

Simulating Kinetic Processes in Time and Space on a Lattice

J. P. Gill¹ *, K. M. Shaw¹, B. L. Rountree⁴,
C. E. Kehl¹, and H. J. Chiel^{1,2,3}

¹ Department of Biology,

² Department of Neurosciences, ³ Department of Biomedical Engineering,
Case Western Reserve University, Cleveland, OH, 44106

⁴ Center for Applied Scientific Computing,
Lawrence Livermore National Laboratory, Livermore, CA 94550

Abstract. We have developed a chemical kinetics simulation that can be used as both an educational and research tool. The simulator is designed as an accessible, open-source project that can be run on a laptop with a student-friendly interface. The application can potentially be scaled to run in parallel for large simulations. The simulation has been successfully used in a classroom setting for teaching basic electrochemical properties. We have shown that this can be used for simulating fundamental molecular and chemical processes and even simplified models of predator–prey interactions. By giving the simulated entities spatial extent in the lattice, the particles do not interpenetrate, and clusters of particles can spatially exclude one another. Our simulation demonstrates that spatial inhomogeneity leads to different results than those that are obtained by using standard ordinary differential equation models, as previously reported.

Key words: education in biomathematics, lattice models, artificial chemistry, simulation, open-source, autopoiesis, origin of life, diffusion, Nernst potential, resting potential, Donnan equilibrium, chemical reactions, kinetics, Michaelis–Menten, enzyme kinetics, Lotka–Volterra, predator–prey model

AMS subject classification: 97M60, 82C22, 92C45, 92D25, 92B05

*Corresponding author. E-mail: jpg18@cwru.edu

1. Introduction

Physical and biological models of natural phenomena that are taught in high schools and undergraduate programs make critical assumptions and simplifications that allow the models to be analyzed and understood in ways that would not otherwise be possible. However, the natural phenomena themselves are not constrained by these assumptions, and the effects on the model of deviating from these assumptions are rarely well understood. For instance, when studying the kinetics of a chemical system using ordinary differential equations (ODEs), it is usually assumed that the distributions of the reacting species are spatially homogeneous. Especially in biological systems, which often have small compartmentalized regions containing macromolecular complexes and small numbers of molecules whose spatial distribution is important for function, these assumptions are violated, and the predictions of kinetic models will deviate significantly [26].

For many introductory physics, chemistry, and biology courses, the spatial and stochastic nature of the real world is ignored, and homogeneous, continuous, and deterministic models are used for teaching basic concepts [12]. Significant deviations between a model and the natural phenomenon it describes cannot be understood until these assumptions are explored.

We present a lattice-based chemical simulation platform that can help students and researchers understand the importance of the spatial and stochastic properties of natural physical, chemical, and biological phenomena. As an educational tool, we have designed our simulator to be user-friendly and interactive. It is capable of running on a student's laptop, and clear visualization tools are included. The simulator can be used for studying basic chemical processes as well as higher-order population interactions. In this paper, we will present examples of simulations of diffusion, basic electrochemical phenomena, enzyme kinetics, and predator-prey interactions.

We created the simulation to use it as a research tool. The focus of our research is to understand how natural, abiotic processes resulted in the emergence of life. We wish to test the hypothesis that the origin of life was a probable, if not inevitable, consequence of forces in play on the prebiotic Earth, and these forces can be deduced from first principles [30]. Such an understanding could bridge the gap between the abiotic and cellular worlds, as well as fill in the earliest chapter of our own origins. We are developing the simulator with the intent of discovering how autopoietic behavior, i.e., behavior of biochemical networks that is self-maintaining [31], may have emerged spontaneously from the fundamental rules of physics and chemistry. Therefore, it has been important to ensure that the simulator remain scalable and computationally efficient. To permit a larger community to contribute to the development of this project, the code is also open-source (licensed under the GNU General Public License, version 3).

We place our work at a *mesoscale* level of granularity: more coarse-grained than molecular dynamics simulations such as CHARMM [4] and NAMD [25], yet finer-grained than cellular simulations such as E-Cell [28] and MCell [5]. Broadly speaking, simulations at this intermediate level of granularity are collectively categorized as *artificial chemistries* [7, 22]. Artificial chemistries are generally lattice-based, e.g., [15], or graph-based, e.g., [6, 27].

Nearly all of the work with artificial chemistries is focused on prebiotic evolution, and this topic informed our initial research as well. Our work differs from previous work in maintaining strong ties to classical chemical systems, allowing us to use these classical systems to validate

the simulation's behavior. As a consequence, we are much closer to the *ab initio* approaches of molecular dynamics than the high-level abstractions necessary for cellular simulation. This approach allows this tool to be far more general in scope and has already proven itself useful in classroom instruction in areas far removed from the topic of autopoiesis.

We further distinguish this work by noting that the design allows both individual student demonstration on modest laptops as well as scaling up to run far larger simulations on super-computer-class machines. While molecular dynamics simulations are written to run in parallel and cellular simulations are moving towards this architecture, we know of no other artificial chemistry simulation that is able to make use of multiple processors.

2. Application Details

Our model exists in two incarnations. The first application, referred to as the Nernst potential simulator, is an electrochemical simulator that can be used for studying the behavior of the electric potential difference across a cell membrane and the forces that drive the system as it approaches a steady state or equilibrium value. The second application is a chemical reaction simulator which is used for building reaction networks and studying their kinetics. Both applications use very similar underlying models.

The Nernst potential simulator has a complete graphical user interface that allows students to easily set up a simulation, run it, and immediately see results (see figure 1). The application presents a view of a compartmentalized space filled with ions. The user can set the initial concentrations of the ions on either side of a central membrane dividing the spaces, set the membrane permeability for each ion species, press "Start," and watch the system evolve. An electric potential will begin to form across the membrane, and its value is plotted. The user can easily explore the effects of each parameter on the behavior of the system. This is an incredibly useful tool for helping students of neurochemistry understand basic concepts.

The graphical user interface for the reaction simulator currently does not expose many features of the simulation engine, but we intend to extend this functionality with time. The user can design a custom network of reactions, populate the simulation with particle types of his or her choosing, and watch the system evolve. This is not only useful for studying chemical systems, but it can also be applied to cellular biology and ecology, as we shall demonstrate.

Both applications have fully functional command line interfaces that make working with them for batch projects or research purposes easy. The applications are open-source and are freely available for use. The Nernst potential simulator can be found at <http://www.github.com/jpg18/nernst>, and the reaction simulator can be found at <http://www.github.com/jpg18/metabolism>.

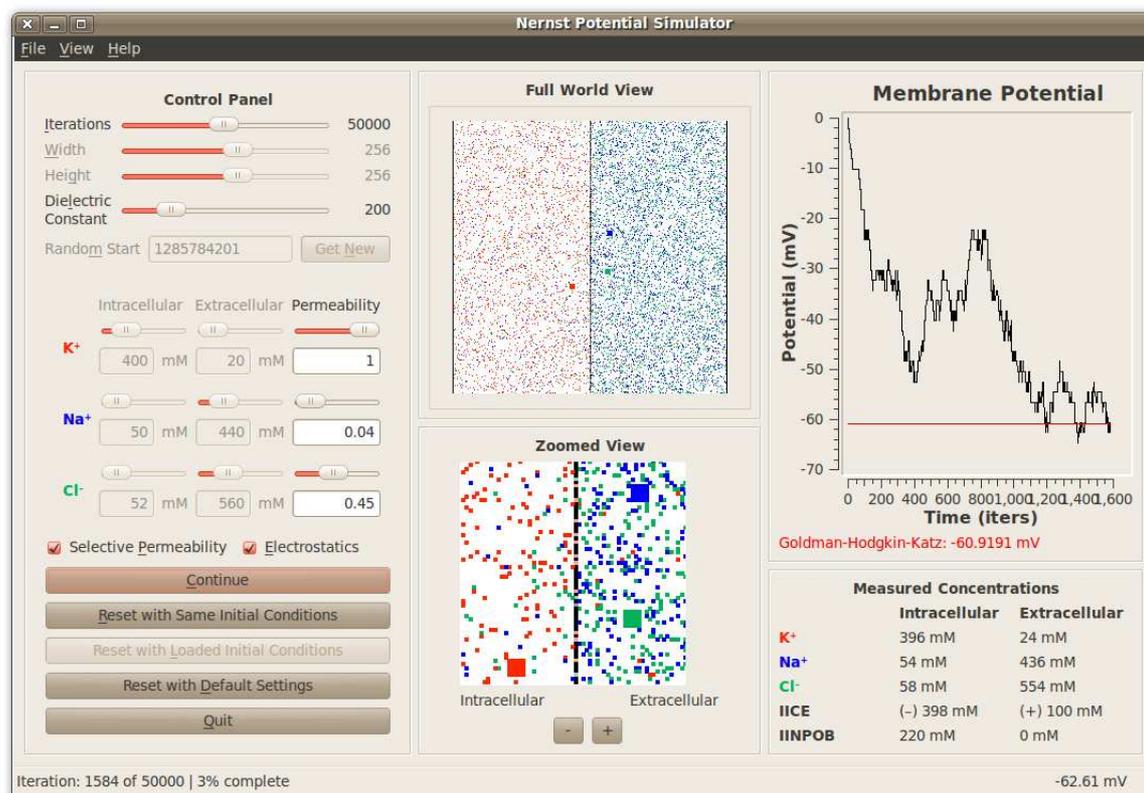


Figure 1: Screenshot of the Nernst potential simulator interface. The user can configure the initial concentration of ion species and other parameters at the control panel on the left. The lattice is displayed in the center in both a full view and a zoomed view. Individual ions can be tracked by clicking on them to highlight them (note the three highlighted ions visible in both view ports). A vertically-oriented membrane is located in the center of the lattice that is selectively-permeable to certain ion types. The potential difference across this membrane and its expected value are plotted on the right, and a table displays the concentrations of each species in real time. Included in the table are the concentrations of the implicitly modeled impermeable charges added to guarantee initial electroneutrality (IICE) and the implicitly modeled impermeable neutral particles added to guarantee initial osmotic balance (IINPOB).

3. Model Details

The underlying model shared by both applications is a lattice-based particle simulator. The “world” in which the particles live and interact is a two-dimensional square lattice. Lattice positions can be empty, but a lattice position cannot be occupied by more than one particle at one time. These particle entities each have a type, and they can represent individual ions, small organic molecules, cells, or even individuals.

The lattice allows the particle movement and interaction algorithms to be very simple and computationally efficient. To simulate the simultaneous movement of all particles without biasing

the movement of any particular particle, we devised the following algorithm: Time advances in the simulator in fixed, discrete steps. Every particle is given the opportunity to claim a nearby lattice position for movement. If more than one particle attempts to claim a single location, none of those particles is permitted to move in the current time-step since each position can only hold one particle. In a second pass through the lattice, particles that can move to their claimed positions are permitted to do so. See figure 2 for details. Note that other algorithms for lattice simulators may more effectively avoid concentration-dependent effects [23].

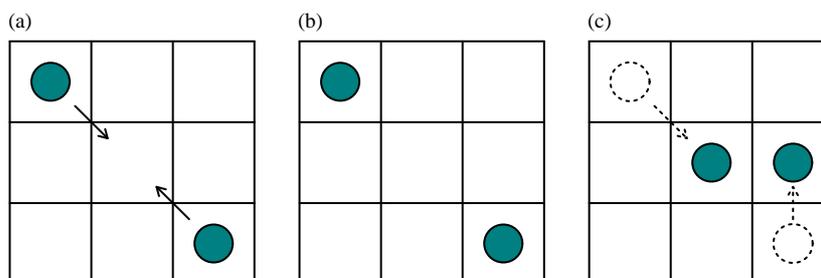


Figure 2: Particle movement algorithm in the simulator. (a) Step 1: Every particle in the lattice stakes a claim for a new position in an adjacent location once each iteration. (b) When a position is claimed by more than one particle, or if a claimed position is already occupied, all associated movement is cancelled. Here, no particles moved this iteration because both claimed the same position for movement. (c) Step 2: Particles move if collisions with other particles will not occur. Here, both particles moved because each claimed a different, vacant space. When particle density is low, collisions are rare, and each particle undergoes a diffusion process in the lattice.

The Nernst potential simulator is equipped with a special set of movement rules for handling ion flow across the membrane. The membrane divides the lattice in half, preventing the particles from moving across it by normal means. Ion-specific channels are embedded in the membrane, and flow across the membrane is possible for ions adjacent to a channel of the appropriate type. We treat ion flow through pores as fast, so that the ion may cross the pore several times in a single time-step. Under this approximation, the probability of finding the ion on a given side of the pore can be determined by a Boltzmann distribution using the relative difference in potential energy created by the membrane potential (see Appendix A). This probability is used for deciding where the particle will be located at the start of the next iteration. When an ion flows through the membrane, it is simply moved to the opposite side of the channel, which must be vacant. The membrane potential is calculated using the difference in charge between the two halves of the lattice and the capacitance of the membrane, which is constant throughout the simulation.

In addition to moving each iteration, particles in the chemical reaction simulator can undergo reactions that change their types. If the simulator is configured with zeroth-order reactions, then each empty lattice position has a constant probability each iteration of being filled with a new particle of the product species. The existing particles in the lattice each choose a neighboring lattice position at random and have a chance to react with the occupant of that position. See figure 3 for details. Appendix B includes a more in-depth look at the reaction algorithm.

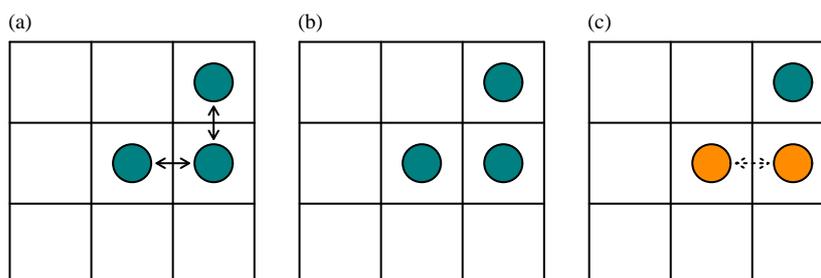


Figure 3: Chemical reaction algorithm in the simulator. (a) Step 1: Every particle in the lattice stakes a claim for a reaction partner once each iteration. First-order reactions are facilitated through self-selection or selection of an empty lattice space. (b) When a reaction partner is claimed by more than one particle, all associated reactions are cancelled. Here, no reactions occurred because the same reaction partner was selected by more than one particle. (c) Step 2: Reactions have a chance of occurring only if no other claims are made on the participating particles. When a reaction occurs, particle types change from those of the reactants to those of the products. Here, a reaction occurred because the reactants were not claimed by any other particles.

