

STEPHAN MULLER

Background

It still surprises me that my Intel project was in the field of biology, because I am not a biologist. I've always thought of myself as a physicist, mathematician or even chemist, dealing in sciences governed by ideas and concepts rather than facts that must be memorized. However, I found myself researching a project combining biology and mathematics. I wish I could say that summer of peering at cells through a microscope in addition to algorithms on a computer screen rekindled (kindled for the first time?) my love of biology. Unfortunately, it didn't. However, the story isn't entirely without a moral. I did learn that my first scientific love, computer science, could be applied to something as seemingly unrelated as the study of human disease. Even if my struggles to properly focus a fluorescence microscope or my clumsiness with keeping tools sterile make me wary of setting foot in a lab in the future, I may collaborate with other scientists on biological applications of math in the future, or on entirely different and entirely unforeseen applications of mathematics or physics. Thus, my advice to high school researchers considering a project in mathematics would be to consider studying an application of it in another branch of science, even one with which you may have a more tenuous relationship. Use your math, physics, engineering, computer science, etc., skills to help solve problems in biology, environmental science, or any number of other fields sorely lacking in quantitative analysis. It likely won't be your last chance to work with applications of more technical sciences, but it may be your last chance to work in the lab or in the field yourself, which I can assure you is more enjoyable at a university than in a high school classroom.

This is all well and good in retrospect, but I suppose I still owe an explanation of how I started researching this unexpected application of computers in the first place. The truth is that the lab was easily available (it was my father's), and the application interested me. I worked at Weill Cornell Medical College in New York City, in a lab that largely studies inflammation. The full importance of inflammation will be discussed later, but I was quickly convinced of its significance as a cause of, as well as treatment for, a wide variety of diseases and conditions. In fact, I was so convinced of this that I was rendered inarticulate by the gravity of the monumental implications while discussing them at a practice interview during the Intel finalist Media Training session, in what has become my sole uncomfortable memory of the entire competition (The sentence "My topic is important because inflammation is an important topic" made sense to me at the time.)

Because of my computer programming experience, I was able to devise my own project relating to and helping the lab's research. Many researchers in the lab had been working on new ways to experimentally study a membrane structure in endothelial cells (the cells that line the walls of blood vessels) so small that it could not be viewed directly with a light microscope. However, nobody had yet made a computer model of this structure, which, if made to replicate the known behavior of the compartment, could be used to study it more accurately than current experimental techniques. This is where I was able to carve out my niche as a programmer in a biology lab.

For the computer segment of the project, I found a computer science mentor, also at Cornell, Dr. Jonathan Victor. Dr. Victor routinely uses computer models to aid his own research in neurology and neuroscience, and was willing to help me with this application. During a small number of hour-long meetings among Dr. Victor, my father and me, I laid out my idea for the

model, and together we discussed the basic structure of the design and made compromises between accuracy to the biological system and computational feasibility. After we filled the whiteboard in Dr. Victor's office with an eclectic combination of integral calculus, computer code and sketches of cells, I went home and programmed the model.

When I mentioned that computer science was my first scientific love, this was not merely a flowery turn of phrase. One of my earliest memories is bringing home my family's first Windows computer. I was two years old, but instantly determined to figure out how it worked and to make it work in new ways. My parents describe the early "programs" I made by manipulating images and text in Microsoft Paintbrush. My mom claims to still have some of these, although thankfully they're probably only compatible with Windows 3.1. I was making real programs from first grade, when I learned QBASIC. Soon, I started making gradually more complex computer games and selling them to my friends for their allowance money. I had learned some new languages and written larger programs over the intervening years, although I still had to spend the beginning of the summer learning MATLAB, Dr. Victor's preferred language. MATLAB isn't the most powerful or universal way of creating a program, but it turned out to be fairly simple to learn and use and is probably the best suited to a computer model like this one.

