]
Binding range for kinesin (μm)	[0.01, 0.05]
Number of nucleators per plasmid	[1, 12]
Number of complexes	[50, 5000]

Table 2.1.1: The parameters that the GA varied for the simulations in Subsection 2.1.

a partitioning was successful: the long axis of the cell was positioned to coincide with the x -axis, with the cell centred at the origin, so that the plane $x = 0$ divided the cell evenly. The partitioning could then be considered successful if at the end of the simulation the plasmids were on opposite sides of this plane.

The two DNA beads are initially placed along the x -axis, at $x = -0.3$ and $x = -0.5$ respectively (units are μm). The simulations were configured with a time step of 0.001 s and a runtime of 100 s, which is the order of magnitude within which the cell should be able to partition its DNA.

The GA was configured with a population size per generation of 50, and was set to terminate (see Step 3 in Figure 4.2.1) either (a) if the maximum fitness in the most recent generation is greater than 0.9 or (b) after 50 generations.

2.1 Infinite energy case

Initially, four properties were decided on for the GA to optimise. These, along with their specified ranges, are listed in Table 2.1.1. The ranges were chosen to be realistic and consistent with empirical data for naturally occurring cells of a similar size (Verhey et al., 1998; Howard et al., 1989; Holy and Leibler, 1994; Matsuyama et al., 2006).

For these simulations, no modifications were made to the base Cytosim code. The GA was first run with a population size of 50 and with 40 simulations for each set of parameters (aka individual). The fitness function was defined as follows: for each of the 40 simulations, the final x -coordinates of the plasmids are examined. If they have opposite signs, then the partitioning was a success and this particular simulation gets a score of 1. Otherwise, the simulation gets a score of 0. A mean

Property	Ind 1	Ind 2
Growing speed of MTs ($\mu\text{m s}^{-1}$)	0.2653	0.1824
Binding range for kinesin (μm)	0.0473	0.0447
Number of nucleators per plasmid	9	11
Number of complexes	2360	4379

Table 2.1.2: Two of the perfect scorers for the infinite energy simulation, demonstrating different optimal solutions.

score over all 40 simulations is then calculated, which corresponds to the fitness of the individual.

The GA terminated at just generation 0 and found four individuals with a perfect fitness score of 1, two of which are shown in Table 2.1.2. Further runs of the GA yielded yet more optimal solutions, indicating that the ranges from Table 2.1.1 are capable of producing positive results. Additions to the Cytosim code, and eventually the cell configuration, were then made to model the cell more realistically, as detailed in the next Subsections.

2.2 Additional energy constraints

The Cytosim code was modified to model simple energy constraints. As described in Subsection 4.1.3), the cell was given an energy income and the energy costs of three key actions were specified.

- The synthesis of a complex was set to cost 400 ATP.
- The movement of a kinesin by one step was set to cost 1 ATP. Since each step is 8 nm long, the movement cost per μm is 125 ATP.
- The polymerisation of an MT was set to cost 13 GTP (ATP equivalent) per monomer length (8 nm). The simulation was set up so that all 200 monomer units are available to the cell right from the start, but this cost accounts for the replenishment of one GTP molecule per attached heterodimer (to replace the GTP that is hydrolysed), with each MT consisting of 13 protofilaments. The movement cost per μm is thus 1625 ATP equivalent.

A very rough, order-of-magnitude estimate of an appropriate income was calculated by taking one of the optimal solutions from Subsection 2.1 (Ind 1 from Table 2.1.2) and estimating the energy used by the cell during the simulation. With $v_{\text{grow}} = 0.2653$, $n_{\text{nuc}} = 9$ and $n_{\text{com}} = 2360$, and the assumption that each kinesin travels a total of $1 \mu\text{m}$, the energy use per second E was estimated to be about 260000 units of ATP¹.

The configuration was slightly changed so that the complexes would now be added to the cell one by one, at a rate in the range $[10, 100] \text{ s}^{-1}$, to roughly match the range for the number of complexes in Subsection 2.1. This parameter replaced the previously used number of complexes.

The GA was run with these modifications, but it was found that the income was too low for the number of successful partitionings to be significant. It is important to emphasise that the estimate obtained above is extremely loose and so it is justified to change the leading digit or increase the energy intake by a reasonable amount. Once the income was increased to at least 800000 ATP per second, the results became much more noteworthy. The results below are all from simulations where the income was 800000 ATP per second.

Again with a population size of 50 and 40 simulations per individual, the GA terminated at generation 16 with a maximum fitness of approximately 0.925². Table 2.2.1 shows the individual with this fitness. The frames in Figure 2.2.1 show a graphical rendering of one successful simulation of Ind 3. The heatmap in Figure 2.2.2 shows the distribution of fitnesses for each generation, and the plotted lines show the maximum fitnesses and average fitnesses across all generations. The increase in both maximum and average fitness is noticeable, demonstrating correct function of the GA in producing generally fitter individuals from one generation to the next. Sometimes the maximum fitness decreased slightly from one generation to the next, due to no individuals being copied into the next generation via elitism (see Figure 4.2.1).

The GA was then rerun, but this time with a population size of 200 and 10 simulations per individual, so that the total number of calculations made by the GA stayed the same as before. This time, an individual (shown in Table 2.2.2) with a perfect fitness of 1 was found in generation 0. Further runs of the GA, with both evolutionary configurations, yielded similar results, with the second configuration consistently producing generations with higher maximum fitnesses earlier on,

¹The full equation used was $E = 2n_{\text{nuc}} \cdot 400(v_{\text{grow}} + 0.01)/0.008 + n_{\text{com}}(400 + 2/0.008)/100$.

²There was a small bug in the Cytosim code which meant that the number of simulations was occasionally slightly less than 40.

Property	Ind 3
Growing speed of MTs ($\mu\text{m s}^{-1}$)	0.4665
Binding range for kinesin (μm)	0.0500
Number of nucleators per plasmid	12
Complex synthesis rate (s^{-1})	12.8571

Table 2.2.1: The fittest individual when the GA was run with a population size of 50 and 40 simulations per individual.




Fig. 2.2.1: Frames of a successful simulation for Ind 3 from Table 2.2.1. The two plasmids are represented by the green balls, which are purely visual as the plasmids are actually modelled as single-point beads. The red dots are complexes, some of which are free-floating in the cytoplasm and some of which are bound to one or two MTs.