This simple model of how chemicals behave in the real world is sufficient to generate very interesting behavior, and it is especially useful for providing students and teachers alike with insight into what drives important physical, chemical, and biological phenomena from the ground up. Classical continuum approximations of diffusion and kinetics are very useful for understanding the behavior of many systems, but if the assumptions of those classical models are violated, then their predictions can deviate drastically from reality. Our model sheds some light on the consequences of violating these assumptions. We will now describe some of these important phenomena in detail.

4. Physical Simulations

4.1. Diffusion

The movement algorithm described above results in the diffusion of the particles in the lattice. In general, diffusion is the tendency of particles to move from areas of higher concentration to lower concentration. A diffusion process can occur when particles are suspended in a fluid in which they are regularly bombarded by thermally-excited solvent particles. Collisions are so frequent that the suspended particles move around randomly, changing direction very rapidly.

The diffusion process can be approximated well by a random walk assuming that collisions between particles are rare. A random walk is independent in each of the x -, y -, and z -dimensions, and the displacement in one dimension of an individual particle undergoing the diffusion process is normally distributed with mean $\mu = 0$ and standard deviation $\sigma = \sqrt{2Dt}$, where D is the diffusion coefficient, and t is the elapsed time of diffusion [8]. Thus, the standard deviation of the

displacements of a collection of diffusing particle grows with the square-root of time.

The diffusion coefficient D can be calculated for particles in our model. The Einstein–Smoluchowski equation states that the diffusion coefficient of a particle diffusing in one-dimension is

$$D = \frac{\lambda^2}{2\tau}, \quad (4.1)$$

where λ is the distance the particle diffused in time τ [8] [14, p. 313]. If we assume that collisions, which cause the particles in our model to fail to move for one time-step, occur so rarely that their effects on diffusion are negligible, then the particles in our model move a distance $\lambda = 1$ lattice position on average every $\tau = 4/3$ iterations, since movement of a particle has a displacement component in the x -dimension for 6 out of the 8 neighboring lattice, and the same is true for the y -dimension. Therefore, when collisions are rare, the diffusion coefficient for particles in our model is $D = 3/8$. Note that it is in this dilute limit that both transport diffusivity (bulk diffusion) and self-diffusion (diffusion of a single particle) are essentially equivalent.

Figure 4a shows the displacement in one dimension of a collection of particles freely diffusing in the simulator at low density ($\rho = 0.5\%$), reflecting self-diffusion. The diffusion coefficient for the model system was used to determine the expected distribution of the displacements. The sample data does not deviate significantly from expectation (see figure 4 for details).

However, when particle density is greater than 1%, significant deviations from the theoretical distribution become detectable (see figure 4b). The data remains normally distributed, and the mean remains indistinguishable from $\mu = 0$, but the standard deviation begins to shrink as the density increases. This is equivalent to a reduction of the self-diffusion coefficient at higher particle densities, and this is not surprising. Collisions, which in our model result in no movement for the colliding particles, are responsible for this reduced diffusion rate. What this means for the biological cell is that if the cytoplasm is very crowded, which is certainly the case [1, pp. 77-78], reactions which rely solely on the diffusion of chemicals to bring reactants together or to remove waste products will be very slow without the help of enzymes or active transport processes.

4.2. Nernst Equilibrium Potential

Normally, diffusion results in the destruction of any concentration gradients that are present in a system. However, if the diffusing particles are ions with electrical charges, the net movement of the ions down their concentration gradient will result in the build up of a charge differential. This charge imbalance can work to oppose the diffusion process under the right conditions, and an electrochemical equilibrium will form in which a concentration gradient persists and an electric potential is established. This is known as the Nernst equilibrium potential. For biological cells with a membrane permeable to an ion species X , a large difference in the concentration of the ion may exist between the intracellular and extracellular spaces due to the activity of ion pumps in the membrane. Passive leak channels will also exist in the membrane, and these allow the ions to diffuse down their concentration gradient until the electrochemical equilibrium is established. The Nernst equation gives the resulting equilibrium potential for ion X as

$$E_X = \frac{RT}{z_X F} \ln \frac{[X]_{out}}{[X]_{in}}. \quad (4.2)$$

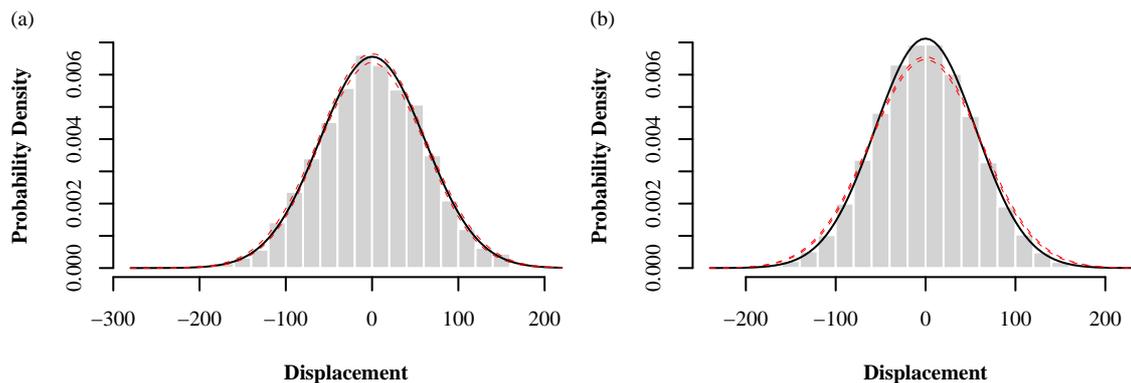


Figure 4: The distribution of the displacements in one dimension for inert particles undergoing a random walk (i.e., self-diffusion) in a large 2-D lattice. The distribution is shown after $t = 5000$ iterations of the simulator (approximately 3 minutes runtime for each simulation on a computer with a 2.1 GHz Intel Core 2 processor). (a) Particle density was low ($\rho = 0.5\%$, $n = 5243$). The solid curve is the normal probability density function fitted to the data. 95% confidence intervals were constructed for the sample mean and sample standard deviation of a hypothetical sample taken from an ideal population with parameters equal to the theoretical mean $\mu = 0$ and theoretical standard deviation $\sigma = \sqrt{2Dt}$. The dashed curves bound the region occupied by normal probability distribution functions whose parameterization satisfy simultaneously both confidence intervals on the mean and standard deviation. Because the sample mean and sample standard deviation of the actual simulated data falls within these confidence intervals (the solid curve lies between the dashed curves), the data does not deviate significantly from expectations. A Kolmogorov-Smirnov test also found that the empirical cumulative density function did not deviate significantly from the expected cumulative density function ($p = 0.511$). (b) Particle density was relatively high ($\rho = 5\%$, $n = 52429$). The sample standard deviation does not fall within the 95% confidence interval (the solid curve does not lie between the dashed curves), and the data does deviate significantly from expectations (Kolmogorov-Smirnov test: $p < 10^{-6}$). This is due to the increased frequency of collisions at higher particle densities.

where R is the molar gas constant, T is the temperature of the system, z_X is the valence of ion X , F is Faraday's constant, and $[X]_{out}$ and $[X]_{in}$ are the concentrations of ion X outside and inside the cell [1, p. 634].

Figure 5 plots the membrane potential measured in our model over the course of a simulation. The membrane potential rapidly drops down to the value predicted by equation (4.2) as the ions move down their concentration gradient. The potential remains at this value indefinitely because the system has reached an equilibrium.

The Nernst potential simulator generates the Nernst equilibrium potential by the unorchestrated diffusion of particles in the lattice. The membrane potential is calculated using an estimate of the membrane capacitance and the difference in charge between the two separated compartments. The plotted membrane potential is a result of nothing more than the simple diffusion process, the presence of an initial concentration gradient, and the simplified ion channel model whose

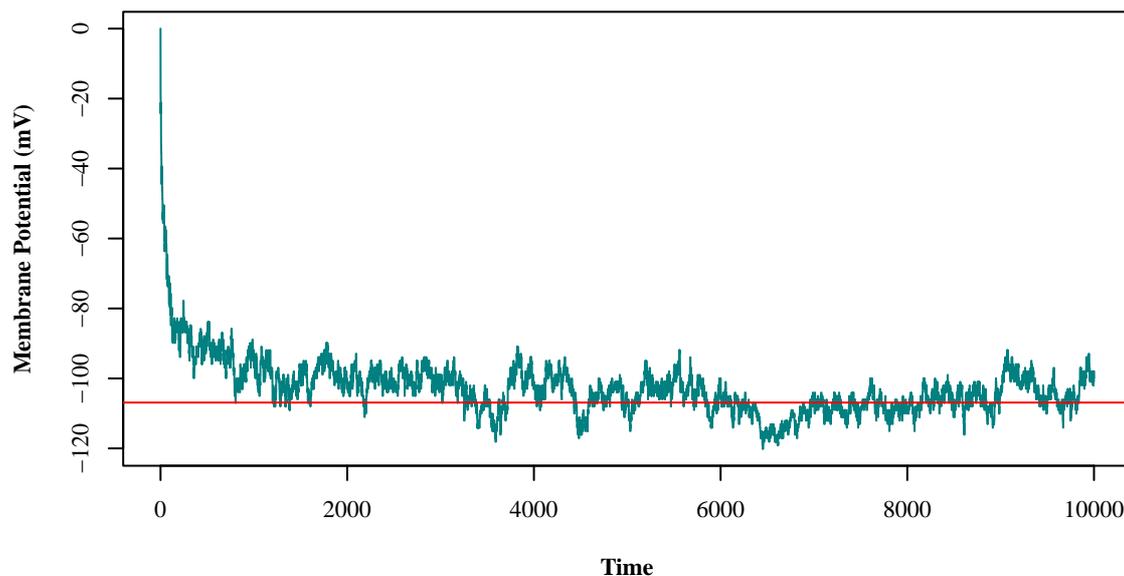


Figure 5: The evolution of the Nernst equilibrium potential across a semipermeable membrane. A simulation of potassium chloride in solution in two adjacent compartments separated by a membrane permeable only to potassium ions was initialized with a concentration ratio of 64:1. The total number of membrane-permeable potassium ions was $n = 56573$. As the potassium ions moved down their concentration gradient, an electric potential began to form across the membrane until an electrochemical equilibrium was established (approximately 8 minutes total runtime on a computer with a 2.1 GHz Intel Core 2 processor for this simulation). Equation (4.2) was used to calculate the Nernst equilibrium potential of potassium and successfully predict the equilibrium potential for the system, indicated by the horizontal line at $V_m = -107$ mV. A strong concentration gradient remained at the end of the simulation (approximately 58:1).

passive fluxes are dependent on potential differences across the membrane. What can be striking for students is that this phenomena arises entirely from the interaction of these simple physical processes using a relatively small number of ions (compared to the number of ions present in a large cell [17, p. 128]). The Nernst potential simulator helps students obtain a concrete understanding of this basic principle of neurochemistry.

4.3. Resting Potential and the Donnan Equilibrium

When multiple ion species are permeable to a cell membrane, the movement across the membrane of each contributes to the membrane potential. The combined effect of the concentration gradients of K^+ , Na^+ , and Cl^- ions on the membrane potential is given by the Goldman–Hodgkin–Katz equation,

$$V_m = \frac{RT}{F} \ln \frac{P_{K^+} [K^+]_{out} + P_{Na^+} [Na^+]_{out} + P_{Cl^-} [Cl^-]_{in}}{P_{K^+} [K^+]_{in} + P_{Na^+} [Na^+]_{in} + P_{Cl^-} [Cl^-]_{out}}, \quad (4.3)$$

where V_m is the membrane potential and P_X is the relative membrane permeability of ion X. Membrane permeability is controlled in the simulation independently for each ion species by changing the number of ion channels that are present in the membrane for that ion type. The resulting membrane potential is known more generally as the resting potential of the cell.

Figure 6 plots the membrane potential in the Nernst potential simulator when K^+ , Na^+ , and Cl^- ions are each present with initial concentration gradients identical to those measured in the resting squid giant axon. The inset shows that the membrane potential rapidly dropped to the value predicted by equation (4.3). Since we did not add active ion pumps to the membrane, the concentration gradients of each species ran down over a very long period of time, and the final resting potential settled at the Donnan equilibrium. See Appendix C for details on calculating the Donnan equilibrium.

The Nernst potential simulator has been successfully used for the last three years in the course “Introduction to Neurobiology” at Case Western Reserve University to teach students the fundamental concepts of electrochemistry. Students have generally found these concepts difficult to grasp, and the ability to work with the simulator has approved their comprehension of these fundamental processes. The problem set based on the simulator, with the answers, is reproduced in Appendix D.

5. Chemical Kinetics

5.1. Zeroth-, First-, and Second-Order Reactions

The reaction simulator is capable of instantiating zeroth-, first-, and second-order reactions. Higher-order reactions can be instantiated with the use of chemical intermediates in a multi-step chain reaction. We validated our reaction simulator by comparing the kinetics of the simulated reactions to the classical integrated ordinary differential equation (ODE) rate laws. In order to do this, we needed first to determine the relationship between our lattice-based reaction algorithm (see figure 3), the probability of a reaction occurring, and the classical concept of a kinetic rate constant so that we could determine the theoretical value of the rate constants of each reaction in our model. The details of these calculations are included in Appendix B. We were able to measure the empirical rate constants in our simulation for each reaction type and compare these to their theoretical values, and we found that under ideal conditions, i.e., when particle density was low and mixing of reactants was introduced, they correspond well.