The first month of my research was spent sitting at home in front of my computer writing the code for the model. This involved some unexpected skills, including some spatial thinking about where line segments would move under certain types of transformations, and long hours pondering the piece of code of which I am most proud, which efficiently searches through a snapshot of the model and evaluates various parameters I wanted to track. The next month, I would take the 40 minute train ride, 10 minute subway ride and 12 block walk (or run, depending

on how soon the next train was) into the lab two or three days a week to do experiments to test some of the predictions that came out of the model. This involved much more learning than the programming. I learned various methods used for studying this structure experimentally, and how to design experiments to determine how it behaves under different conditions. After a couple days, I had the hang of manipulating the equipment quickly enough to drain the fluid from a coverslip filled with cells and replace it with other fluid before the cells dried out. After a month, I was even able to locate the cells under a fluorescence microscope without asking for help or spending three hours turning a focus knob in a dark room saying, “they must be around here somewhere.” I may not need these skills as a mathematician or physicist, but it was nonetheless an enjoyable experience, for the most part.

Spending half of a summer in a lab didn’t make me love biology, but I did realize that a willingness to combine a science with which I was familiar and comfortable and one with which I was neither comfortable nor familiar allowed me to contribute a unique service to a lab filled with people much more knowledgeable and better educated than I am. My perspective as a computer programmer gave me a project that could be almost entirely my own design and execution, but could add something useful to the lab’s research, and that was a very gratifying feeling. There are undoubtedly many other fields out there where collaboration with a mathematician or computer scientist could be beneficial, but has not been attempted yet. All an interested high school student would have to do to be guaranteed a unique, useful and interesting research project would be to find one.

Research

Inflammation is the process by which the body naturally repairs damaged tissue, but inflammation that is uncontrolled or occurring at the wrong time or in the wrong place leads to a number of diseases, including rheumatoid arthritis and atherosclerosis, a type of heart disease. Inflammation occurs when leukocytes (white blood cells) leave the blood stream by passing between endothelial cells, the cells that line the walls of blood vessels, and move into the surrounding tissue. Endothelial cells actively change shape to allow leukocytes to pass between them, but this requires an increase in surface area that would not normally be allowed by a cell membrane. To allow this shape change, extra membrane and molecules that assist in the migration of the leukocytes are released from a compartment in endothelial cells called the Lateral Border Recycling Compartment (LBRC). The membrane in this compartment is constantly recycling in and out of the cell. This recycling is important in inflammation, but is very difficult to study experimentally and much is unknown about it. Indirect experimental techniques can give a general overview of how it behaves, but it cannot be viewed under a light microscope and electron microscopes give only a cross section of a still image. The purpose of this research was to create a computer model based on the known behavior of the LBRC that could be used to study unknown features of the compartment more easily.

The LBRC is a network of interconnecting membrane structures that, based on electron microscope images, is thought to have a shape much like a cluster of grapes. Each “grape” is a spherical collection of membrane, called a vesicle, connected to other vesicles and to the surface of the cell. The compartment appears all along the junctions where two endothelial cells meet and where a white blood cell might squeeze between endothelial cells.

The first step was to determine the general features of the model to be used. The type of model chosen was a cellular automaton, which consists of a grid of cells or spaces. Each space

can have two or more states and can change state based on the states of spaces around it and a set of rules. Given simple guidelines and an initial pattern, cellular automata can create extremely complex, varied patterns over a large number of turns or generations. As a result, these models are often used for modeling complex biological systems and for other similar applications. The general format of a cellular automaton was modified slightly for this model, however. While a cellular automaton typically assigns states to the spaces themselves, this model assigned states to the four edges around each grid space. Each edge could contain membrane or no membrane. A third state was added in which the edge would contain a slightly longer segment of membrane with a slight fold or wrinkle, which could be expanded. The normal membrane segments are represented in black in figures, and the “wrinkled” segments are represented as red. Using this representation, a wide variety of conditions can be created. Single spaces or groups of spaces could be surrounded by membrane to form what is known as a vesicle. This vesicle could be floating in the cell or attached to the surface, which is always represented as a separate row of edges at the top of the visual representation of the model. If two spaces are part of the same