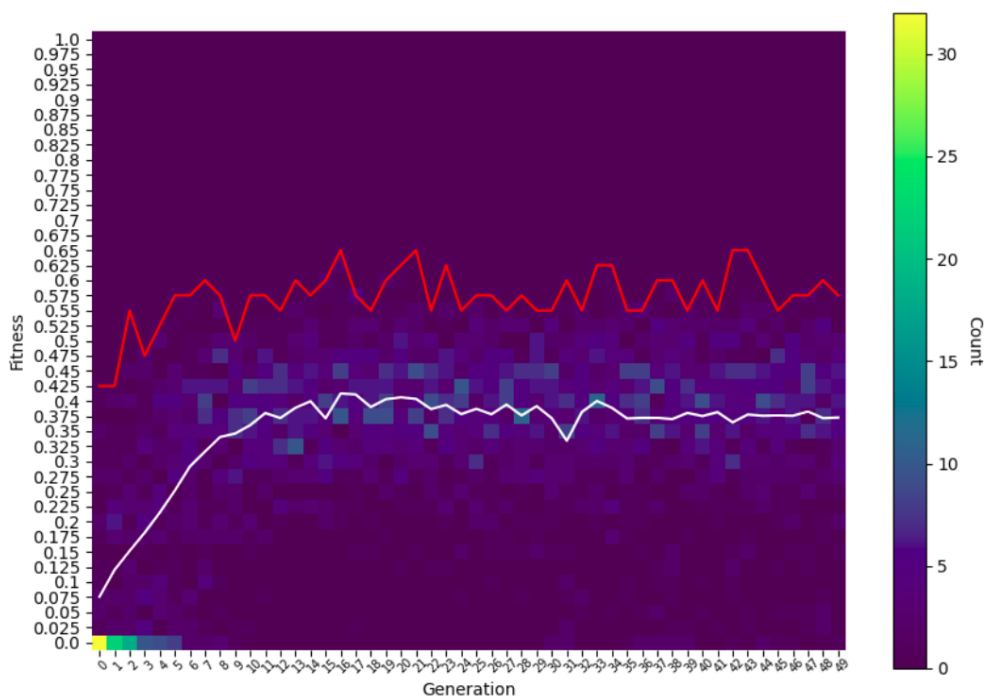
Fig. 2.2.2: Heatmap and fitness plots for runs of the GA with a population size of 50 and 40 simulations per individual. The red line is the maximum fitness plot and the white line is the average (mean) fitness plot.

Property	Ind 4
Growing speed of MTs ($\mu\text{m s}^{-1}$)	0.4400
Binding range for kinesin (μm)	0.0367
Number of nucleators per plasmid	11
Complex synthesis rate (s^{-1})	94.2857

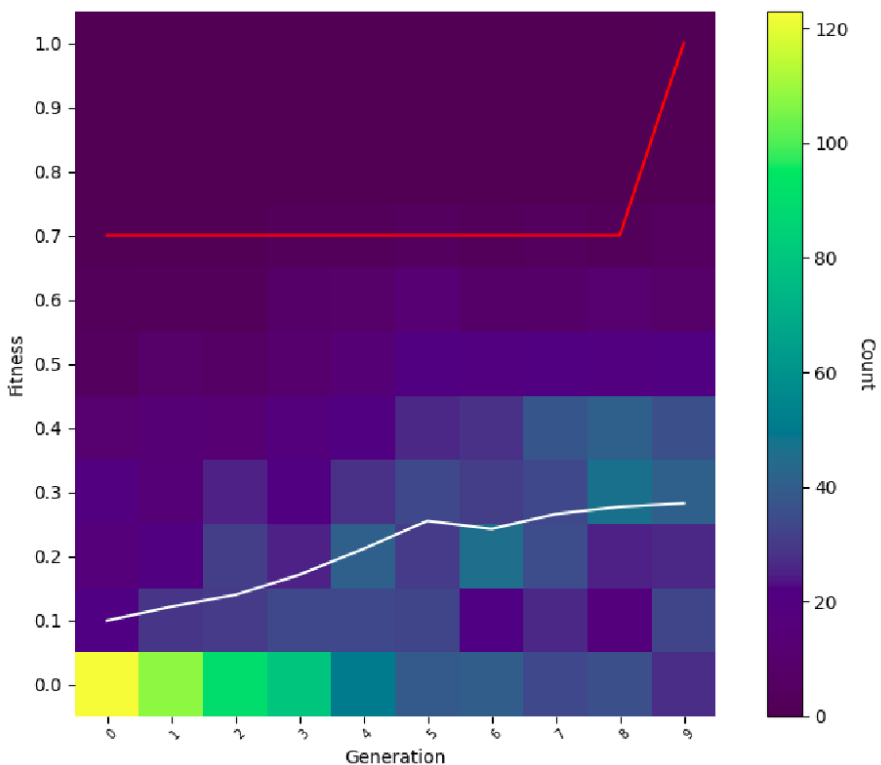
Table 2.2.2: The perfect scorer when the GA was run with a population size of 200 and 10 simulations per individual.

often in generation 0, compared with the first configuration.

To further test both the GA and the suitability of the artificial cell's partitioning mechanism, the viscosity of the cell was increased from 0.1 Pa s to a very high 100 Pa s, rendering the plasmids virtually immobile unless pushed by the cell's MT-complex system (the 'spindle'). Figure 2.2.3 shows the heatmaps and plots for typical GA runs with the first and second evolutionary configurations respectively. With the first configuration, the GA typically terminated at generation 49 (the maximum possible generation before the GA is programmed to stop) with a maximum fitness that rarely exceeded 0.625. While the average fitness increased as generation number, the change (if any) in maximum fitness was much less, with perhaps only a slight increase compared to the very first generations. A more successful GA would be expected to produce a much clearer improvement in the maximum fitness over generations. The difference in success between the first and second configurations is very stark, with the second configuration often producing individuals with very high fitnesses after a small number of generations and a more significant increase in maximum fitness over generations. Even given this unrealistic and extreme viscosity, the maximum fitnesses achieved by the GA indicate that the cell's spindle significantly improves the cell's chances of successful partitioning, since without it it would be expected that the fitness of any individual would be close to 0 for these simulations.



(a) Population size of 50, 40 simulations per individual.



(b) Population size of 200, 10 simulations per individual.

Fig. 2.2.3: Heatmaps and fitness plots for runs of the GA with the cell configured to have a very high viscosity of 100 Pa s. In each Subfigure, the red line is the maximum fitness plot and the white line is the average (mean) fitness plot.


2.3 Partitioning checkpoint

On inspection of the final x -coordinates of the plasmids for many different simulations, it was discovered that sometimes the plasmids were still very close to the origin, and to each other, at the end ($|x| < 0.1$, say). Real cells typically separate their DNA further from each other so that the daughter cells are better defined and successful cytokinesis can occur. The next addition to the Cytosim code was a checkpoint at each time point that checks if (a) the plasmids are on opposite sides of the cell and (b) $|x| > 0.5$ for each plasmid — if so, the cell is considered to have successfully partitioned its DNA. If this still has not happened after 100 s, the cell is deemed to have failed. This was inspired by real checkpoints that are found at key moments in the cell cycle in most species, with the cell cycle unable to pass a checkpoint until the conditions there are met (Hartwell et al., 1974).


The fitness function for the GA was also modified: now, if partitioning is unsuccessful, the simulation gets a score of 0. If partitioning is successful, the time t taken to partition determines the score s ($0 < s < 1$) via a linear scale, with a smaller t corresponding to a higher score. Explicitly, $s = 1 - t/100$.