5.2. Diffusion-Limited Reactions

The integrated ODE rate laws assume that a reaction system is spatially homogeneous. For a model that is not continuous but has discrete entities, the closest approximation to a homogeneous reaction system would be a well-mixed reaction system. The kinetics of second-order reactions can become limited by the rate of diffusion of the reactants. The simple reaction $A + B \longrightarrow \emptyset$, in which reactants of two species combine to annihilate one another, reveals in the simulator that the

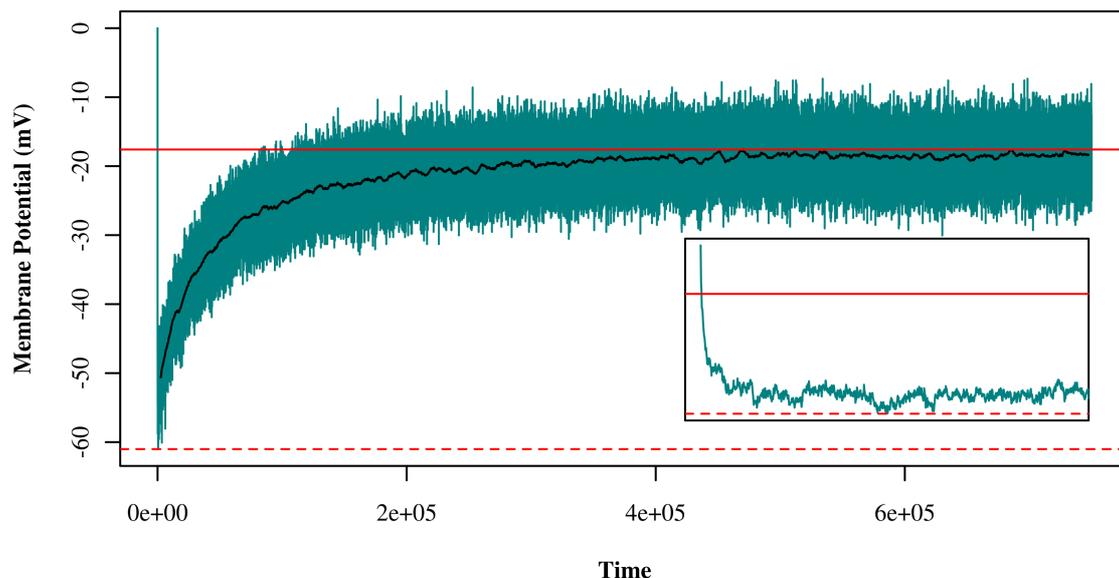


Figure 6: The initial evolution of the steady state resting potential across a semipermeable membrane followed by the Donnan equilibrium. A simulation of $n = 129523$ ions in solution separated by a semipermeable membrane was initialized with concentration gradients [17, p. 128] and relative membrane permeability values [17, p. 133] of potassium, sodium, and chloride ions identical to those measured in the squid giant axon at rest. Passive diffusion alone rapidly led to the development of the resting potential (see inset, which plots the membrane potential up to $t = 1000$ iterations; elapsed runtime was approximately 1 minute on a computer with a 2.1 GHz Intel Core 2 processor). Equation (4.3) gives the membrane potential of a neuron at rest and successfully predicts the resting potential for the system, indicated by the dashed, horizontal line at $V_m = -61$ mV. Because the simulator does not currently incorporate active ion pumps, the concentration gradients that contribute to the resting potential are not maintained. Consequently, the resting potential was only a steady state, and the concentration gradients broke down until an electrochemical equilibrium was established (approximately 10 hours total runtime on a computer with a 2.1 GHz Intel Core 2 processor for this simulation). The Donnan equilibrium is indicated by the solid, horizontal line at $V_m = -18$ mV (see Appendix C for details). The black curve centered on the data indicates the moving average of the membrane potential across 5000 time-steps.

minority reactant in each local domain is rapidly depleted until the domain contains only a high-concentration of the remaining species (see figure 7). The domains grow as smaller domains are destroyed by the encroachment of neighboring domains, and eventually the reactants are exhausted. However, because these domains form and many particles are isolated from particles of the other species, the rate of the reaction is reduced significantly relative to what the ODE rate laws predict. This is true for any second-order reaction involving two different reacting species.

If a reaction system is well-mixed, then these domains can never form, and the kinetics are more accurately predicted by the ODE rate laws. A switch exists in our model that allows mixing to be simulated by randomly shuffling the positions of all particles in the lattice every iteration.

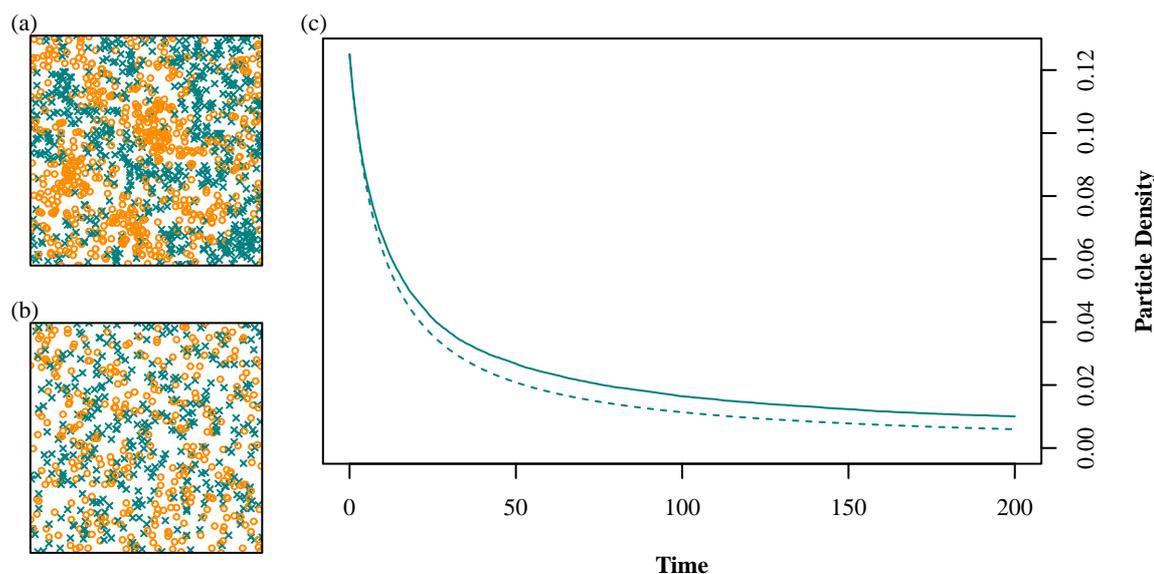


Figure 7: Domain formation in a simple diffusion-limited particle annihilation system. (a) When particles of two species react to destroy one another and the kinetics are fast enough such that the reaction is diffusion-limited, local domains with high concentration of a single species tend to form despite an initial uniform distribution of the reacting species. This snapshot of the lattice was taken after 200 iterations. (b) When automatic mixing is enabled for the same system, domains cannot form. This snapshot was taken after the same interval. (c) The existence of these domains when mixing is disabled reduces the rate of the reaction in the simulation relative to the rate predicted by the ODE rate law for the system. The solid curve indicates the density of one of the reacting species in the simulator when mixing is disabled, and the dashed curve indicates the integrated ODE rate law. When mixing is enabled, the simulation is nearly indistinguishable from the integrated rate law (not shown).

Figure 7 shows the absence of domain formation and the increased rate of reaction when mixing is enabled in the simulator.

The existence of this phenomenon has been known for more than three decades [24, 29]. The effect of domain formation on reaction kinetics is weaker when the reaction system is represented in a three-dimensional space because of the increased degrees of freedom of movement for the particles, but it is still present [18, 21, 19]. In living cells, many reactions occur on two-dimensional surfaces, and, because of significant crowding in the cell, classical ODE rate laws may not be good models for many phenomena seen *in vivo* [26]. Our simulator demonstrates plainly why this is true.

5.3. Michaelis–Menten Enzyme Kinetics

The Michaelis–Menten model of enzyme kinetics is very useful for understanding how enzymes and their substrates interact to form a chemical product [1, p. 164]. The model reaction system is



where E is the enzyme, S is the substrate chemical, ES is an intermediate enzyme–substrate complex, and P is the final product.

This system can be easily implemented in our simulator, and the results are compared with the integrated ODE rate laws in figure 8. Much of the deviation from the integrated rate laws can be explained as diffusion-limited kinetics, and mixing of the particles removes some of the discrepancy. The remainder is in part a consequence of the lattice implementation of the world, in which the dissociation reactions of the enzyme–substrate complex require an empty lattice position to be located nearby. If empty positions are scarce, this can reduce the rates of these reactions.

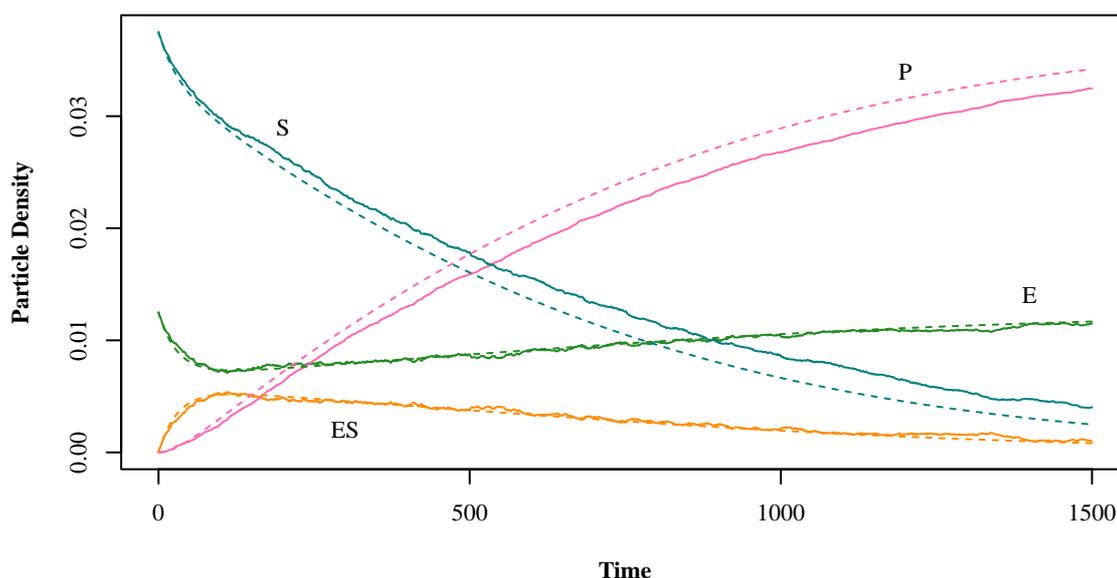


Figure 8: Behavior of the Michaelis–Menten enzyme kinetics system in the simulator. The reaction system described in equation (5.1) was built in the simulator with 2344 substrate (“S”) particles and 781 enzyme (“E”) particles initially placed in the lattice. The solid curves indicate the particle densities for each species in the simulator, and the dashed curves indicate the integrated ODE rate laws. Automatic mixing was disabled for this simulation. With mixing, the simulation results are closer to the integrated rate laws (not shown).

Unlike the famous Michaelis–Menten equation, which describes the rate of product formation as a function of the substrate concentration, our simulation at no time assumes that the system is in a steady state or that the amount of substrate-bound enzyme remains roughly constant. In fact, the fast, transient behavior that is ignored by the Michaelis–Menten equation can easily be seen in

figure 8. Note that a significant amount of substrate is bound to the enzyme before the steady state is reached, and by this time some substrate has already been converted to product. Calculations that use the Michaelis–Menten equation with the *initial* concentration of substrate to estimate the rate of product formation should be careful that this value has not changed too significantly by the time the steady state is established.

6. Ecological Dynamics: Lotka-Volterra Predator-Prey System

The Lotka–Volterra differential equation model of predator–prey interactions is an ecological model that exhibits oscillatory population growth and decline for both predator and prey species [13, p. 155]. The system is sometimes represented as

$$\begin{aligned}\frac{dR}{dt} &= (r - \alpha F) R, \\ \frac{dF}{dt} &= (\varepsilon\alpha R - \delta) F,\end{aligned}$$

where R and F are the numbers of rabbits (prey) and foxes (predator), r is the growth rate of the rabbit population, α is the attack rate of fox and rabbit interactions, ε represents the net gain the fox population sustains from each predator–prey interaction, and δ is the mortality rate of the foxes. In the absence of foxes, the rabbit population grows exponentially. In the absence of rabbits, the fox population declines exponentially. When both species are present, their interactions create a neutrally stable limit cycle of oscillatory behavior, in which the rabbits experience a boom in population that fuels a growth spurt in the fox population. This drives the number of rabbits down, and in turn the fox population plummets. The cycle continues indefinitely.

This ODE system can be converted into a reaction system. If we let $\varepsilon = 1$, then



is an equivalent reaction system that captures the traditional Lotka–Volterra model. The behavior of this system when implemented in our simulator is shown in figure 9. Regardless of whether mixing is enabled in the simulator, the oscillations of the system die down eventually, in contrast to the integrated rate laws, which exhibit a neutrally stable limit cycle. The stochastic nature of our simulator, along with the limited space in the lattice that makes exponential growth impossible, contributes to this.