Figure 0 Basic structures that can be represented by the model. Black lines represent “smooth” membrane, red lines represent “wrinkled” membrane. **a.** Two vesicles side by side, but not connected. **b.** A vesicle that consists of two spaces fused together. It is near the surface of the cell but not fused. **c.** A vesicle opening onto the surface of the cell.

vesicle, there will be no membrane between them. If a vesicle is attached to the surface, there will be no surface membrane at the attachment point. Some of these structures are represented in Figure 1.

Next, a set of transitions, rules by which edges would change state, was added. There were initially six transitions: Move (vesicles move through the cell or around another vesicle if they are attached), Flow (a wrinkle flows to an adjacent edge), Bud (two adjacent wrinkles expand to form a bud off of the vesicle), Unbud (a bud collapses into two wrinkled edges), Fuse (two adjacent but not attached vesicles fuse together) and Unfuse (two fused spaces separate). Each transition could be assigned a “probability,” a value that would determine how often each transition would occur relative to the others.

For the initial simulations, the transitions were placed in the order of how likely each was thought to be within the actual cell. This order was fuse, bud, unfuse, unbud, move, flow, from least probable to most probable. Fusion of membrane is thought to require more energy than flowing and thus be more difficult and less likely to occur. A number of simulations were run with the transitions assigned probabilities in this order, but with different ratios of probabilities. The simulations were started with an initial condition similar to what a still image of the compartment is thought to look like based on electron microscope images, and simulations were judged on how well they were able to maintain this realistic structure. In the more successful initial simulations, the transitions were all fairly close together in probability. Simulations in which the transitions had very different probabilities were less successful.

After these initial simulations, more simulations were run using the more successful probabilities with a different initial condition. In these simulations, the model would start with all of the membrane present as wrinkled edges on the surface and no membrane inside the compartment. The hope was that the model could use the same rules to form the compartment from scratch. The simulations were fairly unsuccessful in doing this. At this point, the model was adjusted to be more realistic given experimental knowledge of the compartment. For example, it

was known experimentally that at any given time, approximately 70% of the membrane is on the surface and 30% is internal, while these simulations resulted in a much larger percentage of the membrane being internalized. To fix this, an algebraic analysis technique for cellular automata was used. In this technique, the current value of a certain quantity (in this case, percentage of membrane that is on the surface) is graphed against the most likely value of this quantity on the




Figure 2 Analysis for achieving the best ratio of Bud to Unbud probabilities. Along the x axis is the current percentage of membrane on the surface. Along the y axis is the percentage of membrane likely to be on the surface in the next generation. Where the graph crosses the dotted red line, this value is stable.

next generation of the model. To make this graph, past simulations were analyzed to determine how much membrane typically internalizes and externalizes when different percentages of membrane are present on the surface. Based on this, the most likely percentage on the next generation was calculated for each current percentage. The resulting graph is shown in Figure 2. A value is considered to be stable where the graph crosses a 45° upward diagonal line ($y=x$), thus where the value on the next generation is likely to be the same as the value on the current generation, and this value will continue to be the same in successive generations. The goal in this case was to make the percentage of membrane on the surface stable at 70%, the desired value. The relative probabilities of budding and unbudding were adjusted until the value became stable at 70%.

A second problem with the simulations was that vesicles tended to bud off from the surface and then fuse together, forming long, narrow “troughs” attached to and just below the

surface. These structures were not seen in electron microscope images. In another related problem, individual vesicles would unfuse from the surface and move into the cell while, based on experiments, it was known that the entire compartment remains connected to the surface. Both structures were created by fusion and unfusion of membrane, so the probabilities of these two transitions were decreased and the compartment was encouraged to expand through budding instead of fusion of smaller compartments.