The GA was run multiple times with both evolutionary configurations from before. Figure 2.3.1 shows typical results for the two configurations respectively. Regardless of the evolutionary configuration, most of these runs did not produce a significant improvement in maximum fitness over generations, and any improvements in average fitness appeared to be small. Closer inspection of some of the fittest individuals (see Table 2.3.1) showed a very mixed picture and significant inconsistencies: for example, simulations of Ind 5 were sometimes very short, with some successful partitionings occurring in less than 10 s (see Figure 2.3.2), but often the DNA did not successfully separate at all. One interesting feature that was discovered when viewing some of the graphical outputs was that in some successful simulations, especially when the number of nucleators was low, no complexes ever attached to two MTs and hence there was no pushing action on the MTs from the motor proteins. Despite this, successful partitioning was still sometimes achieved, apparently because the growing MTs were able to push the plasmids towards the ends of the cell.

Similarly to what was found in Subsection 2.2, the second evolutionary configuration again performed much better than the first in finding individuals with notably higher fitnesses more quickly.



(a) Population size of 50, 40 simulations per individual.



(b) Population size of 200, 10 simulations per individual.

Fig. 2.3.1: Heatmaps and fitness plots for runs of the GA with a partitioning checkpoint added to the cell. In each Subfigure, the red line is the maximum fitness plot and the white line is the average (mean) fitness plot.

Property	Ind 5	Ind 6
Growing speed of MTs ($\mu\text{m s}^{-1}$)	0.4612	0.4682
Binding range for kinesin (μm)	0.0340	0.0180
Number of nucleators per plasmid	1	1
Complex synthesis rate (s^{-1})	20.0000	62.8571
Fitness	0.024100	0.097320

Table 2.3.1: One of the fittest individuals from each of two runs of the GA. Ind 5 is from a run with the first evolutionary configuration, and Ind 6 is from a run with the second evolutionary configuration.

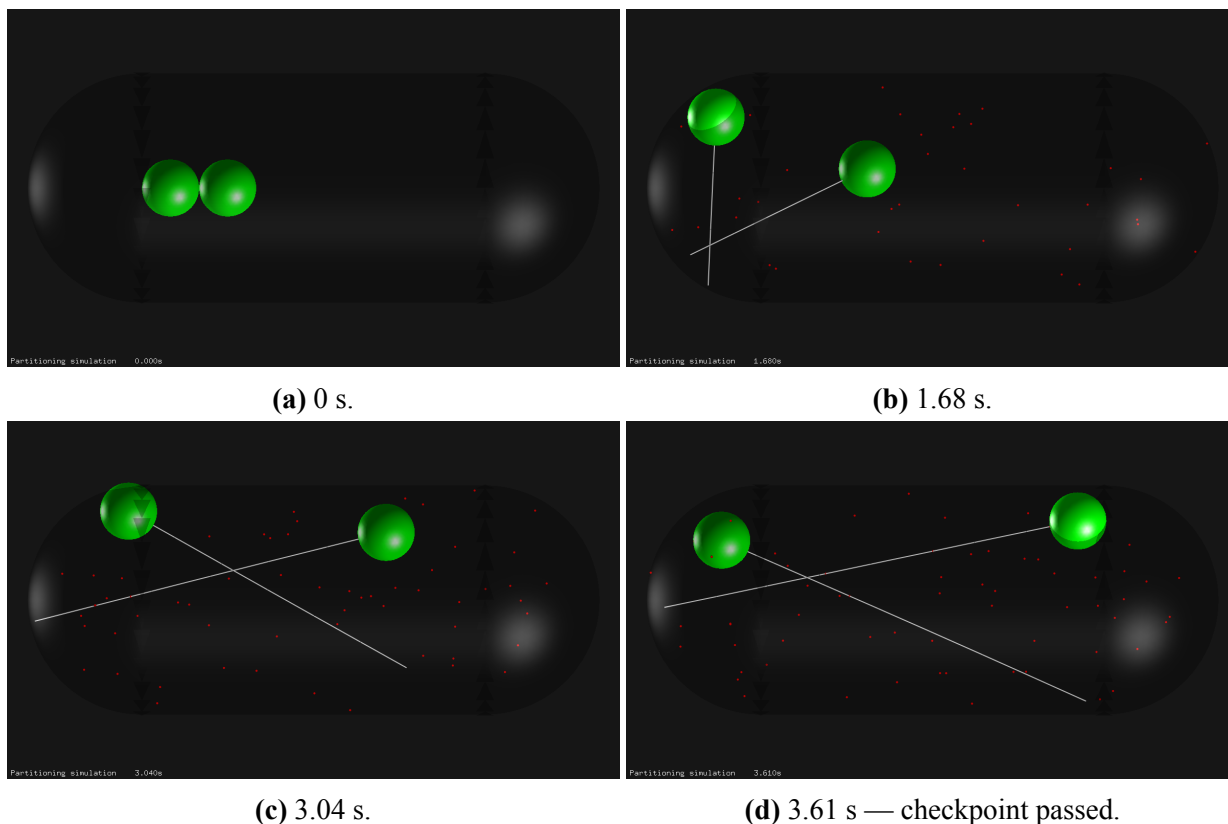


Fig. 2.3.2: Frames of a successful simulation for Ind 5 from Table 2.3.1. Interestingly, no complexes attached to both MTs! The single MT attached to each plasmid seemed to be enough to push the plasmids towards the ends of the cell, fortunately to opposite ends.

3 Discussion

A large number of simulations, all of the same basic cell but with somewhat different configurations, was successfully run and produced results of varying quality. Given that the use of relatively specialised cellular simulation tools such as Cytosim is currently burgeoning, and that up until now GAs have rarely been used for such problems as DNA partitioning, the results from this project are very promising and naturally inspire further avenues of investigation.

3.1 Conclusions

Even with relatively simple individual mechanisms, the components of a cell exhibit self-organisation that allows the cell to perform important tasks

In the artificial cell, each component has a small number of simple actions: for example, a kinesin can bind to an MT, unbind from an MT or walk along an MT; and an MT can grow or shrink. Given the right parameters, these components work together and communicate indirectly with each other to produce well-organised systems. In most of the simulations, it was observed that eventually the MTs attached to each plasmid tended to be longer on the side facing the other plasmid, and tended to align themselves with the long axis (x -axis) of the cell. The shape of the resulting structure consisting of connected plasmids, MTs and complexes is very reminiscent of the structures observed in real cells. The simulations confirm that no sophisticated communication or interaction between a cell's components is required for complex larger-scale behaviour to occur.

Computational models are a very useful and efficient way to investigate cellular dynamics

It was demonstrated that computer simulations can produce artificial cells with a strong likeness to natural ones. Making adjustments to a simulation is much easier, faster and cheaper than trying to engineer a real cell, and parallel simulations can be fast-forwarded and run much more speedily than physical experiments. While wet laboratory work is undeniably essential in biological research,

computational models can be of great assistance in improving our understanding of various cellular mechanisms.

The artificial cell considered in this project is reasonably successful at partitioning its DNA

In many of the runs in Subsections 2.1 and 2.2, parameters that resulted in (almost) perfect partitioning behaviour were found very early on, even when energy constraints were introduced. For the normal viscosity runs in Subsection 2.2 with the first evolutionary configuration, the typical maximum fitness found was somewhat lower but still above 0.8, which is comparable to (perhaps slightly worse than) success rates typically found in nature. The effectiveness of the cell's spindle was very clearly demonstrated by the results from the high viscosity runs, where despite the unrealistic conditions a maximum fitness of over 0.5 was still attained.