The greatest cause of the deviations, however, between the simulation results without mixing and the integrated rate laws becomes apparent when looking at the simulated world itself (see figure 10). When mixing is enabled, the foxes and rabbits are distributed evenly throughout the lattice. In contrast, when mixing is disabled, domains form that provide refuge for the prey species

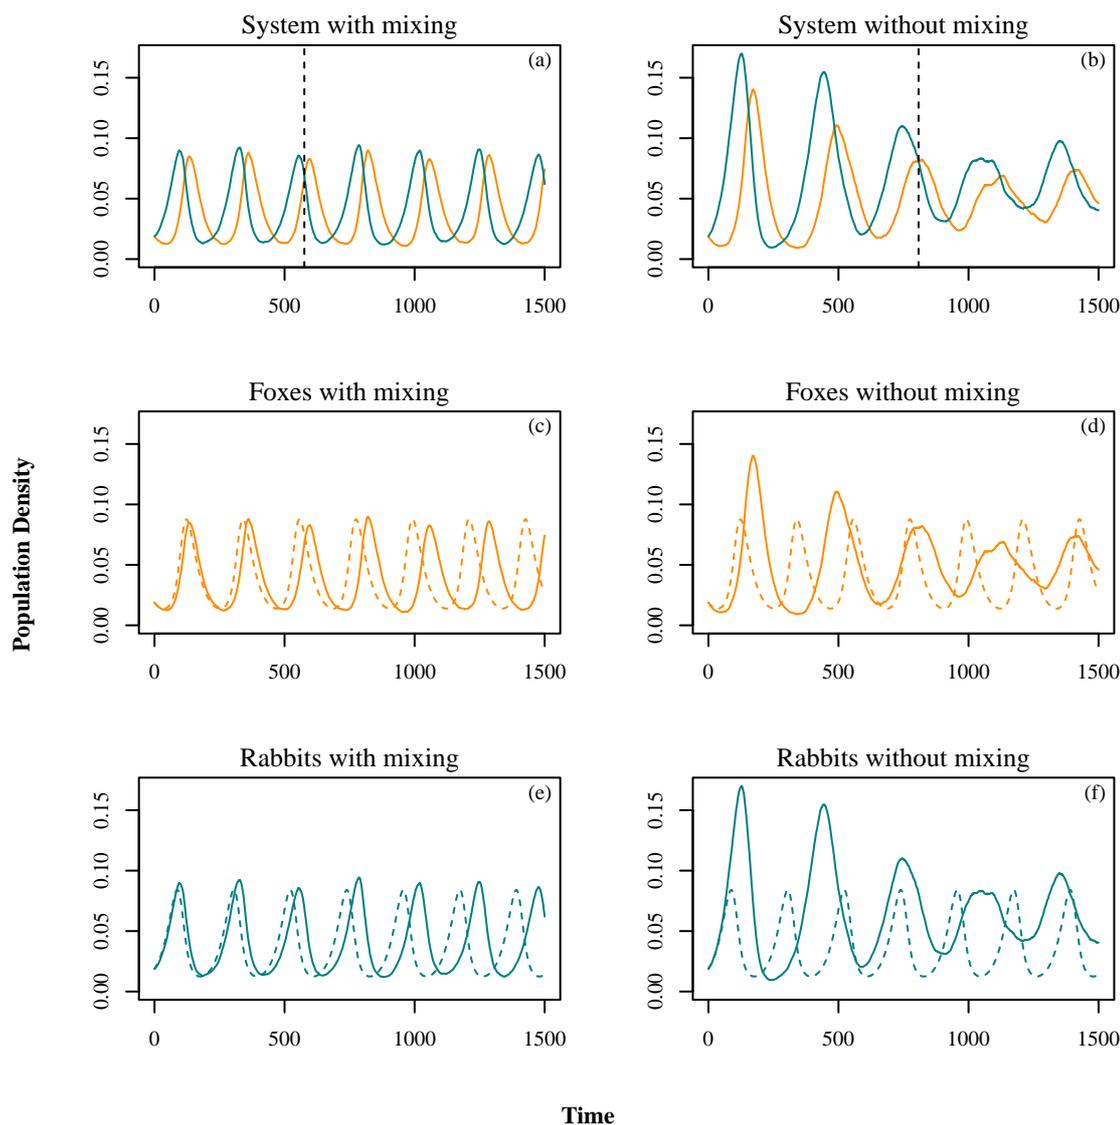


Figure 9: Behavior of the Lotka–Volterra predator–prey system in the simulator. The reaction system described in equation (6.1) was built in the simulator with equal initial quantities of foxes and rabbits and run with and without automatic mixing. The overall system behavior (a) with and (b) without mixing is plotted. The dashed, vertical lines indicate the times at which the lattice snapshots displayed in figure 10 were taken. The simulated behavior of the fox population is compared to the integrated ODE rate laws, shown as dashed curves, (c) with and (d) without mixing. Similarly, the simulated behavior of the rabbit population is compared to the integrated ODE rate laws (e) with and (f) without mixing. Because the ODE rate laws assume a spatially homogeneous medium, simulations with mixing compare more favorably with the integrated rate laws.

and limit the predator–prey interaction to occurring only on the edges of the domains. This creates massive fluctuations in both species populations that are not present in the homogeneous case considered by the ODE rate laws.

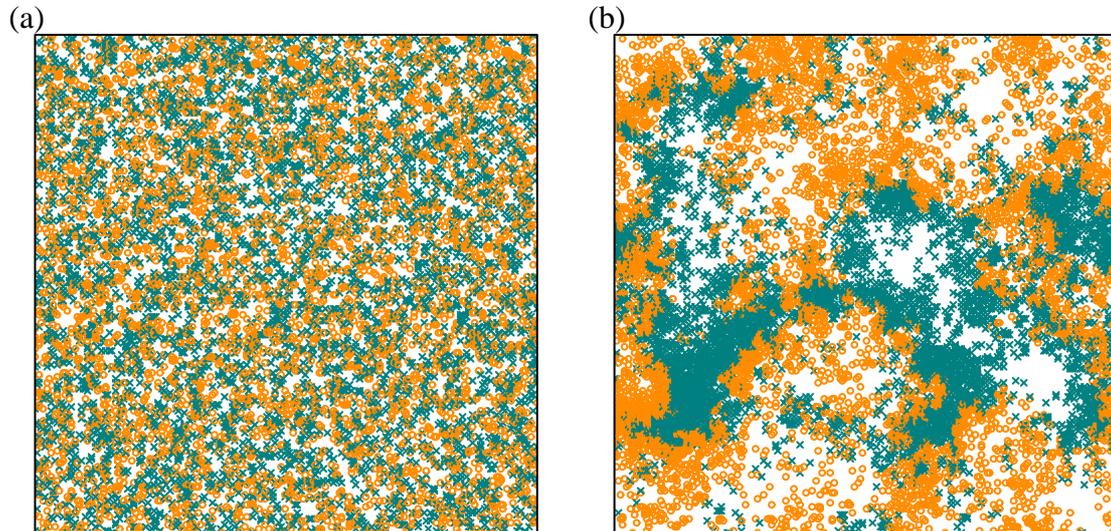


Figure 10: Snapshots of the Lotka–Volterra predator–prey system in the simulator. The reaction system described in equation (6.1) was built in the simulator with equal initial quantities of foxes (“o”) and rabbits (“x”) and run (a) with and (b) without automatic mixing. The snapshots were taken of the lattice at approximately the same time during the third oscillation of the system when foxes and rabbits were equally abundant (see figure 9). The spatial distributions of the species reveal that when mixing is disabled, species-exclusive domains form which temporarily provide refuge for the rabbits and dampen the growth of the foxes.

Our simulation demonstrates that a simple, classical ecological model such as the Lotka–Volterra system can deviate very significantly from what occurs in nature. Predator–prey interactions studied in real-world land ecology often occur in a two-dimensional space, like the lattice, in which we have seen species-exclusive domains form that can dramatically affect the dynamics of the system. Additionally, the classical continuous ODE Lotka–Volterra system—which lacks discrete quantities of rabbits and foxes—can never obtain extinction of either species. This is possible both in nature and in our simulation. Furthermore, our model shows that clustering or flocking behavior can be produced by nothing more than this simple physical phenomenon, and neither natural selection nor intelligent decision-making on the part of the organisms involved are necessarily required for behavior of this very simple sort to occur in nature.

7. Conclusions

We have striven to validate the phenomena that our model has demonstrated using analytical methods. The model shows that departures from the continuum approximations of classical models based on ODEs can be significant under certain conditions, especially when chemicals in solution are not well-mixed, as is the case in many biological systems. In particular, we have successfully demonstrated that taking into account the actual space occupied by particles and their nonuniform distribution in space can lead to departures from classical models of chemical kinetics and the Lotka–Volterra predator–prey system.

Our simulator highlights the deviations from reality caused by many of the common assumptions used by classical physical, chemical, and biological models. The simulator is both instructive on the consequences of these assumptions and illuminating for understanding the modeled phenomena as it occurs in nature. The simulator demonstrates that a very small number of particles (relative to the number of particles present in a biological cell) are sufficient to produce interesting collective behavior.

We have also shown that a single simulation framework can be used to simulate a wide variety of physical, chemical, and biological phenomena. The Nernst potential simulator has been used successfully in a classroom setting. The reaction simulator has not been used in a classroom yet, but we have plans to do so.

Although excellent simulations exist for teaching students many of the concepts we present here, to our knowledge, none of these simulations models the underlying mechanisms from the interaction of component parts. These easy-to-use and attractive simulations are based on equations that assume spatial homogeneity, deterministic behavior, and that all components are represented as continuous variables. For example, a Nernst–Goldman equation simulator with a student-friendly user interface is freely accessible to students and instructors, and allows users to manipulate the Nernst and the Goldman–Hodgkin–Katz equations to see the effects of changes in temperature, ion concentrations, and membrane permeabilities [3]. The simulation immediately tracks changes in parameters, showing that all membrane potentials are directly computed from the underlying equations. Although one window shows schematic movement of ions, this display can be paused without affecting the simulation’s ability to update, demonstrating that it is merely an illustration and not the basis of the computation. In contrast, our simulation is capable of showing the transient behavior as the membrane approaches its final potential, fluctuations in the potential especially in small compartments that contain few ions, and the run-down of the membrane potential in the absence of ion pumps, because the simulation is done at the level of the individual ions.

Similarly, elegant and easy-to-use simulations exist for Michaelis–Menten enzyme kinetics [10, 33]. These are based on either simulating the underlying ordinary differential equations describing the full four-dimensional enzyme system or the simplified equation for the steady state. Interestingly, there is a simulation of chemical kinetics that incorporates stochasticity using the Gillespie algorithm [9], but this still assumes spatial homogeneity unlike our simulation.

At the level of populations, several excellent simulations exist for the Lotka–Volterra predator–prey system, such as [32, 34]. These simulators are based on the ordinary differential equations, however, and thus they cannot show the importance of spatial inhomogeneity for the creation of

refugia for the different interacting populations.

In some cases, the relatively small size of the simulated world in our model makes it difficult to generalize to larger systems. The discretization of space inherent in using a lattice and the small number of participating particles used in our simulation contribute to the amount of noise seen in our data. For many oscillators, noise is known to cause a decrease in the frequency of oscillations [11]. The simulation results shown in figures 9c and 9e may be improved by a reduction of noise. By scaling the simulation up and increasing the number of elements in the system, the amount of noise should decrease, and we expect our simulation results to improve.

A basic assumption of our simulation is that every particle attempts to move or to react simultaneously at every iteration. Previous work in simulations of lattice gasses and zeolites has demonstrated that ensuring a Poisson distribution for rates of movement on a lattice in which double occupancy is prohibited such as ours is essential for ensuring that transport diffusivity remains concentration independent [20, 23]. At low densities, the differences are unlikely to be significant and will not affect the pedagogical value of the simulation. The rationale for the current simulation approach was that it is very computationally efficient at medium and higher densities because it accesses memory in a predictable manner, and it is thus highly scalable. In future studies, we will explore the effect of using Poisson distributed movement or reaction events [23].

8. Future Directions

To extend the teaching applications of our simulator, we are interested in adding voltage-gated ion channels so that action potentials can be simulated in the Nernst potential simulator. We would also like to apply the lessons of the importance of spatial distributions in our model to the classical susceptible–infectious–recovered (SIR) model of epidemiology.

We would also like to extend the model to three-dimensions. By restricting the movement of the simulated particles to a two-dimensional space, additional error is undoubtedly introduced when attempting to model three-dimensional phenomena. Visualizing a three-dimensional model will be an added challenge.

As a research tool, we want to continue to pursue designing an environment in which autopoietic, life-like behavior can emerge as a stepping stone toward understanding the origin of life. In order to do this, we believe we will need a generative chemistry that is capable of producing novel substances in the simulator that can be selected for by natural selection. We would like this generative chemistry to ultimately allow for the spontaneous formation of micelles and membranes from amphipathic molecules, for template-directed processes to succeed in the self-description and replication of chemical systems, and for enzyme catalysis to be driven by the shapes and properties of the molecules involved. Designing a generative chemistry that still allows our simulator to remain both a faithful representation of reality and a computationally efficient model will be difficult. Despite these challenges, we are excited by the promise of this approach.

Acknowledgements

We are grateful to the Howard Hughes Medical Institute (HHMI) for providing support through the Summer Program in Undergraduate Research (SPUR) at Case Western Reserve University (CWRU) and to the National Science Foundation (NSF) for providing support through its Interdisciplinary Training for Undergraduates in Biological and Mathematical Sciences (UBM) program (DUE-0634612). We would also like to thank Dr. Robin Snyder for her advice and guidance.

References

- [1] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, P. Walter. *Molecular biology of the cell*. Garland Science, New York, 4th ed., 2002.
- [2] P. Atkins, J. de Paula. *Physical chemistry*. W. H. Freeman, New York, 7th ed., 2002.
- [3] M. Branch, S. Wright. The Nernst/Goldman equation simulator. <http://www.nernstgoldman.physiology.arizona.edu/>.
- [4] B. R. Brooks, C. L. Brooks, A. D. Mackerell, L. Nilsson, R. J. Petrella, B. Roux, Y. Won, G. Archontis, C. Bartels, S. Boresch, A. Caffisch, L. Caves, Q. Cui, A. R. Dinner, M. Feig, S. Fischer, J. Gao, M. Hodoscek, W. Im, K. Kuczera, T. Lazaridis, J. Ma, V. Ovchinnikov, E. Paci, R. W. Pastor, C. B. Post, J. Z. Pu, M. Schaefer, B. Tidor, R. M. Venable, H. L. Woodcock, X. Wu, W. Yang, D. M. York, M. Karplus. *CHARMM: the biomolecular simulation program*. *J. Comput. Chem.*, 30 (2009), No. 10, 1545–1614.
- [5] H. Casanova, F. Berman, T. Bartol, E. Gokcay, T. Sejnowski, A. Birnbaum, J. Dongarra, M. Miller, M. Ellisman, M. Faerman, G. Obertelli, R. Wolski, S. Pomerantz, J. Stiles. *The virtual instrument: support for grid-enabled MCell simulations*. *Int. J. High Perform. C.*, 18 (2004), No. 1, 3–17.
- [6] P. S. di Fenizio, P. Dittrich, W. Banzhaf. *Spontaneous formation of proto-cells in an universal artificial chemistry on a planar graph*. In: J. Keleman, P. Sosik, editors. *Advances in Artificial Life*. 6th European Conference, ECAL 2001, 2001 Sep 10–14, Prague, Czech Republic. *Lect. Notes Comput. Sc.*, 2159 (2001), 206–215.
- [7] P. Dittrich, J. Ziegler, W. Banzhaf. *Artificial chemistries - a review*. *Artif. Life*, 7 (2001), No. 3, 225–275.
- [8] A. Einstein. *Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen*. *Ann. Phys.-Berlin*, 17 (1905), 549–560.
- [9] B. M. Frezza. Deterministic versus stochastic chemical kinetics. <http://demonstrations.wolfram.com/DeterministicVersusStochasticChemicalKinetics/>.

