

Stochastic Simulation to Visualize Gene Expression and Error Correction in Living Cells

Kevin Y. Chen

Department of Chemistry, University of Cambridge,
Lensfield Road, Cambridge CB2 1EW, UK

Daniel M. Zuckerman

Dept of Biomedical Engineering, Oregon Health & Science Univ,
Portland, OR 97239,
and Philip C. Nelson

Department of Physics and Astronomy, University of Pennsylvania,
Philadelphia PA 19104

August 26, 2018

Abstract

Stochastic simulation can make the molecular processes of cellular control more vivid than the traditional differential-equation approach by generating typical system histories instead of just statistical measures such as the mean and variance of a population. Simple simulations are now easy for students to construct from scratch, that is, without recourse to black-box packages. In some cases, their results can also be compared directly to single-molecule

experimental data. After introducing the stochastic simulation algorithm, this article gives two case studies, involving gene expression and error correction, respectively. Code samples and resulting animations showing results are given in the online supplement.

1 Introduction

Physical processes unfold over time. Our minds grasp physical mechanisms largely via narrative. So it is not surprising that some of the most vivid physics demonstrations also play out over time. Simulations of physics that unfold over time are similarly powerful; interactive simulations are better; and simulations created by the student can be best of all. This view is gaining ground in introductory courses [1], but the benefits of animated simulation extend farther than this. Here we wish to show that the behavior of strongly nonequilibrium statistical systems can be illustrated via stochastic simulations that are simple enough to serve as undergraduate projects. Recently developed, free, open-source programming resources sidestep the laborious coding chores that were once required for such work. In particular, we believe that the error-correction mechanism known as kinetic proofreading can be more clearly understood when a student views its linear temporal sequence, as opposed to solving deterministic rate equations. Coding this and other simple processes opens the door for the student to study other systems, including those too complex for the rate-equation approach to yield insight.

2 Double-well hopping

2.1 The phenomenon

We tell students that a simple chemical reaction, for example isomerization of a macromolecule, can be regarded as a barrier-passing process. A micrometer-size bead in a double optical trap serves as a mesoscopic model system with this character [2], and it is well worthwhile for students to watch it undergo a few dozen sharp transitions in between episodes of Brownian motion near its two stable positions (see supplementary video 1). A simple model for this behavior states that the hopping transitions occur at random times drawn from an exponential distribution. That is, many rapid transitions are interspersed with a few long pauses.

2.2 Simulation: Waiting times drawn from an exponential distribution

With this physical motivation, students can explore how to generate simulated waiting times of the sort just described. Any computer math system has a pseudorandom number generator that generates floating-point numbers uniformly distributed between 0 and 1. Many students are surprised (and some are intrigued), to learn that applying a nonlinear function to samples from a random variable yields samples with a different distribution, and in particular that $y = -\tau \ln x$ is exponentially distributed, with mean τ , if x is uniform on $(0,1]$ [3].

Starting from that insight, it takes just one line of code to generate a list of simulated waiting times for transitions in a symmetric double well; finding the cumulative sums of that list gives the actual transition times (see supplementary

computer code 1).

The freely accessible VPython programming system (or its Web-based version Glowscript) makes it very easy to create an animation of an object whose spatial position is supplied as a function of time [4]. The only challenging part is to pass from a list of irregularly-spaced transition times to particle positions at each of many (regularly-spaced) video frames (see supplementary computer code 1). The payoff is immediate: Visually, the simulated hopping has a very similar character to the actual Brownian hopping of a bead in a double trap (see supplementary video 2).

2.3 Upgrade to 1D random walks

It may be of interest to make a small modification of the code: Instead of hopping between two wells (reversing direction on every step), consider one-dimensional diffusion on a symmetric many-well potential, for example, one of the form $U(x) = \sin(x)$. In such a potential, for each transition the system must also decide whether to increase or decrease a “position” coordinate. The resulting random walk will display the same long-time scaling behavior as any unbiased 1D walk, but with trajectories that undergo hops at random times, not periodic steps as in the simplest realization [3].