After the partitioning checkpoint was added for the simulations in Subsection 2.3, the performance of the cell was not as consistent. There are a few possible reasons for this. Firstly, it is possible that the cell's mechanism has certain limitations that prevent it being able to consistently push its plasmids far enough apart to pass the checkpoint — indeed, examination of frames of some of these simulations showed that regular unbinding of complexes and MT catastrophes did often stall the separation of the plasmids, decreasing the effectiveness of the cell's mechanism. Secondly, the checkpoint may have been too strict, and a greater success rate might be observed if the $|x| > 0.5$ threshold were relaxed a little. Thirdly, the fitness function could certainly be improved, as the linear scale used meant that a high score could only be obtained if the cell managed to partition its DNA within tens of seconds, which is not very realistic. A function with a flatter peak such as $s = 1 - (t/100)^3$ or similar could probably produce better scaled fitnesses. All of these possibilities could be investigated by making further tweaks to the configurations of the cell and the GA.

The effectiveness of a GA is significantly affected by the way it is configured

The stark difference in performance between the two evolutionary configurations in Subsections 2.2 and 2.3 shows that it is important to strike a good balance between population size and number

of replicates (or other parameters that affect the total runtime of the GA). Despite the second configuration having a smaller number of simulations per individual, the much larger population size cast out a much wider net and drastically increased the ability of the GA to find more optimal solutions. However, if the number of simulations is too low, then the errors in the calculated fitnesses are high, potentially leading to inaccurate results. It would certainly be interesting to experiment with different evolutionary configurations to see if an even more optimal one can be found for these simulations.

Other parameters in the evolutionary configuration, or certain details of the GA, could also be tweaked, which might lead to further improvements in the GA. For example, it would be interesting to see what would happen if the elitism proportion were increased and the mutation proportion decreased (see Figure 4.2.1 for further explanation), or vice versa. Some GAs perform crossover at multiple sites rather than just one (Mitchell, 1998, Chapter 1) — it could be worth investigating what effect this change would have.

Perhaps one suboptimal feature of the GA used in this project is that, for some of the runs, while the average fitness is often significantly improved over generations, the improvement in maximum fitness tends to be less great. Further investigation would be required to determine to what extent this issue is caused by the artificial cell's partitioning mechanism, and to what extent by the GA. For example, increasing the elitism proportion would, at the very least, reduce the decreases in maximum fitness that sometimes occur from one generation to the next.

3.2 Other further work

In addition to the suggestions in Subsection 3.1 above, other natural next steps would be to try modelling more complex cells with more advanced spindles, eventually building up to cells on the complexity level of mammalian cells, as well as experimenting with more complex artificial mechanisms for DNA partitioning. It would certainly be possible to take small steps that build on this project: for example, the addition of rescuer proteins that stop MTs from shrinking was considered, but time constraints meant that this was not properly investigated. In order to accurately model the mitotic spindle in a human cell, for example, more modifications to Cytosim would have to be made, resulting in a more accurate, and hopefully more useful, simulation tool.

4 Methods

4.1 Cytosim

4.1.1 Background

*Cytosim*¹ is a software suite designed for simulating cytoskeletal dynamics. It is capable of modelling large systems of flexible filaments and associated proteins using a combination of mechanical and stochastic calculations. The Cytosim project was initiated in 1995 by François Nédélec, who is also the current principal maintainer, and it is currently being developed by the Nédélec group at the Sainsbury Laboratory, Cambridge University. The software is run by first compiling the C++ code and then inputting a `.cym` configuration file, which specifies the properties of the system and its components, including

- the time length and time step of the simulation;
- spaces, which represent cells;
- objects such as `fibers`, which are linear strings of points, and `beads`, `spheres` and `solids`, which can represent DNA and other molecules;
- `hands`, which can represent proteins (or parts of proteins) with various functions and must be placed in a `single` or `couple` object;
- any events that occur;
- what outputs to print.

Key executable files that can be compiled from the Cytosim code include `sim`, which runs simulations, and `play`, which produces graphics for a simulation either live or using data from a `sim` run.

¹<https://cytosim.org/>

4.1.2 Main assumptions and techniques

Cytosim makes a number of simplifying assumptions for more efficient simulation (Nedelec and Foethke, 2007). Fundamentally, it uses a fixed time step, agent-based simulation method. At each time point, each component of the simulation (including events and objects) has its step function called which determines its state for the next period of time. Points are taken to be zero-dimensional, and `fibers` are one-dimensional sequences of points. As is appropriate for the movements in the cytoplasm, the system is assumed to be at low Reynolds number, meaning inertia is neglected in all calculations. The boundary of the cell is taken to be frictionless, so that the force on any object from the boundary acts perpendicularly to the boundary, and the forces from the boundary are assumed to be linear (Belmonte et al., 2017, Appendix).

Each point-like object experiences Brownian dynamics described by an *overdamped Langevin equation*, $\xi \frac{dx}{dt} = f(x, t) + B(t)$, where x is the position vector of the object, ξ is a drag coefficient, $f(x, t)$ accounts for the deterministic forces and $B(t)$ represents a random fluctuation. Multi-point objects such as `fibers` are modelled using points that represent discrete segments. Molecular links are assumed to be Hookean. For objects attached to `fibers`, detachment rates are determined using Bell's law (as outlined in Subsection 1.2), and the speed v of an attached motor under a load is related to its unloaded speed v_0 (as specified in the configuration file) by $v = v_0(1 - f_{\parallel}/f_m)$, where f_{\parallel} is the force parallel to the `fiber` and f_m is the specified stall force. Catastrophes for a dynamic `fiber` are modelled using a rate k_{cat} which is related to a specified GTP hydrolysis rate, k_{hyd} , by $k_{\text{cat}} = 3k_{\text{hyd}}^2/k_{\text{gro}}$, where k_{gro} is the growth rate of the `fiber` in (no. of monomer units) s^{-1} . Stochastic events that depend on rates or probabilities are simulated by Gillespie or modified Gillespie approaches.

To solve the Langevin equations at each time step, Cytosim uses *implicit integration*, which has a 1000-fold efficiency improvement over explicit numerical integration.

4.1.3 Modifications

In order to run the simulations in Subsection 2.2 onwards, various additions were made to the base Cytosim code. A `cash` variable was added which stores the total amount of energy, in units of molecules of ATP/GTP, available to the cell, and an new property called `income` was added to

Action	Name of property	Events that require this action
Synthesis of a single or complex	synthesis_cost	Any attempt to add a single or complex to the cell
Movement of a hand which is attached to a fiber by one step	movement_cost (per μm)	Any attempt by an attached hand (of a single or complex) to move along a fiber
Polymerisation of a fiber	unit_energy_cost (per μm)	Any attempt to increase the length of a fiber, either ‘naturally’ or with assistance from a nucleator

Table 4.1.1: The three energy-costing actions that are accounted for in the simulations in Subsection 2.2 onwards. As with all properties defined in Cytosim, the values of the new properties are specified in the .cym configuration file, which makes the simulations highly customisable.

represent the amount of energy gained by the cell per second. It is assumed that, for the simple artificial cell, there are three actions that require a significant amount of energy. These are detailed in Table 4.1.1 together with the names of the new properties that represent the energy costs involved. For each energy-costing event, the wealth of the cell is checked and the event occurs if and only if the cost of the event is not greater than the cash.