- [10] B. M. Frezza. Michaelis-Menten enzyme kinetics and the steady-state approximation. <http://demonstrations.wolfram.com/MichaelisMentenEnzymeKineticsAndTheSteadyStateApproximation/>.
- [11] R. F. Galán. *Analytical calculation of the frequency shift in phase oscillators driven by colored noise: implications for electrical engineering and neuroscience*. Phys. Rev. E, 80 (2009), No. 3, 036113.
- [12] R. Grima, S. Schnell. *Modelling reaction kinetics inside cells*. Essays Biochem., 45 (2008), 41–56.
- [13] W. S. C. Gurney, R. M. Nisbet. *Ecological dynamics*. Oxford Univ. Press, New York, 1998.
- [14] B. Hille. *Ionic channels of excitable membranes*. Sinauer Associates, Sunderland, MA, 3rd ed., 2001.
- [15] T. J. Hutton. *Evolvable self-reproducing cells in a two-dimensional artificial chemistry*. Artif. Life, 13 (2007), No. 1, 11–30.
- [16] D. Johnston, S. M. Wu. *Foundations of cellular neurophysiology*. MIT Press, Cambridge, 1994.
- [17] E. R. Kandel, J. H. Schwartz, T. M. Jessell. *Principles of neural science*. McGraw-Hill, 4th ed., 2000.
- [18] K. Kang, S. Redner. *Fluctuation-dominated kinetics in diffusion-controlled reactions*. Phys. Rev. A, 32 (1985), No. 1, 435–447.
- [19] Z. Konkoli. *Diffusion controlled reactions, fluctuation dominated kinetics, and living cell biochemistry*. In: S. B. Cooper, V. Danors, editors. *Computational Models from Nature*. 5th Workshop on Developments in Computational Models, DCM 2009, 2009 Jul 11, Rhodes, Greece. EPTCS, 9 (2009), 98–107.
- [20] R. Kutner. *Chemical diffusion in the lattice gas of non-interacting particles*. Phys. Lett. A, 81 (1981), No. 4, 239–240.
- [21] F. Leyvraz, S. Redner. *Spatial structure in diffusion-limited two-species annihilation*. Phys. Rev. A, 46 (1992), No. 6, 3132–3147.
- [22] B. McMullin. *Thirty years of computational autopoiesis: a review*. Artif. Life, 10 (2004), No. 3, 277–295.
- [23] P. H. Nelson, A. B. Kaiser, D. M. Bibby. *Simulation of diffusion and adsorption in zeolites*. J. Catal., 127 (1991), No. 1, 101–112.
- [24] A. A. Ovchinnikov, Y. B. Zeldovich. *Role of density fluctuations in bimolecular reaction kinetics*. Chem. Phys., 28 (1978), 215–218.

- [25] J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, K. Schulten. *Scalable molecular dynamics with NAMD*. J. Comput. Chem., 26 (2005), No. 16, 1781–1802.
- [26] S. Schnell, T. E. Turner. *Reaction kinetics in intracellular environments with macromolecular crowding: simulations and rate laws*. Prog. Biophys. Mol. Bio., 85 (2004), 235–260.
- [27] H. Suzuki. *An approach toward emulating molecular interaction with a graph*. Aust. J. Chem., 59 (2006), No. 12, 869–873.
- [28] K. Takahashi, N. Ishikawa, Y. Sadamoto, H. Sasamoto, S. Ohta, A. Shiozawa, F. Miyoshi, Y. Naito, Y. Nakayama, M. Tomita. *E-Cell 2: multi-platform E-Cell simulation system*. Bioinformatics, 19 (2003), No. 13, 1727–1729.
- [29] D. Toussaint, F. Wilczek. *Particle-antiparticle annihilation in diffusive motion*. J. Chem. Phys., 78 (1983), No. 5, 2642–2647.
- [30] J. Trefil, H. J. Morowitz, E. Smith. *The origin of life*. Am. Sci., 97 (2009), No. 3, 206–213.
- [31] F. Varela, H. Maturana, R. Uribe. *Autopoiesis: the organization of living systems, its characterization and a model*. Biosystems, 5 (1974), No. 4, 187–196.
- [32] E. W. Weisstein. Predator-prey equations. <http://demonstrations.wolfram.com/PredatorPreyEquations/>.
- [33] T. Weisstein. Michaelis-Menten enzyme kinetics. http://bioquest.org/esteem/esteem_details.php?product_id=246.
- [34] T. Weisstein, R. Salinas, J. R. Jungck. Two-species model. http://bioquest.org/esteem/esteem_details.php?product_id=203.

A Ion Channel Model

In the Nernst potential simulator, two compartments are separated by a vertical membrane embedded with ion channels. Each time-step, every ion in the lattice is given the opportunity to move, and ions moved to positions adjacent to an ion channel of matching type are given the opportunity to flow across the membrane to the other compartment. We use the Boltzmann distribution to determine the probability of flow across the membrane.

According to the Boltzmann distribution, for a collection of particles in thermodynamic equilibrium with its surroundings, the probability of selecting a particle in energy state i with energy E_i is

$$p_i = \frac{e^{-\beta E_i}}{\sum_j e^{-\beta E_j}},$$

where $\beta = 1/(k_B T)$, k_B is the Boltzmann constant, and T is the temperature. [2, p. 632].

We assume that ion flow through pores across the membrane is very fast relative to the length of a time-step. Consequently, if an ion moves into a position adjacent to an ion channel of matching type after the movement portion of a time-step, the probability of finding the ion on either side of the membrane at the end of the time-step is independent of the side on which the ion started, and it depends only on the energies of the ion when placed on either side of the membrane and the temperature. For a particular ion,

$$p_L = \frac{e^{-\beta E_L}}{e^{-\beta E_L} + e^{-\beta E_R}} \quad \text{and} \quad p_R = \frac{e^{-\beta E_R}}{e^{-\beta E_L} + e^{-\beta E_R}},$$

where p_L and p_R are the probabilities of the ion being found on the left or right side of the membrane, and E_L and E_R are the energies of the ion when it is located on the left or right side of the membrane.

By taking the ratio of these probabilities, the energy difference, $\Delta E = E_R - E_L$, can be used instead of the absolute energies, i.e.,

$$\frac{p_L}{p_R} = e^{\beta \cdot \Delta E}.$$

Since $p_L + p_R = 1$,

$$p_L = \frac{e^{\beta \cdot \Delta E}}{1 + e^{\beta \cdot \Delta E}}.$$

The energy difference ΔE is calculated from the work required to move a charge through an electric field, $\Delta E = V \cdot q$, where V is the membrane potential and q is the charge on the ion. The membrane potential is calculated as $V = Q/C$, where Q is the charge separation between compartments and C is the capacitance of the membrane. Thus, p_L and p_R are calculated and a random number allows the simulation to determine where an ion should be placed at the end of a time-step.

Alternatively, flow of ions across the membrane can be modeled as slow relative to diffusion away from the pore and the length of a time-step. By using a similar derivation and setting an arbitrary maximum rate of ion flow through pores, an alternative method of determining when ions

cross the membrane can be devised. It is likely that ion flow through pores is somewhat slower than free diffusion for real biological systems [14, pp. 313, 400]. However, we have tried implementing both methods in the Nernst potential simulator, and we found that they have the same long-term behavior. The first method described in detail above allows ion movement across the membrane to occur more frequently, resulting in faster evolution of the system that is desirable for pedagogical reasons. For these reasons, we chose to implement ion flow through pores as it is first described above.

B Calculating Reaction Rate Constants for our Model

The reaction algorithm used in our model is a bit more complicated than is presented in figure 3. The full details of the reaction algorithm are as follows. Each particle picks either its own position or one of the four lattice positions to its east, southeast, south, or southwest with equal probability (ignoring the other neighboring positions ensures that each possible reactant pair is only checked once). If the space is a neighboring position, there are at most two possible second-order reactions that could occur, or if the space chosen was the particle's own position, there are at most two possible first-order reactions that could occur. One of these two possibilities is chosen at random. If the chosen reaction exists in the predefined reaction table, then a random number is checked against a probability P specific to the reaction to determine whether the "colliding" particles have enough energy and the proper orientation to react. Thus, the probability of a particular first-order reaction occurring or a particular second-order reaction occurring with a particular neighboring particle is $P/10$ (since only one of five possible positions and one of two possible reactions will be chosen). If the chosen reaction is successful, the particle types of the reactants are changed to those of the products.

In order to limit particles from reacting more than once each time-step, if a particle would be claimed for participation in more than one reaction within a single time-step, all reactions claiming that particle are simply cancelled.

Furthermore, reactions that do not have a balanced number of reactant and product particles must either create or destroy particles by using empty lattice positions. If a reaction has two reactants and only one product (synthesis), then one of the two reactant particles must change type to match the product, and the other is converted to an empty space. If a reaction has one reactant and two products (decomposition), then when the reactant particle is selecting a reaction partner as described above, it must select an empty space. If the reaction is successful, its type will change, and a particle will be created in the empty space with an appropriate type.

Each empty space is also given an opportunity to select a reaction partner in the same manner as described above. This allows zeroth-order reactions to occur when the number of empty spaces remains roughly constant. Empty spaces can also select non-empty spaces to facilitate decomposition reactions. In these cases, it should be noted that empty spaces can be thought of as containing solvent particles which are assumed to have a steady concentration and may be capable of reacting with one another or undergoing conformational changes.

Thus, we have six fundamental reaction types:

1. $\emptyset \longrightarrow A$
2. $\emptyset + \emptyset \longrightarrow A + B$ (shorthand: $\emptyset \longrightarrow A + B$)
3. $A \longrightarrow B$
4. $A + \emptyset \longrightarrow B + C$ (shorthand: $A \longrightarrow B + C$)
5. $A + B \longrightarrow C + \emptyset$ (shorthand: $A + B \longrightarrow C$)
6. $A + B \longrightarrow C + D$

Empty spaces are represented by the symbol \emptyset , and types A, B, C, and D can be any particle types. The first two are considered zeroth-order because they have no particle reactants, and therefore their rate laws do not depend on the concentration of any particles. The first can be thought of as solvent isomerization, and the second is a reaction between solvent particles or the decomposition of a single solvent particle. The third and fourth reactions are first-order because each has only one particle reactant. The third can be thought of as particle isomerization, and the fourth is the standard decomposition reaction. The fifth and sixth reactions are second-order because each has two reactants. The fifth is the standard synthesis reaction, and the sixth is a standard substitution or replacement reaction.

If we assume that particle density is low, then the number of empty spaces will remain roughly constant in the near future, and particle diffusion will not be hindered by collisions. If we also assume that individual particles are very rarely claimed for multiple reactions (a scenario which results in no reaction occurring at all) and that reactants are always well-mixed, then we can estimate the number of reactions that will occur in the artificial chemistry each iteration. Table 1 displays the expected number of reactions per iteration for each of the six reaction types (and two special cases). From these expressions, the rate laws can be deduced. When the number of particles of a particular type produced by a reaction is exactly one, then the number of reactions per iteration divided by the size of the lattice ($x \times y$ if x and y are the dimensions of the lattice) is the rate of change in the concentration of the chemical, or the rate law for the reaction. Furthermore, if we know the rate law, we can extract from it the rate constant k . The rate laws and rate constants for our six reactions (and the two special cases) are displayed in table 1.

C Finding the Donnan Equilibrium in a Finite System

The Donnan equilibrium is an electrochemical equilibrium that forms across a cell membrane when some charges exist in the system that are impermeable to the membrane. For instance, these charges could be attached to large proteins in the cell. If permeable ions are also present with concentration gradients between the intracellular and extracellular compartments, then these ions will move down their gradients until the chemical gradient is counterbalanced by a charge differential. Consequently, the permeable ions are not able to diffuse freely, and the membrane potential will settle at a non-zero equilibrium.

Table 1: Theoretical rate laws for reactions in the artificial chemistry. The expected number of reactions per iteration is derived assuming that particle density is low, reactants are well-mixed, and multiple reaction partner claims are rare enough for their effects to be negligible. x and y are the dimensions of the lattice. P is the probability that reactants will have both sufficient energy to react and the correct orientation when colliding. Rate laws are related to the expected number of reactions per iteration by the factor xy , the size of the lattice, when time t is measured in iterations. Concentrations are measured in particles per lattice square. The rate constant k is taken directly from the rate law.