A third problem was that the compartments being formed by the model were small and simple while the compartments seen experimentally are large and complex. An assumption was made that interactions between molecules on the surfaces of adjacent cells would inhibit membrane from budding off from the surface, so fewer buds would form but those that formed would become larger compartments. To implement this assumption, separate transitions were created for budding and unbudding from the surface, and these transitions were given lower probabilities than standard budding and unbudding.

Finally, transitions would not occur in exactly the ratios they were assigned, since some transitions were possible more often than others. For example, fuse requires a much more complex, and rarer, pattern of membrane (two adjacent segments of smooth membrane attached to other segments of smooth membrane) whereas flow requires a much simpler pattern (a wrinkled edge adjacent to a smooth one). Even if fuse and flow are given equal probabilities, flow will occur more often because there are more opportunities. To compensate for this, the model was programmed to force the transitions to occur in the assigned ratios by checking on each generation how often the transitions had occurred in the past. The probabilities of underrepresented transitions would be increased and overrepresented transitions would be temporarily stopped.

With these modifications, the model predicted the known behavior of the compartment fairly well. However, to further confirm the model it was necessary to test some predictions of the model for behavior that was not currently known. Two such predictions were tested experimentally using Human Umbilical Vein Endothelial Cells (HUVEC) grown in the laboratory. The model was able to use the set of transitions to create the compartment from membrane entirely on the surface, but it was not known that the actual cells could do this. If they could, the model predicted the kinetics with which this would occur. The endothelial cells were separated using ethylenediamine tetraacetic acid (EDTA), which disrupts the junctions between them. It was thought that this would bring the compartment membrane entirely to the surface. This was confirmed using a fluorescently tagged antibody which bound to surface molecules on