In addition to the above, for the simulations in Subsection 2.3, extra lines of code were added to the main step function, which is called at each time point, that check whether (a) the two DNA beads are on opposite sides of the cell (i.e. their x -coordinates have opposite signs) and (b) the absolute values of their x -coordinates are both greater than a certain threshold (hard-coded as 0.5). If both conditions are true, then the DNA is considered to be successfully partitioned, so the program prints the time taken to partition and then restarts the simulation (unless there have been enough replicates). Otherwise, the simulation runs for the specified time length as before.

Code snippets for some of these additions can be found in Appendix A.

4.2 Genetic algorithm

4.2.1 Background

A *genetic algorithm* (GA) was extensively used in this project as the optimisation method for finding parameters that enable the highest success rate for completing a cellular task. This technique is inspired by the process of evolution by natural selection and draws many parallels with it, hence the process by which it works is known as *artificial evolution* (Mitchell, 1998; Rupp and Nédélec, 2012). Given a starting *population* A_0 consisting of sets of randomly generated parameters (within specified bounds) and a *fitness function* f , GAs aim to produce fitter successive *generations* A_i , $i \geq 1$. The GA used in this project encodes the values of the parameters as follows: for any *individual* $x_j \in A_0$, the value of each parameter p is encoded as a bit string (a *gene*) whose length is specified. Given specified bounds on p , the value of p is calculated from the bit string via a linear scale. The bit strings for all the different parameters are then concatenated into a *genome* for convenience.

The steps of the GA are detailed in Figure 4.2.1. The intention of the elitism and crossover operations is to create a fitter generation by retaining some of the attributes of the fittest individuals in the current population but also allowing the mixing of genes (or parts of genes) to try to form individuals with the ‘best of both worlds’. The mutation operation aims to prevent the GA from converging to a local maximum by introducing perturbations and thus increasing diversity at each bit position. A more rigorous mathematical justification for the validity of GAs, including the *Schema Theorem*, can be found in Chapter 1 of Mitchell (1998).

4.2.2 Implementation

The GA used in this project is written in Python and was originally developed by Maud Formanek and François Nédélec. The author of this report helped to fix bugs and refine the code, as well as writing the code for the fitness functions relevant to this project. The code for the fitness functions can be found in Appendix A.

1. Start with a population A_0 consisting of individuals x_j , where the parameter values for each x_j are encoded in a randomly generated genome.
2. Given the current population A_i , form the offspring population A_{i+1} by calculating the fitness $f(x_j)$ of each $x_j \in A_i$ and then performing the following operations in order:
 - (a) *Elitism*: choose an integer $0 < S < |A_i|$, where $|A_i|$ is the population size, and select the fittest S individuals of A_i to be copied into A_{i+1} .
 - (b) *Crossover*: choose an integer $0 < P < |A_i| - S$. Select $2P$ individuals, called *parents*, from A_i with replacement using *fitness rank selection*, i.e. the probability of an individual being selected scales linearly as its rank. Now recursively choose two different parents from the selection and form a *child* by taking the first k bits from the first parent and the remaining bits from the second parent. Add all these children to A_{i+1} .
 - (c) *Mutation*: choose an integer $0 < Q < |A_i| - S - P$. Select Q individuals from A_i with replacement using fitness rank selection. For each of these individuals, mutate each bit with a specified probability and add the resulting individual to A_{i+1} .
 - (d) If at this point $|A_{i+1}| < |A_i|$, complete the new population with random genomes like for A_0 .
3. Repeat Step 2 until one of the following conditions is satisfied:
 - the increase in maximum fitness from A_i to A_{i+1} is below a specified threshold;
 - the maximum fitness of an individual in A_{i+1} is above a specified threshold;
 - $i + 1$ reaches a specified upper bound.

Fig. 4.2.1: The steps of the GA used in this project. The first stopping condition in Step 3 was in fact disabled in all the simulations by setting the threshold to 0.

REFERENCES

- B. Alberts. *Molecular Biology of the Cell*. Garland Science, Taylor and Francis Group, 2015. ISBN 9780815345244. URL <https://books.google.co.uk/books?id=tLSuoQEACAAJ>.
- Hector Aldaz, Luke M Rice, Tim Stearns, and David A Agard. Insights into microtubule nucleation from the crystal structure of human γ -tubulin. *Nature*, 435(7041):523–527, 2005.
- Julio M Belmonte, Maria Leptin, and François Nédélec. A theory that predicts behaviors of disordered cytoskeletal networks. *Molecular systems biology*, 13(9):941, 2017.
- Christopher S Campbell and R Dyché Mullins. In vivo visualization of type ii plasmid segregation: bacterial actin filaments pushing plasmids. *The Journal of cell biology*, 179(5):1059–1066, 2007.
- Alessandro De Simone, François Nédélec, and Pierre Gönczy. Dynein transmits polarized actomyosin cortical flows to promote centrosome separation. *Cell Reports*, 14(9):2250–2262, 2016. ISSN 2211-1247. doi: <https://doi.org/10.1016/j.celrep.2016.01.077>. URL <https://www.sciencedirect.com/science/article/pii/S2211124716300936>.
- Marileen Dogterom and Bernard Yurke. Measurement of the force-velocity relation for growing microtubules. *Science*, 278(5339):856–860, 1997. doi: 10.1126/science.278.5339.856. URL <https://www.science.org/doi/abs/10.1126/science.278.5339.856>.
- Colby P. Fees and Jeffrey K. Moore. A unified model for microtubule rescue. *Molecular Biology of the Cell*, 30(6):753–765, 2019. doi: 10.1091/mbc.E18-08-0541. URL <https://doi.org/10.1091/mbc.E18-08-0541>. PMID: 30672721.
- Matthew J Footer, Jacob WJ Kerssemakers, Julie A Theriot, and Marileen Dogterom. Direct mea-