Reaction Type	E [Reactions Per Iteration]	Rate Law	k
Zeroth Order $\emptyset \xrightarrow{k} A$	$\frac{P}{10} \cdot xy$	$\frac{d[A]}{dt} = \frac{P}{10}$	$\frac{P}{10}$
Zeroth Order $\emptyset \xrightarrow{k} A + B$	$\frac{P}{10} \cdot 4 \cdot xy$	$\frac{d[A]}{dt} = \frac{d[B]}{dt} = \frac{2P}{5}$	$\frac{2P}{5}$
First Order $A \xrightarrow{k} B$	$\frac{P}{10} \cdot A$	$-\frac{d[A]}{dt} = \frac{d[B]}{dt} = \frac{P}{10} [A]$	$\frac{P}{10}$
First Order $A \xrightarrow{k} B + C$	$\frac{P}{10} \cdot 8 \cdot A$	$-\frac{d[A]}{dt} = \frac{d[B]}{dt} = \frac{d[C]}{dt} = \frac{4P}{5} [A]$	$\frac{4P}{5}$
Second Order $A + B \xrightarrow{k} C$	$\frac{P}{10} \cdot \frac{8}{xy} \cdot A \cdot B$	$-\frac{d[A]}{dt} = -\frac{d[B]}{dt} = \frac{d[C]}{dt} = \frac{4P}{5} [A][B]$	$\frac{4P}{5}$
Second Order $2A \xrightarrow{k} B$	$\frac{P}{10} \cdot \frac{4}{xy} \cdot A^2$	$-\frac{1}{2} \frac{d[A]}{dt} = \frac{d[B]}{dt} = \frac{2P}{5} [A]^2$	$\frac{2P}{5}$
Second Order $A + B \xrightarrow{k} C + D$	$\frac{P}{10} \cdot \frac{8}{xy} \cdot A \cdot B$	$-\frac{d[A]}{dt} = -\frac{d[B]}{dt} = \frac{d[C]}{dt} = \frac{d[D]}{dt} = \frac{4P}{5} [A][B]$	$\frac{4P}{5}$
Second Order $2A \xrightarrow{k} B + C$	$\frac{P}{10} \cdot \frac{4}{xy} \cdot A^2$	$-\frac{1}{2} \frac{d[A]}{dt} = \frac{d[B]}{dt} = \frac{d[C]}{dt} = \frac{2P}{5} [A]^2$	$\frac{2P}{5}$

The Nernst potential simulator models three ion types: K^+ , Na^+ , and Cl^- ions. At equilibrium, the Nernst equilibrium potentials of each ion species must be equal to one another [16, p. 19]. Thus, we can write,

$$E_{K^+} = E_{Na^+} = E_{Cl^-}.$$

From equation (4.2), this simplifies to

$$\frac{[K^+]_{out} - \Delta[K^+]}{[K^+]_{in} + \Delta[K^+]} = \frac{[Na^+]_{out} - \Delta[Na^+]}{[Na^+]_{in} + \Delta[Na^+]} = \frac{[Cl^-]_{in} + \Delta[Cl^-]}{[Cl^-]_{out} - \Delta[Cl^-]} = r,$$

where $[X]_{in}$ and $[X]_{out}$ are the initial concentrations of species X inside and outside of the cell, and $\Delta[X]$ is the amount of species X that transfers from outside of the cell to within as the Donnan equilibrium is established.

If the system is large and the expected change in concentrations of each ion species is small, or if only an approximation of the equilibrium potential is needed, then a solution can be found by assuming that the intracellular and extracellular spaces remain constantly neutral. However, when a small number of particles constitute the entire system, as in our simulator, then this assumption may not hold, and an exact solution is needed.

Charged particles impermeable to the membrane are implicitly represented in the simulator on both sides of the membrane to balance the initial charges and to ensure that both compartments are electrically neutral at the start of the simulation. Consequently, the initial membrane potential is always set to 0 mV, and any potential difference that develops as the Donnan equilibrium is established is due to the movement of the permeable ions. The membrane potential at equilibrium will then be

$$\frac{RT}{F} \ln r = \frac{F}{C} (\Delta[K^+] + \Delta[Na^+] - \Delta[Cl^-]),$$

where R is the molar gas constant, T is the temperature of the system, F is Faraday's constant, and C is the capacitance of the membrane. This system of equations can be solved numerically to find the equilibrium potential.

D Problem Set with Answers

The problems are presented in *italics*, and the answers are printed in upright font. Starred problems (*) are extra credit for undergraduate students and are required regular credit for graduate students.

Problem 1

(a)

Download the Nernst Potential Simulator (note that versions for Windows, Mac, and Linux are available on the website; please download the one appropriate for your computer). Each of the four exercises on this problem set is also posted (as Problem1.init, Problem2.init, and Problem3.init). Please download these files as well.

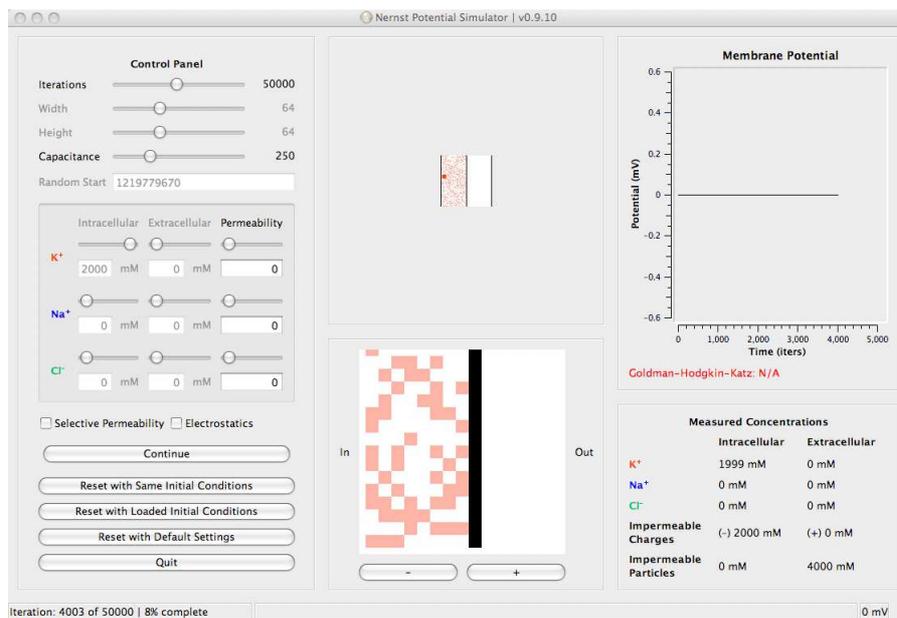
(b)

Double click to open the simulator. Under the File menu, select Load Initial Conditions. For this problem, please load Problem1.init. Note that the permeability, which is a measure of the number of channels that will allow a particular kind of ion to cross the membrane, has been set to zero for all of the ions. When the permeability is set to 1, the maximum number of channels are available to those ions.

(c) (10 points)

Start the simulation (press Start). Click on one of the red ions, which will enlarge it. Describe what you observe about the motion of the single ion, as well as the movements of all of the ions in the intracellular compartment. Predict the long term behavior of the system, and observe it for at least 2,000 iterations to see if your long term prediction is likely to be correct.

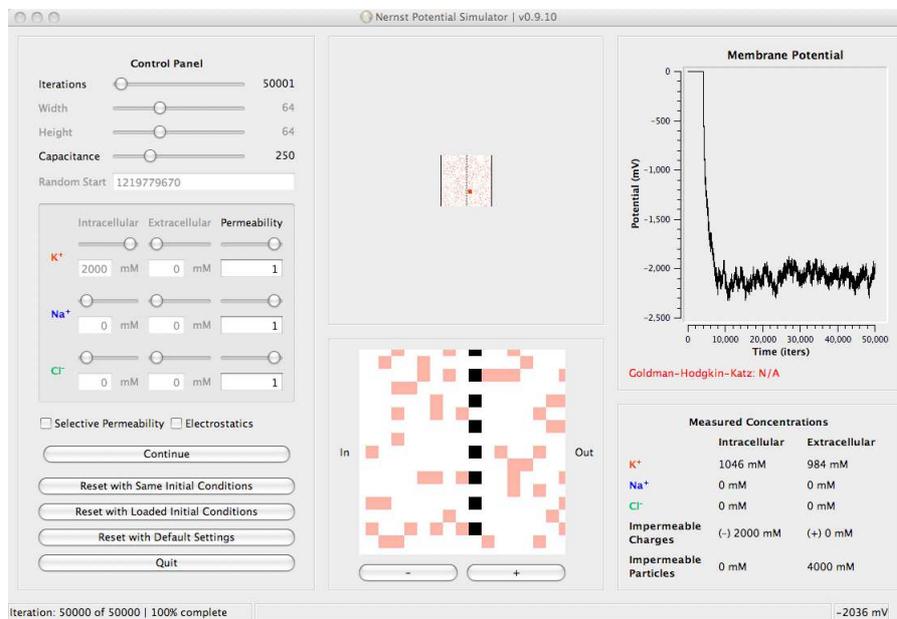
Both the single red ion and all of those surrounding it move randomly, recoiling when they encounter either the central membrane or the sides of the container. If they reach the top of the container, they may reappear at the bottom, and vice versa, because there is no barrier to their movement in this direction. Over the long term, the random movements will persist, and there will be no macroscopic change in the system.

**(d) (10 points)**

Pause the simulation (press Pause). Change the permeability for the K⁺ ions from 0 to 1 (either by moving the slider, or typing 1 into the box). Also change the permeability of the Na⁺ and Cl⁻ ions to 1. Note that the Selective Permeability box is unchecked. Thus, increasing the permeability to Na⁺ and Cl⁻ under these conditions will provide many opportunities for potassium ions to move through the membrane. You can see the many channels shown in the zoom window at the bottom middle of the simulation.

Now continue the simulation (press Continue). Observe what happens to the individual ion that you have selected, and to all of the ions over time. Also observe the change in the intracellular and extracellular concentrations of K⁺ listed in the table under the graph to the right. Predict the final intracellular and extracellular concentrations of K⁺ ions. Run the system for up to another 50,000 iterations to see if your long term prediction is likely to be correct.

Initially, the potassium ions move from the intracellular compartment to the extracellular compartment. As they do so, however, the likelihood of potassium ions moving back increases. The predicted final concentrations on the two sides of the membrane should be equal, or 1000 mM intracellularly, and 1000 mM extracellularly. The final concentrations differ somewhat from the predicted concentrations, but this is likely to be due to the fluctuations caused by the small number of model ions used in the simulation.



(e) (10 points)

Take the simulation back to the initial conditions loaded from the file by pressing the “Reset with Loaded Initial Conditions” button. Now switch the intracellular concentration of K⁺ from 2000 to 0, and the extracellular concentration of K⁺ from 0 to 2000, so that the K⁺ ions are all in the extracellular compartment. Change the permeability of all the ions from 0 to 1. Start the simulation, and again click on one of the red ions. Predict the long term intracellular and extracellular concentrations of K⁺ ions, and observe the system for about 50,000 iterations to see if your long term prediction is correct.

Now the potassium ions have a higher probability of moving from the extracellular to the intracellular compartment. Once again, the likelihood of ions moving back in the opposite direction steadily increases. The predicted final concentrations on the two sides of the membrane should again be equal, or 1000 mM intracellularly, and 1000 mM extracellularly. Once again, the final concentration differs slightly from the predicted concentrations because of the fluctuations due to the small number of model ions.

(f) (10 points)

Use these results to describe the long term behavior of any system in which a single ion has different concentrations in the intracellular and extracellular spaces, and is allowed to freely diffuse across a permeable membrane. First answer this assuming that the ion is uncharged. Then, discuss what will happen if the first ion is associated with a second ion of opposite and equal charge that has the same initial concentration on each side of the membrane.

Assuming an uncharged particle, it is very clear that diffusion of the particles from areas of higher concentration to lower concentration across a permeable membrane leads to an equalization of particles on both sides of the membrane, so that the concentrations become the average of the two original concentrations in the two compartments. If the system consists of two ions of opposite charges, their concentrations will equalize across the membrane, and they will move, on average, together so as to maintain electroneutrality.

Problem 2**(a) (10 points)**

Reset the simulation to its default values, and load the file *Problem2A.init*. Note that the model membrane is now permeable solely to potassium, and that both selective permeability and electrostatics are active; thus, you will note that the channels in the membrane only allow potassium ions through, and the program changes the probability that ions may cross the membrane based on the relative number of charges on either side of the membrane. Before starting the simulation, please calculate the predicted potential across the membrane using the Nernst equation. Show your work.

Assuming that the system is at a temperature of 20 ° centigrade, we can use the following equation:

$$E_{K^+} = 58 \log \frac{[K^+]_{out}}{[K^+]_{in}} = 58 \log \frac{20 \text{ mM}}{400 \text{ mM}} = 58 \log \frac{1}{20} = -75.5 \text{ mV}.$$

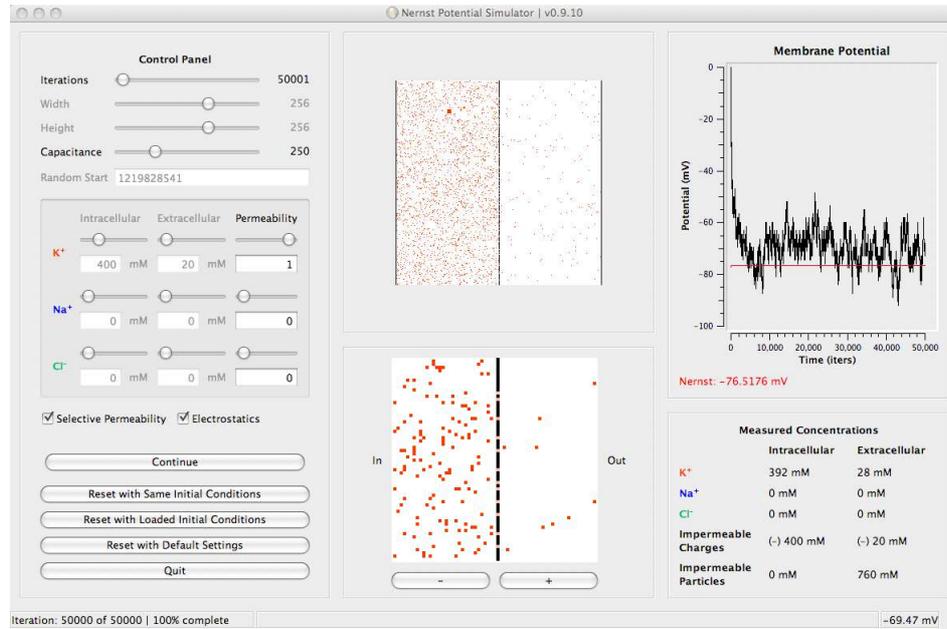
Note that the simulation used 25 °C, and so obtains slightly different answers. Also note that human body temperature is approximately 37 °C, so that the constant would be about 62 mV, rather than 58 mV.