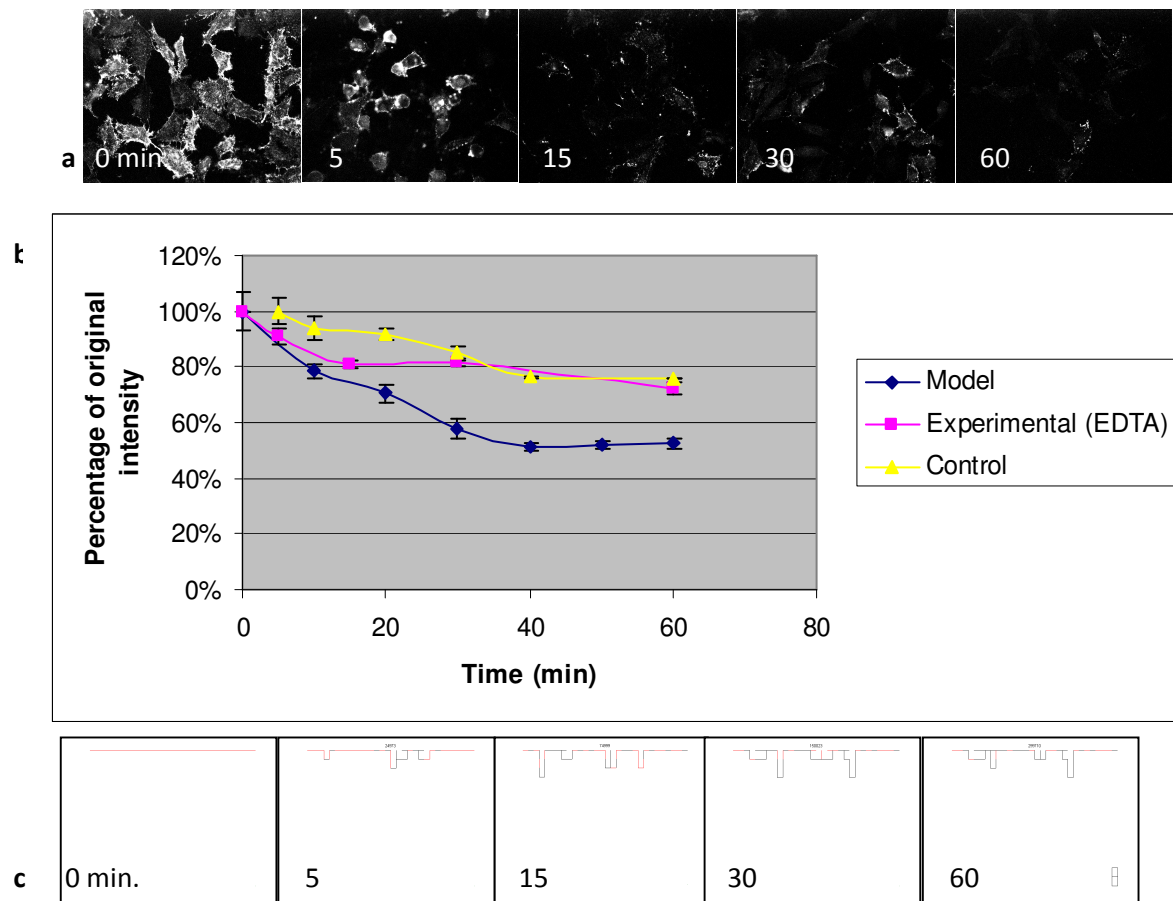


Figure 3 a. Photographs of experimental cells at various timepoints. **b.** The results of the re-formation of the compartment after junctional disengagement, along with control, starting with 70% of membrane in the compartment and the prediction of the computer model. The experimental results are the mean \pm standard error of 4 or more image fields from the same group of cells, and the model results are of 10 simulations. **c.** Graphic representations provided by the computer model at various timepoints.

any membrane that came to the surface during the treatment. Areas in which the cells had separated were bright when viewed under a fluorescence microscope, indicating membrane had come to the surface. Now the membrane was allowed to reinternalize. At various timepoints, membrane still on the surface was labeled fluorescently. Fluorescence levels at these timepoints could be quantified to determine what percentage of the original membrane remained on the surface. These results are shown in Figure 3, along with the prediction of the model for this situation. The model came close, but didn't exactly predict the behavior of the cells.

The next prediction tested was the assumption that surface interactions would inhibit budding from the surface. Since this assumption improved the model, if it could be confirmed experimentally, this would support the model. This prediction was tested by performing an internalization experiment similar to the one described above on confluent cells (cells that had come together and formed junctions with other endothelial cells) and subconfluent cells that had not yet come into contact with other endothelial cells. Subconfluent cells with no surface interactions showed a much greater amount of internalization, similar to the results obtained