- surement of force generation by actin filament polymerization using an optical trap. *Proceedings of the National Academy of Sciences*, 104(7):2181–2186, 2007.
- Niels Galjart. Plus-end-tracking proteins and their interactions at microtubule ends. *Current Biology*, 20(12):R528–R537, 2010.
- Ethan C Garner, Christopher S Campbell, and R Dyche Mullins. Dynamic instability in a dna-segregating prokaryotic actin homolog. *Science*, 306(5698):1021–1025, 2004.
- F Gittes, B Mickey, J Nettleton, and J Howard. Flexural rigidity of microtubules and actin filaments measured from thermal fluctuations in shape. *Journal of Cell Biology*, 120(4):923–934, 02 1993. ISSN 0021-9525. doi: 10.1083/jcb.120.4.923. URL <https://doi.org/10.1083/jcb.120.4.923>.
- Leland H Hartwell, Joseph Culotti, John R Pringle, and Brian J Reid. Genetic control of the cell division cycle in yeast: A model to account for the order of cell cycle events is deduced from the phenotypes of yeast mutants. *Science*, 183(4120):46–51, 1974.
- T E Holy and S Leibler. Dynamic instability of microtubules as an efficient way to search in space. *Proceedings of the National Academy of Sciences*, 91(12):5682–5685, 1994. doi: 10.1073/pnas.91.12.5682. URL <https://www.pnas.org/doi/abs/10.1073/pnas.91.12.5682>.
- Hirokazu Hotani and Tetsuya Horio. Dynamics of microtubules visualized by darkfield microscopy: treadmilling and dynamic instability. *Cell motility and the cytoskeleton*, 10(1-2): 229–236, 1988.
- J Howard, AJ Hudspeth, and RD Vale. Movement of microtubules by single kinesin molecules. *Nature*, 342(6246):154–158, 1989.
- Jacob WJ Kerssemakers, E Laura Munteanu, Liedewij Laan, Tim L Noetzel, Marcel E Janson, and Marileen Dogterom. Assembly dynamics of microtubules at molecular resolution. *Nature*, 442(7103):709–712, 2006.
- Masahide Kikkawa. Big steps toward understanding dynein. *Journal of Cell Biology*, 202(1): 15–23, 2013.

- Chapin Korosec and Nancy Forde. Engineering nanoscale biological molecular motors. *Physics in Canada*, 73:78–81, 04 2017.
- Akihisa Matsuyama, Ritsuko Arai, Yoko Yashiroda, Atsuko Shirai, Ayako Kamata, Shigeo Sekido, Yumiko Kobayashi, Atsushi Hashimoto, Makiko Hamamoto, Yasushi Hiraoka, et al. Orfome cloning and global analysis of protein localization in the fission yeast *Schizosaccharomyces pombe*. *Nature biotechnology*, 24(7):841–847, 2006.
- M. Mitchell. *An Introduction to Genetic Algorithms*. Complex Adaptive Systems. MIT Press, 1998. ISBN 9780262631853. URL <https://books.google.co.uk/books?id=gjoiEAAAQBAJ>.
- D.O. Morgan. *The Cell Cycle: Principles of Control*. Primers in biology. New Science Press, 2007. ISBN 9780878935086. URL https://books.google.co.uk/books?id=_7ygQAOK1DUC.
- Francois Nédélec and Dietrich Foethke. Collective Langevin dynamics of flexible cytoskeletal fibers. *New Journal of Physics*, 9(11):427, 2007.
- Raja Paul, Roy Wollman, William T. Silkworth, Isaac K. Nardi, Daniela Cimini, and Alex Mogilner. Computer simulations predict that chromosome movements and rotations accelerate mitotic spindle assembly without compromising accuracy. *Proceedings of the National Academy of Sciences*, 106(37):15708–15713, 2009. doi: 10.1073/pnas.0908261106. URL <https://www.pnas.org/doi/abs/10.1073/pnas.0908261106>.
- Julien Pernier, Shashank Shekhar, Antoine Jegou, Bérengère Guichard, and Marie-France Carlier. Profilin interaction with actin filament barbed end controls dynamic instability, capping, branching, and motility. *Developmental cell*, 36(2):201–214, 2016.
- Beat Rupp and François Nédélec. Patterns of molecular motors that guide and sort filaments. *Lab on a chip*, 12(22):4903–4910, 2012.
- Jeanne Salje and Jan Löwe. Bacterial actin: architecture of the *parMRC* plasmid DNA partitioning complex. *The EMBO Journal*, 27(16):2230–2238, 2008.
- Virginie Stoppin, Marylin Vantard, Anne-Catherine Schmit, and Anne-Marie Lambert. Isolated plant nuclei nucleate microtubule assembly: The nuclear surface in higher plants has centrosome-like activity. *The Plant Cell*, 6(8):1099–1106, 1994.

- Ronald D Vale and Ronald A Milligan. The way things move: looking under the hood of molecular motor proteins. *Science*, 288(5463):88–95, 2000.
- Kristen J Verhey, Donna L Lizotte, Tatiana Abramson, Linda Barenboim, Bruce J Schnapp, and Tom A Rapoport. Light chain–dependent regulation of kinesin’s interaction with microtubules. *The Journal of cell biology*, 143(4):1053–1066, 1998.
- Kristen J Verhey, Neha Kaul, and Virupakshi Soppina. Kinesin assembly and movement in cells. *Annual review of biophysics*, 40:267–288, 2011.
- Sam Walcott. The load dependence of rate constants. *The Journal of chemical physics*, 128(21):06B601, 2008.
- Claire E Walczak, Shang Cai, and Alexey Khodjakov. Mechanisms of chromosome behaviour during mitosis. *Nature reviews Molecular cell biology*, 11(2):91–102, 2010.

A Code Snippets

A.1 Modifications to Cytosim

Below are some of the significant blocks of code added to Cytosim.

Additional energy constraints

`simul.cc`

```
33 :
34 /// The amount of energy the cell has, in units of ATP equivalent
35 real cash;
36
37 /// Checks if the cell has enough 'cash' to perform the action
38 /// - if so, it returns true and subtracts the passed amount
39 bool checkWealth(real cost)
40 {
41     if (cash >= cost)
42     {
43         cash -= cost;
44         return true;
45     }
46     else
47     {
48         return false;
49     }
50 }
51 :
```

couple_set.cc

```
248 :
249 void CoupleSet::newObjects(ObjectList& res, const std::string& name, Glossary&
    opt)
250 {
251     CoupleProp * p = simul_.findProperty<CoupleProp>("couple", name);
252     Couple * obj = p->newCouple(&opt);
253
254     if ((p->synthesis_cost > 0) && !checkWealth(p->synthesis_cost))
255     {
256         return;
257     }
258 :
```

fiber_site.h

```
110 :
111     /// move to a different abscissa on the current fiber
112     void moveTo(real a)
113     {
114         hAbs = a;
115         reinterpolate();
116     }
117
118     /// moveTo but taking into account energy cost
119     void moveTo(real a, real movement_cost)
120     {
121         if ((a * movement_cost > 0) && !checkWealth(a * movement_cost))
122         {
123             return;
124         }
125
126         moveTo(a);
```

```

127     }
128     :

```

nucleator.cc

```

24     :
25 void Nucleator::makeFiber(ObjectList& objs, Simul& sim, Vector pos, std::
      string const& fiber_type, Glossary& opt)
26 {
27     ObjectMark mk = 0;
28     Rotation rot(0, 1);
29
30     Fiber * fib = sim.fibers.newFiber(objs, fiber_type, opt);
31     Hand const* h = hMonitor->otherHand(this);
32
33     real cost = fib->length() * fib->prop->unit_energy_cost;
34
35     if ((cost > 0) && !checkWealth(cost))
36     {
37         objs.destroy();
38         return;
39     }
40     :