(b) (10 points)

Now start the simulation briefly, immediately pause it, and click on a red ion on the left side (the intracellular side) so that you can readily track the movements of the ion. Then continue the simulation. Describe the behavior of the single ion, and of all the potassium ions. Also, explain the changes in the graph of the membrane potential (black line), and contrast it to the predicted value (the red line). Predict what will happen over long periods of time. Run the simulation for 50,000 iterations to see if your prediction is correct. Explain the results.

As the potassium ions begin to move from the intracellular compartment (in which they are at a higher concentration) to the extracellular compartment (in which their concentration is lower), they add positive charges to the outside of the membrane, and increase negative charges on the

inside of the membrane (as seen in the graph of the membrane potential, which drops from 0 to about -75 mV). This electrical gradient opposes the outward movement of positive charges, and thus decreases the probability that ions will flow from the intracellular to the extracellular compartment. At the same time, it increases the probability that ions will flow along the electrical gradient that they themselves have created from the extracellular to the intracellular compartment. Thus, the predicted behavior is that the black line will fluctuate around the red line. The prediction is reasonably well matched by the actual behavior of the system.



(c) (10 points)

Press the “Reset with Default Settings” button. Before running the simulation, use the Goldman-Hodgkin-Katz (GHK) equation to calculate the predicted membrane potential. Show your work. Note that the zoom window shows channels selectively permeable to potassium, sodium or chloride ions.

Again assuming a temperature of 20°C ,

$$\begin{aligned}
 V_m &= 58 \log \frac{P_{K^+} [K^+]_{out} + P_{Na^+} [Na^+]_{out} + P_{Cl^-} [Cl^-]_{in}}{P_{K^+} [K^+]_{in} + P_{Na^+} [Na^+]_{in} + P_{Cl^-} [Cl^-]_{out}} \\
 &= 58 \log \frac{1 \times 20 + 0.04 \times 440 + 0.45 \times 52}{1 \times 400 + 0.04 \times 50 + 0.45 \times 560} \\
 &= 58 \log \frac{20 + 17.6 + 23.4}{400 + 2 + 252}
 \end{aligned}$$

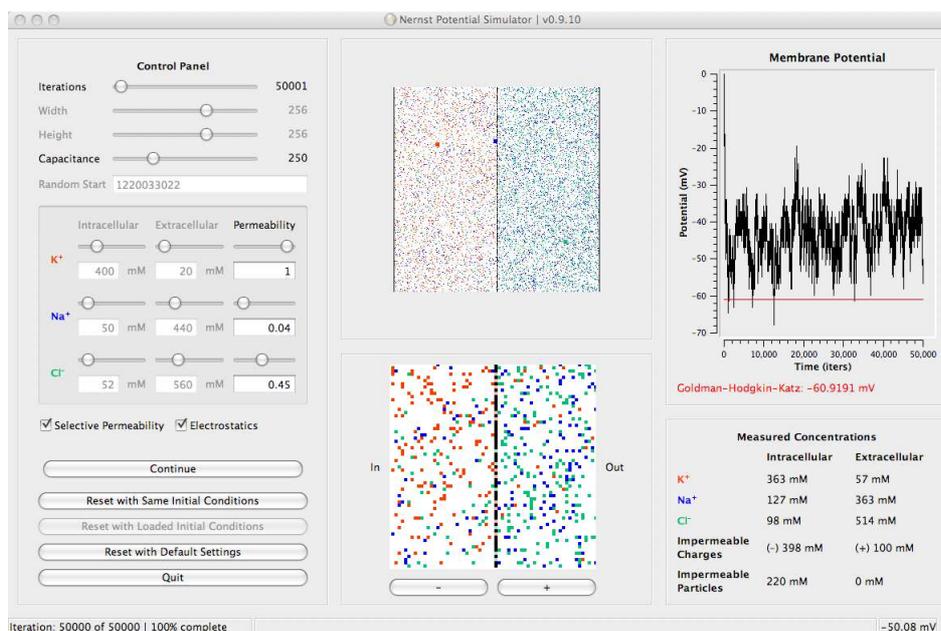
$$= 58 \log \frac{61}{654}$$

$$= -59.8 \text{ mV}$$

(d) (10 points)

Briefly start the simulation, pause it, and select a potassium (red), sodium (blue) and chloride (green) ion to track. Describe the movements of the selected ions and the graph of the membrane potential relative to the predicted potential (red line). Predict what will happen over time. Run the simulation for about 50,000 iterations to see if your prediction is correct. Explain the results.

The membrane potential starts at zero. The high permeability of the potassium ions allows them to move from the intracellular to the extracellular compartment; as they do so, the membrane potential rapidly becomes more negative. At the same time, because of the permeability of the membrane to the sodium ions, some of them can move from the extracellular to the intracellular compartment, reducing the charge separation across the membrane. Thus, the membrane sits somewhat positive to the Nernst equilibrium potential for potassium, and thus settles at -60 mV rather than -75 mV . The simulation does go down towards the red line predicted by the GHK equation, and then fluctuates above it. It shows a clear trend to be above the red line; the next question addresses this discrepancy.

**(e)* (10 points)**

The simulation represents only a very small number of the very large numbers of ions that are actually present in real nerve cells. To see how this may affect the results, press the “Reset with Default Settings” button, and then, under the File menu, select Load Initial Conditions, and load

the file *Problem2B.init*. The initial conditions are the same, but the “world’s” thickness has been greatly reduced. This would be similar to a dendritic spine surrounded by a limited extracellular space. Predict what may happen under these circumstances to the concentration gradients, and to the long term potential across the membrane. Now run the simulation for 50,000 iterations, and see what happens. At the end of the simulation, calculate the new predicted potential from the measured concentrations and the GHK equation. Explain the results, contrasting the mechanisms generating the resting potential with those that generate the Nernst potential for a single ion (parts a and b of this question). How is it that the results you obtained are not observed in most nerve cells?

Because the resting potential is a steady state, not an equilibrium, the charge changes due to the flow of potassium ions out of the cell are balanced by a net inward flow of sodium ions. Over time, however, these flows will cause the concentration of potassium ions inside the cell to decrease, and the concentration of sodium ions inside the cell to increase. As a consequence, the concentration differences begin to break down over a long enough time, and go to a Donnan equilibrium.

The *total number* of ions inside and outside of the membrane will not change. We are given the initial concentrations, whose sum is easy to compute. Thus, $[K^+]_i + [K^+]_o = 420$ mM, $[Na^+]_i + [Na^+]_o = 490$ mM, and $[Cl^-]_i + [Cl^-]_o = 612$ mM.

At the outset, and throughout the simulation, it will also have to be true that there is net charge conservation. Given that the impermeable charges are 398 mM on the inside, and -100 mM on the outside, the following must be true:

$$\begin{aligned} [K^+]_i + [Na^+]_i - [Cl^-]_i &= 398 \text{ mM, and} \\ [K^+]_o + [Na^+]_o - [Cl^-]_o &= -100 \text{ mM.} \end{aligned}$$

For the membrane to be in equilibrium, it must be the case that the Nernst potentials for the potassium, sodium, and chloride ions will ultimately all be equal to one another. Setting the Nernst potentials equal, and stripping away the constants and taking the anti-logarithms yields

$$\frac{[K^+]_o}{[K^+]_i} = \frac{[Na^+]_o}{[Na^+]_i} = \frac{[Cl^-]_i}{[Cl^-]_o}.$$

Solving these equations yields:

$$\begin{aligned} [K^+]_i &= 279 \text{ mM,} & [Na^+]_i &= 325 \text{ mM,} & [Cl^-]_i &= 206 \text{ mM,} \\ [K^+]_o &= 141 \text{ mM,} & [Na^+]_o &= 165 \text{ mM,} & [Cl^-]_o &= 406 \text{ mM.} \end{aligned}$$

Substituting any one of these internal/external pairs of concentrations into the Nernst potential yields the final Donnan equilibrium value. For example, using potassium, the predicted final potential is (assuming a temperature of 20°C)

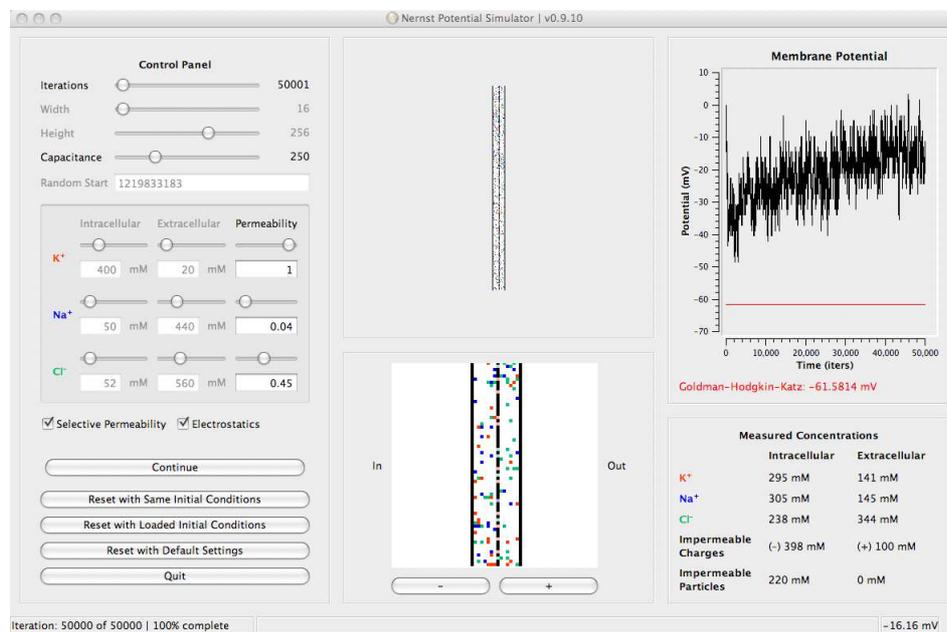
$$58 \log \frac{141 \text{ mM}}{279 \text{ mM}} = -17.2 \text{ mV}$$

For the simulation shown in the figure below, the new predicted potential after 50,000 iterations

based on the measured concentrations is

$$\begin{aligned}
 V_m &= 58 \log \frac{1 \times 141 + 0.04 \times 344 + 0.45 \times 238}{1 \times 295 + 0.04 \times 238 + 0.45 \times 344} \\
 &= 58 \log \frac{141 + 13.76 + 107.1}{295 + 9.52 + 154.8} \\
 &= 58 \log \frac{261.86}{459.32} \\
 &= -14.1 \text{ mV}
 \end{aligned}$$

The plot of the voltage over time shows many fluctuations, but the voltage never descended all the way to the predicted initial GHK potential, and is steadily rising, suggesting that over enough iterations it will reach -17 mV .



Two factors prevent this from occurring in most nerve cells. First, the number of ions is usually far higher. A millimolar concentration of an ion in one cubic microliter contains approximately 6×10^{23} ions/mole $\times 10^{-3}$ moles/liter $\times 10^{-6}$ liters/microliter = 6×10^{14} ions, so that far more time must elapse for the concentration gradients to run down. Second, most neurons contain ATP-dependent ion pumps, which steadily restore the concentration gradients. This question clearly illustrates the difference between the Nernst potential, which is a *stable equilibrium*, and the resting potential, which is an *energy-requiring steady state*.

Problem 3**(a) (10 points)**

Reset the simulation to its default values, and load the file *Problem3A.init*. Predict the final membrane potential using the GHK equation. Show your work.

Again assuming the temperature is 20 °C,

$$\begin{aligned}
 V_m &= 58 \log \frac{0.05 \times 20 + 1 \times 440 + 0.02 \times 52}{0.05 \times 400 + 1 \times 50 + 0.02 \times 560} \\
 &= 58 \log \frac{1 + 440 + 1.04}{20 + 50 + 11.2} \\
 &= 58 \log \frac{442.04}{81.2} \\
 &= +42.7 \text{ mV}
 \end{aligned}$$

(b) (10 points)

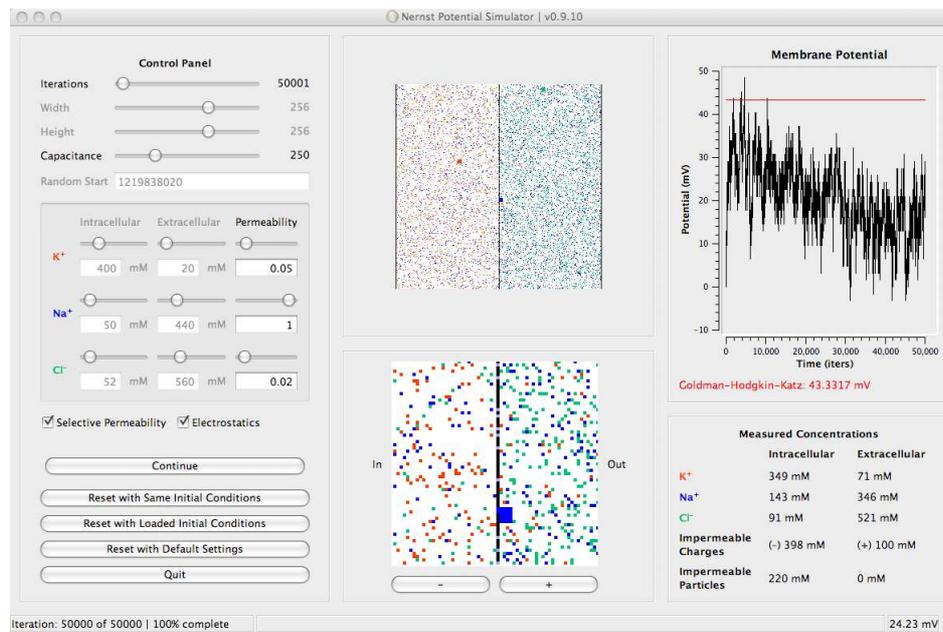
Now start the simulation briefly, immediately pause it, and click on a potassium (red), sodium (blue) and chloride (green) ion to track. Describe the movements of the selected ions and the graph of the membrane potential relative to the predicted potential (red line). Predict what will happen over time. Run the simulation for 50,000 iterations to see if your prediction is correct. From the actual measured intracellular and extracellular ion concentrations, and the GHK equation, what should the final membrane potential be? What has happened to account for the change in the initial predicted membrane potential, and that which you observe after 50,000 iterations?