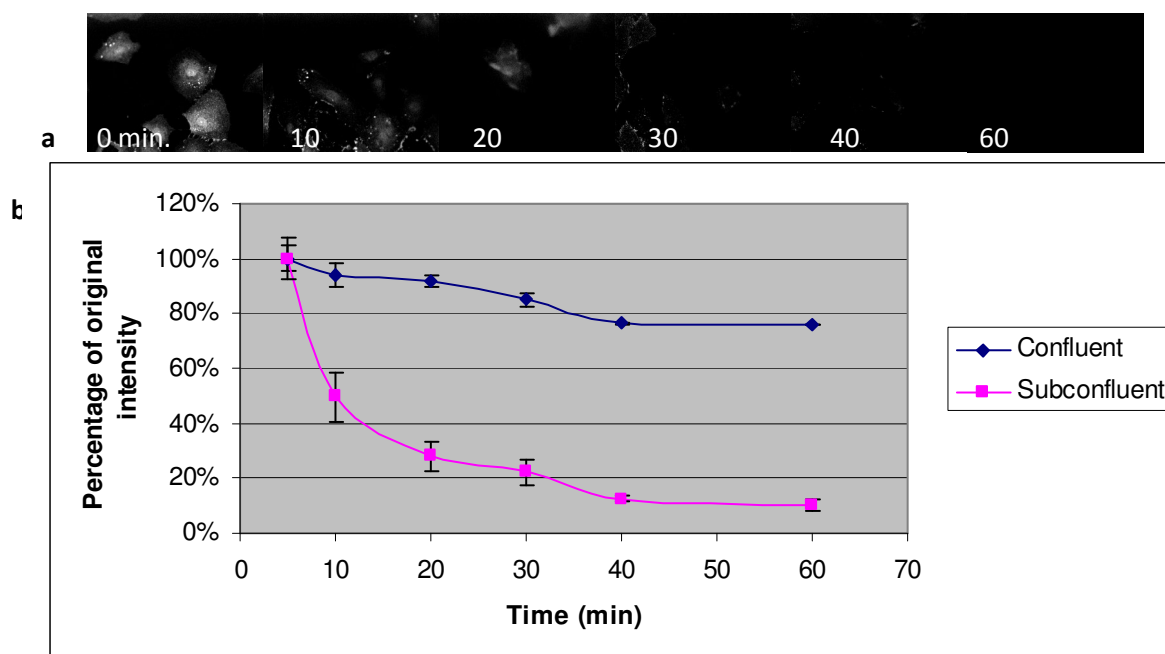


Figure 4 a. Photographs of subconfluent cells stained with fluorescent dye. The decrease in brightness indicates that the membrane labeled with P1.1 is moving inside the cell and cannot be labeled with the fluorescent antibody. **b.** The results of internalization in both confluent and subconfluent cells. Results are the mean \pm standard error of 5 or more image fields from the same group of cells.

from the model before this assumption was implemented. This indicated that surface interactions were important for the proper formation of this compartment, as predicted by the computer model. These results are shown in Figure 4.

After this paper was submitted to Intel, the model was revised to bring its predictions closer to the experimental results through several additional modifications. First, the code for adjusting the probabilities of transitions based on past occurrences was improved. Second, the probabilities of budding and unbudding from the surface were made dependent on the tightness of the junction. A junction with a great deal of wrinkled membrane would have weaker surface interactions than a smooth junction. Third, fusion and unfusion of membrane were removed as transitions. These events are not known definitively to occur, and removing them improved the model. The final results from the model and experiment are shown in Figure 5.




Figure 5 The modifications made to the model bring the predictions of the model much closer to the experimental results.

With the modifications, the model is now able to accurately reflect the known features of the compartment. It was also able to successfully predict two unknown features of the compartment, and has already revealed new findings about the LBRC. First, the compartment is able to re-form itself using the same mechanisms it uses to maintain itself. This was predicted by the model since the same transitions can maintain and form the compartment, and confirmed in

the cells. Second, fusion and unfusion are not necessary for the formation of these compartments, but are vital for other types of membrane recycling. This implies the LBRC may use a different mechanism from most membrane compartments and so it may be possible to block inflammation by preventing the LBRC from forming without affecting normal cell function. This model could be used in the future to predict the behavior of the compartment under a variety of other situations, replacing experiments which can be costly and require the use of human or animal tissue, and cannot currently detect some features of the compartment that can be studied using the model. The model could also be used to develop and test methods of inhibiting the formation of the LBRC to treat inflammatory diseases.

References

1. Gutowitz, H., Victor, J.D., and Knight, B.W. (1987). Local structure theory for cellular automata. *Physica D*, 28, 18-48
2. Mamdouh, Z, *et. al.* (2003). Targeted recycling of PECAM from endothelial surface-connected compartments during diapedesis. *Nature*, 42, 748-753
3. Muller, W.A. (2003). Leukocyte–endothelial-cell interactions in leukocyte transmigration and the inflammatory response. *Trends in Immunology*, 24(6), 326-333
4. Robbins, Stanley, *et. al.* (1974). *Pathologic Basis of Disease*. Philadelphia: W.B. Saunders Company
5. Weisstein, Eric W. (2006). *Life*. Retrieved June 2007, from *MathWorld--A Wolfram Web Resource*. Web Site: <http://mathworld.wolfram.com/Life.html>