```

fiber.cc

```

163     :
164     else if ( inc > 0 )
165     {
166         real cost = inc * prop->unit_energy_cost;
167
168         if ((prop->unit_energy_cost > 0) && !checkWealth(cost))

```

```

169     {
170         return;
171     }
172     :

```

Partitioning checkpoint

simul_step.cc

```

108     :
109     void Simul::step()
110     {
111         #if SPATIAL
112             // See if the beads have separated enough - if so, end the simulation
113             Bead* dna1 = beads.first();
114             Bead* dna2 = dna1->next();
115             real dna1_x = dna1->position().x();
116             real dna2_x = dna2->position().x();
117
118             if ((dna1_x * dna2_x < 0)
119                 && std::fabs(dna1_x) > 0.5
120                 && std::fabs(dna2_x) > 0.5)
121             {
122                 printf("%% dna1_x: %f\tdna2_x: %f\n", dna1_x, dna2_x);
123                 printf("%% time %f\n", prop.time);
124                 abortRun = 1;
125             }
126         #endif
127     :

```

A.2 Fitness functions

arena.py

```
1 try:
2     import os
3     import sys
4     import re
5     import numpy as np
6     from typing import List
7 except ImportError as e:
8     sys.stderr.write("Error loading module: %s\n" % str(e))
9     sys.exit()
10
11 spatial = False # set to True if including partitioning checkpoint
12
13 # simulation executable must be in current working directory:
14 simex = os.path.abspath('sim_spatial') if spatial else os.path.abspath('sim')
15 template = 'config.cym.tpl'
16 genetic_config = 'genetics.config'
17 max_time = 100 # units: s
18
19
20 def update_max_time():
21     global max_time
22
23     with open(template) as f:
24         for line in f:
25             if line.strip().startswith('time_step'):
26                 time_step = float(line.split()[2])
27             elif line.strip().startswith('run'):
28                 max_time = float(line.split()[1]) * time_step
29                 break
30
31
32 def calculate_fitness(data: List[str], target):
```

```
33     """
34     Calculate fitness expressing how close `data` is to `target`
35     Higher fitness is better
36     """
37     data = data[0].split('\n')
38
39     part_scores = []
40
41     if spatial:
42         for i, line in enumerate(data):
43             if line.startswith('% report'):
44                 # Figure out if the partitioning was successful
45                 if i >= 2 and data[i-2].startswith('% time'):
46                     # Partitioning successful
47                     t = float(data[i-2].split()[2])
48                     score_this_time = 1 - t/max_time
49                     # Account for floating point errors
50                     score_this_time = np.round(score_this_time, decimals=5)
51                 else:
52                     # Partitioning unsuccessful
53                     score_this_time = 0
54
55                 part_scores.append(score_this_time)
56     else:
57         # Remove empty lines
58         data_copy = data.copy()
59         for line in data:
60             if not line:
61                 data_copy.remove(line)
62         data = data_copy
63
64         for i in range(2, len(data), 4):
65             dna1 = data[i]
66             dna2 = data[i+1]
67             dna1_x = float(dna1.split()[2])
68             dna2_x = float(dna2.split()[2])
```

```
69
70     score_this_time = (-1 * np.sign(dna1_x) * np.sign(dna2_x) + 1) / 2
71     # Account for floating point errors
72     score_this_time = np.round(score_this_time, decimals=5)
73     part_scores.append(score_this_time)
74
75 part_score = np.round(np.mean(part_scores), decimals=5)
76 return part_score
```

A.3 Heatmaps and fitness plots

make_fitness_plot.py

```
1 #!/usr/bin/env python
2 # M. Yin 2022
3
4 """
5 Reads the run data from a specified directory and generates a heatmap
6 based on the distribution of fitness scores for each generation of the run.
7
8 Syntax:
9
10     make_fitness_plot.py [directory containing run data]
11
12     For correct naming and labelling, do not end the directory path with a
13     '/'.
14
15 Example:
16
17     make_fitness_plot.py run0
18
19 """
20 import os
21 import sys
```

```

21 import matplotlib.pyplot as plt
22 import numpy as np
23
24
25 spatial = True
26
27
28 # Function for generating heatmap
29 def heatmap(data, row_labels, col_labels, ax=None,
30             cbar_kw={}, cbarlabel="", **kwargs):
31     """
32     Create a heatmap from a numpy array and two lists of labels.
33
34     Parameters
35     -----
36     data
37         A 2D numpy array of shape (M, N).
38     row_labels
39         A list or array of length M with the labels for the rows.
40     col_labels
41         A list or array of length N with the labels for the columns.
42     ax
43         A `matplotlib.axes.Axes` instance to which the heatmap is plotted.  If
44         not provided, use current axes or create a new one.  Optional.
45     cbar_kw
46         A dictionary with arguments to `matplotlib.figure.colorbar`.  Optional
47     .
48     cbarlabel
49         The label for the colorbar.  Optional.
50     **kwargs
51         All other arguments are forwarded to `imshow`.
52     """
53
54     if not ax:
55         ax = plt.gca()

```

```

56     # Plot the heatmap
57     im = ax.imshow(data, **kwargs)
58
59     # Create colorbar
60     cbar = ax.figure.colorbar(im, ax=ax, **cbar_kw)
61     cbar.ax.set_ylabel(cbarlabel, rotation=-90, va="bottom")
62
63     # Show all ticks and label them with the respective list entries.
64     ax.set_xticks(np.arange(data.shape[1]), labels=col_labels)
65     ax.set_yticks(np.arange(data.shape[0]), labels=row_labels)
66
67     # Let the horizontal axes labeling appear on top.
68     ax.tick_params(top=False, bottom=True,
69                   labeltop=False, labelbottom=True)
70
71     # Rotate the tick labels and set their alignment.
72     plt.setp(ax.get_xticklabels(), rotation=45, rotation_mode="anchor",
73             size=7)
74
75     # Turn spines off and create white grid.
76     ax.spines[:].set_visible(False)
77
78     ax.set_xticks(np.arange(data.shape[1]+1)-.5, minor=True)
79     ax.set_yticks(np.arange(data.shape[0]+1)-.5, minor=True)
80     ax.tick_params(which="minor", bottom=False, left=False)
81
82     return im, cbar
83
84
85 def main(dir: str):
86     """
87     Generates a heatmap based on the distribution of fitness scores for each
88     generation of the run
89
90     Args:
91         dir (str): the directory where all the run data is saved

```



```

92     """
93
94     if not dir:
95         dir = os.getcwd()
96     elif dir.endswith('/'):
97         dir = dir[:-1]
98     dir = os.path.abspath(dir)
99
100    # Get required parameters from evolve.config
101    pam = {}
102    with open(dir + '/evolve.config', 'r') as f:
103        for s in f:
104            # Remove commented out lines
105            if not s.startswith('%'):
106                exec(s, {}, pam)
107    pop_size = pam.pop('population_size')
108    gen_max = pam.pop('generation_max')
109    config_file = pam.pop('config_file')
110
111    # The number of times each simulation is run, as specified by the last
112    # line in the config.cym.tpl file
113    # e.g. 'restart 9' --> num_reps = 10
114    with open(dir + '/' + config_file, 'r') as f:
115        for line in f:
116            if line.startswith('restart'):
117                num_reps = int(line.split()[1]) + 1
118                break
119
120    # Extract all the data we need
121
122    x = [] # The generation numbers
123    y = [] # The fitness scores
124    averages = [] # Average fitness for each generation
125    max_fitnesses = [] # Maximum fitness for each generation
126
127    for gen in range(gen_max):