Initially, the membrane potential is at zero. Because of the large permeability to sodium ions, and their high extracellular concentration, the sodium ions begin to move into the intracellular compartment, bringing with them positive charges, and the membrane rapidly depolarizes. Of course, this creates an electrical gradient that favors the outward movement of potassium ions, and the potassium ions leave the cell, keeping the membrane from ever reaching the sodium Nernst equilibrium potential (which is about +55 mV).

Interestingly, although the membrane does depolarize to the potential predicted by the GHK equation, over multiple iterations it tends to be lower than the predicted potential. This is again due to the relatively small number of ions being simulated, which leads to a much more rapid rundown of the concentration gradients than would be predicted for an actual nerve cell.

(c)* (10 points)

Devise a set of concentrations across the membrane, and permeabilities at rest and at the peak of the action potential, such that movement of chloride ions is primarily responsible for the resting



potential, and movement of potassium ions is primarily responsible for the depolarization at the peak of the action potential. In addition to presenting your values, and validating them using the GHK equation, enter them into the simulation and demonstrate that they can generate a stable resting potential and a stable depolarization.

There are many possible solutions to this problem. In order that movement of chloride ions be primarily responsible for the resting potential, it should be the case that the membrane will be most permeable to chloride ions, and much less permeable to either sodium or potassium. In addition, to allow potassium to play an important role in depolarizing the membrane, the Nernst potential for potassium ions should be very positive. This implies that the concentration of K⁺ should be high outside of the cell, and low inside the cell. Finally, the absence of Na⁺ from playing a major role in either the resting potential or the action potential suggests that its permeability should be low.

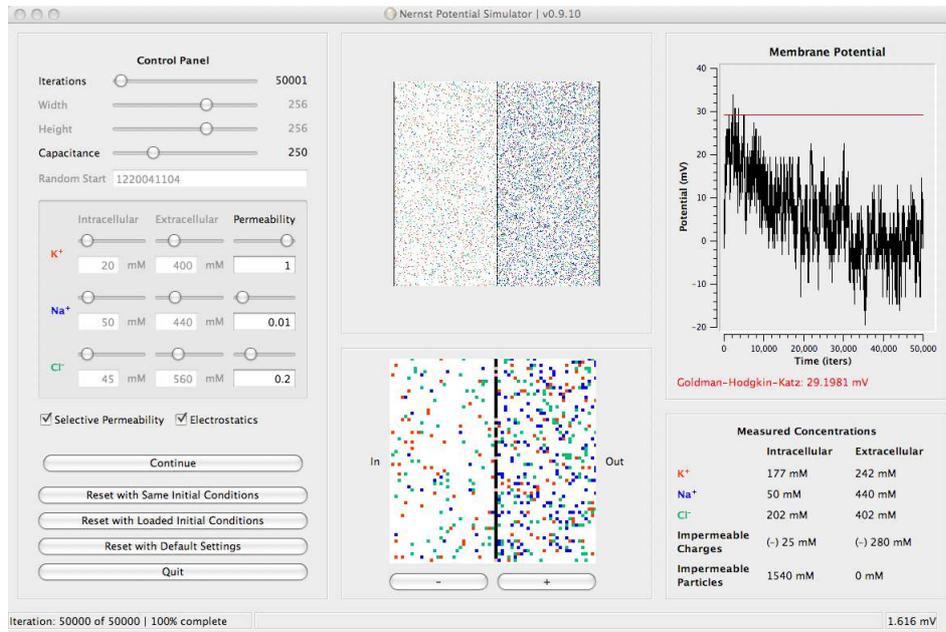
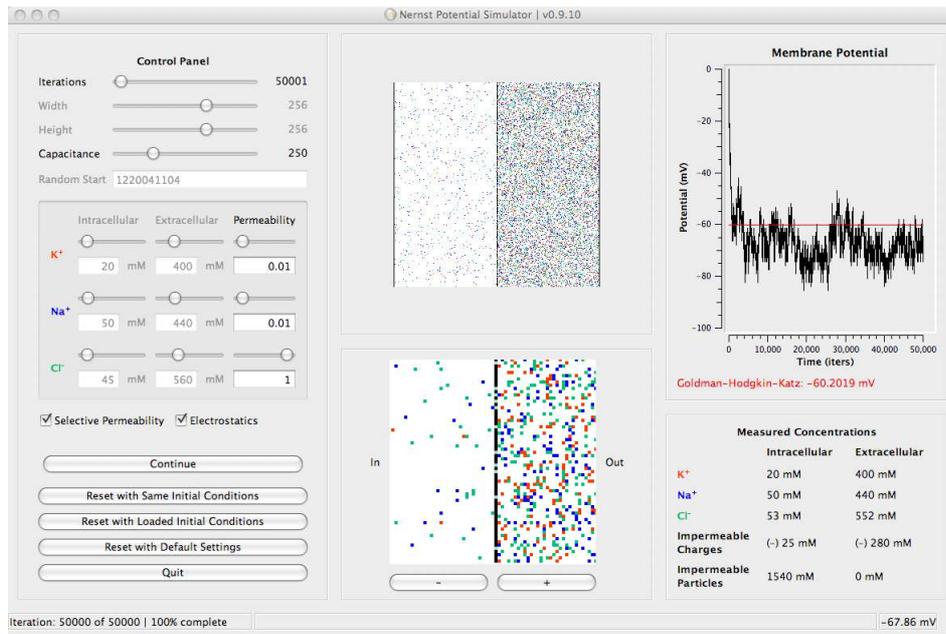
The values I chose to solve the problem are shown in the screen shots below. During the resting potential, permeability to Cl⁻ is high and to K⁺ is low; at the peak of the action potential, permeability to K⁺ is high and to Cl⁻ is low.

Here is the system during the resting potential:

Here is the system at the peak depolarization of the action potential:

Problem 4

A neurons internal and external concentrations of potassium, sodium and chloride have been measured. The internal concentration of potassium ions is 168 mM, of sodium ions is 50 mM, and of



chloride ions is 41 mM. The external concentration of potassium ions is 6 mM, of sodium ions is 337 mM, and of chloride ions is 340 mM. At rest, $P_{K^+} : P_{Na^+} : P_{Cl^-}$ is 1 : 0.019 : 0.381.

(a) (10 points)

Predict V_{rest} from the Goldman-Hodgkin-Katz equation.

Assuming a temperature of 20 °C,

$$\begin{aligned} V_{rest} &= \frac{RT}{F} \ln \frac{P_{K^+} [K^+]_{out} + P_{Na^+} [Na^+]_{out} + P_{Cl^-} [Cl^-]_{in}}{P_{K^+} [K^+]_{in} + P_{Na^+} [Na^+]_{in} + P_{Cl^-} [Cl^-]_{out}} \\ &= 58 \log \frac{1 \times 6 + 0.019 \times 337 + 0.381 \times 41}{1 \times 168 + 0.019 \times 50 + 0.381 \times 340} \\ &= -59.6 \text{ mV} \end{aligned}$$

(b) (10 points)

The resting membrane conductances of a neuron are $g_{K^+} = 0.57 \mu\text{S}$, $g_{Na^+} = 0.11 \mu\text{S}$, and $g_{Cl^-} = 0.32 \mu\text{S}$. What is the resting potential based on the electrical equivalent circuit equation?

$$\begin{aligned} E_{K^+} &= 58 \log \frac{6}{168} = -83.9 \text{ mV} \\ E_{Na^+} &= 58 \log \frac{337}{50} = +48.1 \text{ mV} \\ E_{Cl^-} &= 58 \log \frac{340}{41} = -53.3 \text{ mV} \\ \\ V_{rest} &= \frac{g_{K^+} E_{K^+} + g_{Na^+} E_{Na^+} + g_{Cl^-} E_{Cl^-}}{g_{K^+} + g_{Na^+} + g_{Cl^-}} \\ &= \frac{0.57 \times -83.9 + 0.11 \times 48.1 + 0.32 \times -53.3}{0.57 + 0.11 + 0.32} \\ &= -59.6 \text{ mV} \end{aligned}$$

The two ways of calculating the resting potential are in good agreement.

(c)* (10 points)

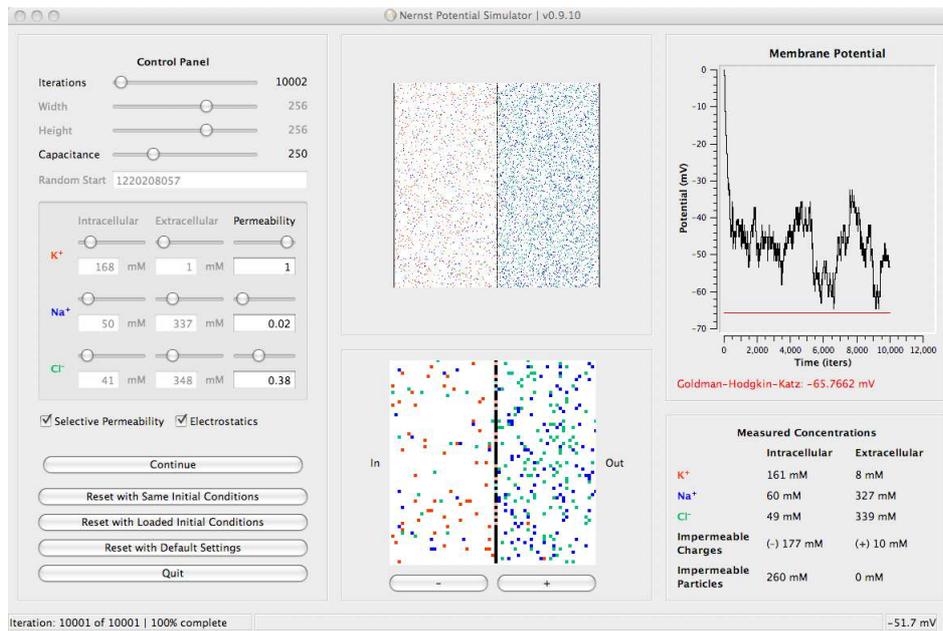
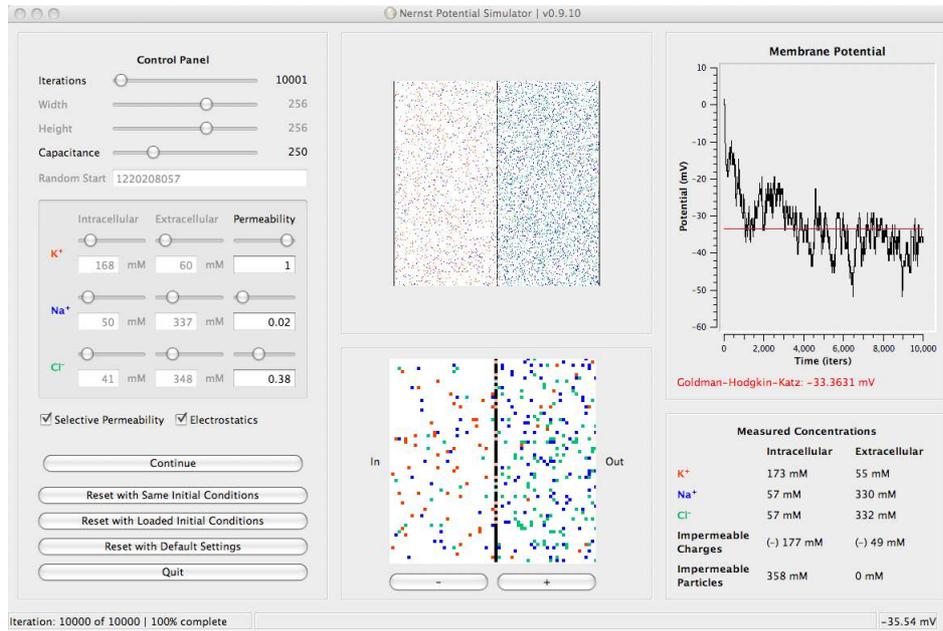
What would be the effect of a tenfold increase in external potassium concentration on the resting membrane potential? What would be the effect of a tenfold decrease in external potassium concentration on the resting membrane potential? Predict the effects using the GHK equation, and confirm your predictions by running the simulation. Explain.

Tenfold increase: external K^+ goes from 6 mM to 60 mM. Result:

$$\begin{aligned} V_{rest} &= 58 \log \frac{1 \times 60 + 0.019 \times 337 + 0.381 \times 41}{1 \times 168 + 0.019 \times 50 + 0.381 \times 340} \\ &= -32.5 \text{ mV} \end{aligned}$$

The increase in external potassium reduces the concentration gradient for potassium, and thus its Nernst potential (which becomes $58 \log (60/168) = -25.9 \text{ mV}$). This reduces the chemical gradient, and thus the opposing electrical gradient at which the flow of potassium will be equal and

opposite is reduced. As a consequence of the high permeability of the cell to potassium, the entire cell depolarizes.



Tenfold decrease: external K^+ goes from 6 mM to 0.6 mM. Result:

$$\begin{aligned} V_{rest} &= 58 \log \frac{1 \times 0.6 + 0.019 \times 337 + 0.381 \times 41}{1 \times 168 + 0.019 \times 50 + 0.381 \times 340} \\ &= -65 \text{ mV} \end{aligned}$$

The decrease in external potassium increases the concentration gradient for potassium. The increased outward flow of potassium ions must be countered by a stronger electrical gradient. Thus, the potassium ions Nernst potential becomes $58 \log (0.6/168) = -141.9 \text{ mV}$, and because of the high permeability to potassium, the entire cell hyperpolarizes (note that the simulation, as currently programmed, cannot represent fractional concentrations, so I have used 1 mM rather than 0.6 mM).