```

```

128     try:
129         with open(f'{dir}/gen{gen:04}/generation.txt') as file:
130             scores_for_gen = []
131
132             for l in file:
133                 if l == '':
134                     break
135                 # Add fitness score to list
136                 score = float(l.split()[1])
137                 y.append(score)
138                 scores_for_gen.append(score)
139     except FileNotFoundError:
140         print(f'Found {gen} gen* directories')
141         gen_max = gen
142         break
143
144     averages.append(np.round(np.mean(scores_for_gen), decimals=5))
145     max_fitnesses.append(np.round(max(scores_for_gen), decimals=5))
146     x += [gen] * pop_size
147
148     averages = np.array(averages)
149     max_fitnesses = np.array(max_fitnesses)
150
151     fig = plt.figure(figsize=(13, 5))
152     ax = fig.subplots()
153
154     # Count the number of occurrences of each fitness score
155     num_rows = 21 if spatial else (num_reps + 1)
156     upper = 0.02 if spatial else 1
157
158     counts = np.zeros((num_rows, gen_max))
159     fitness_scores = np.round(np.linspace(0, upper, num_rows), decimals=5)
160     for gen in range(gen_max):
161         gen_counts, _ = np.histogram(y[gen*pop_size:(gen+1)*pop_size],
162                                     fitness_scores)
163         for i, cnt in enumerate(gen_counts):

```

```
164         counts[-(i+1), gen] = cnt
165     fitness_scores = fitness_scores[:, :-1]
166
167     ax.set_title(f'Heatmap of generational fitnesses for '
168                f'{os.path.basename(dir)}\n')
169     heatmap(counts, fitness_scores, range(gen_max), ax,
170            cbarlabel='Count', cmap='viridis')
171     ax.set_xlabel('Generation')
172     ax.set_ylabel('Fitness')
173
174     # Draw average line
175     ax.plot(range(gen_max), (upper - averages) * (num_rows - 1), 'w-')
176     # Draw max fitness line
177     ax.plot(range(gen_max), (upper - max_fitnesses) * (num_rows - 1), 'r-')
178
179     os.chdir(dir)
180     filename = f'{os.path.basename(dir)}_heatmap.png'
181     plt.savefig(filename)
182     print(f'Image saved as {filename}')
183
184
185 if __name__ == '__main__':
186     if len(sys.argv) > 1 and sys.argv[1].endswith('help'):
187         print(__doc__)
188     else:
189         main(sys.argv[1])
```

B Configuration Files

The following files were used for the GA runs in Subsection 2.2 with the first evolutionary configuration. The other configuration files were similar.

Cell template file

`config.cym.tpl`

```
1 % Simulating partitioning
2
3 set simul system
4 {
5     time_step = 0.001
6     viscosity = 0.1
7     display = (label=(Partitioning simulation); size=1200, 700)
8     income = 800000
9 }
10
11 set space cell
12 {
13     shape = capsule
14 }
15
16 new cell
17 {
18     radius = 0.4
19     length = 2
20 }
21
22 %%% Microtubules
23 set fiber microtubule
24 {
25     rigidity = 20
26     segmentation = 0.2
```

APPENDIX B. CONFIGURATION FILES

```
27   confine = inside, 100
28   unit_energy_cost = 1625 % 13/0.008
29
30   activity = dynamic
31   unit_length = 0.008
32   growing_speed = [[growing_speed]]
33   shrinking_speed = -0.5
34   hydrolysis_rate = 1.0 % catastrophe_rate approximately 1/8.3 s-1
35   growing_force = 1.7
36   total_polymer = 200
37 }
38
39 %% Motors
40 set hand kinesin
41 {
42   binding_rate = 10
43   binding_range = [[binding_range]]
44   unbinding_rate = 0.3
45   unbinding_force = 2.5
46
47   activity = move
48   unloaded_speed = 0.02 % mitotic kinesin very slow
49   stall_force = 6
50   movement_cost = 125 % 1/0.008
51
52   display = (color=red, red;)
53 }
54
55 set couple complex
56 {
57   hand1 = kinesin
58   hand2 = kinesin
59   stiffness = 100
60   diffusion = 0.05
61   specificity = antiparallel
62   length = 0.025
```

APPENDIX B. CONFIGURATION FILES

```
63     fast_diffusion = 0
64     synthesis_cost = 400
65 }
66
67 %%% DNA (plasmids)
68 set bead dna
69 {
70     confine = inside, 100
71     display = (color=green;)
72 }
73
74 %%% Nucleator
75 set hand nucleator
76 {
77     unbinding = 0, inf % rate, force
78     activity = nucleate
79     nucleate = 1, microtubule, (length=0.01; plus_end=grow;)
80     % rate, name of filament, new microtubule
81     display = (color=gray;)
82 }
83
84 set single grower
85 {
86     hand = nucleator
87     stiffness = 1000
88 }
89
90 %%% Create objects
91 new dna
92 {
93     radius = 0.1
94     attach = [[int(growers_per_dna)]] grower
95     position = -0.5 0 0
96 }
97
98 new dna
```

```

99 {
100     radius = 0.1
101     attach = [[int(growers_per_dna)]] grower
102     position = -0.3 0 0
103 }
104
105 new event
106 {
107     activity = (new 1 complex)
108     rate = [[motor_syn_rate]]
109 }
110
111 run 100000 system
112 {
113     nb_frames = 4
114 }
115
116 report dna:position *
117 restart 39

```

Evolutionary configuration

evolve.config

```

1 config_file = "config.cym.tpl" # name of the config template file
2 population_size = 50           # size of the population
3 njobs = 32                     # number of processes
4 FIT_DIFF_MAX = 0               # convergence criteria on fitness difference
    between two successive generations
5 FIT_MAX = 0.9                  # convergence on absolute fitness value
6 elitism = 0.1                  # fraction of population that continues to new
    generation
7 crossover = 0.7                # fraction of population that is created from
    mating of two parents

```

```
8 mutation = 0.2                # fraction of population that is created from
    mutation of parents
9 mutation_rate = 0.04          # rate of change per bit of genome
10 n_plus = 1.3                 # for rank selection, expected number of
    offspring for best individual
11 resurrect = 1                 # reuse data from old generation when one
    genome reappears
12 generation_max = 50          # max number of generations to be calculated
```

Parameter ranges

genetics.config

```
1 growing_speed = ([0.05, 0.5], 8)
2 binding_range = ([0.01, 0.05], 4)
3 growers_per_dna = ([1, 12], 4)
4 motor_syn_rate = ([10, 100], 6)
```