3 Birth-death process

We can now generalize from situations with essentially only one kind of transition (or two symmetric kinds), to the more interesting case where several inequivalent choices are possible, and where the relevant probabilities depend on the current

state. This general situation can describe a chemical reaction that requires, and depletes, molecules of some substrate.

Most science students know that living cells synthesize each of their messenger RNAs (mRNAs) from a single copy (or a small fixed number) of the corresponding gene. Even a constitutive (unregulated) gene must wait for the transcription apparatus to arrive, bind, and begin transcription [5]. We consider a situation in which that apparatus is in short enough supply that this waiting is the primary determinant for the initiation of transcription. Once a mRNA transcript has formed, it has a limited lifetime until it is degraded by other cellular machinery. We assume that this process, too, relies on chance encounters with degradation enzymes. Moreover, each of many species of mRNA must all share the attentions of a limited number of degradation enzymes, so each mRNA copy has a fixed probability per unit time to be removed from the system.

The physical hypotheses in the preceding paragraph amount to a model called the “birth-death process,” which has many other applications in physics and elsewhere. As in the 1D walk, we characterize the system’s state by an integer, in this case the population of the mRNA of interest. Synthesis is a transition that increases this number, with a fixed probability per unit time k_s (called the “mean rate” of synthesis). Degradation is a transition that decreases it, with a probability per unit time that is the current population n times another constant k_d (the “rate constant” for degradation).

3.1 Simulation

D. Gillespie extended and popularized a simple but powerful method, the “stochastic simulation algorithm,” for simulating systems of this sort [6]. In the case just described, the algorithm repeatedly executes the following steps (see supplementary computer code 2):

- Determine the probability per time k_{tot} for *any* of the allowed transitions to occur by summing all the mean rates. In a birth-death process, we have $k_{\text{tot}} = k_s + nk_d$.
- Draw a waiting time from the exponential distribution with mean given by the reciprocal of k_{tot} , via the method in Sect. 2.2.
- Determine which of the allowed processes happens at that transition time. In the birth-death process, we make a Bernoulli trial with probability $p = k_s/k_{\text{tot}}$ to increase population n by one, and $1 - p$ to decrease it.
- Update n and repeat.

The beauty of this algorithm, besides its correctness [7], is that no computation is wasted on time steps at which nothing happened: By definition, there is a state transition at *every* chosen time.

3.2 Convergence to the continuous, deterministic approximation

Students will probably find it reasonable that, when n is sufficiently large, we may neglect its discrete character. Students who have been exposed to probability ideas may also find it reasonable that in this case, the relative fluctuations of n from one

realization to the next will be small, and so n effectively behaves as a continuous, deterministic variable, subject to the differential equation $dn/dt = -k_d n + k_s$. That equation predicts exponential relaxation from an initial value n_0 to the steady value $n_* = k_s/k_d$ with e-folding time k_d :

$$n(t) = n_* + (n_0 - n_*)e^{-k_d t}. \quad (1)$$

The simulation bears out this expectation (Fig. 1a,b).

Actually, mRNA populations in living cells are often *not* large. Nevertheless, although individual realizations of $n(t)$ may differ significantly, the *ensemble average* of many such trajectories does follow the prediction of the continuous/deterministic idealization (Fig. 1c). Within individual cells, there will be significant deviation around that mean behavior (Fig. 1c again). In particular, the “steady” state will have fluctuations of n that follow a Poisson distribution (Fig. 1d). That key result is more memorable for students when they discover it empirically in a simulation than it would be if they just watched the instructor prove it with abstract mathematics (by solving a master equation [3]).

State fluctuations of the sort just mentioned may suffice to pop a more complex system out of one “steady” state and into a very different one. Indeed, even the simplest living cells do make sudden, random state transitions of this sort. Such unpredictable behavior, not seen in the differential-equation approach, seems to be useful to bacteria, implementing a population-level “bet hedging” strategy [8, 9, 10].

A real bacterium is not simply a beaker of reagents. Bacteria periodically divide, partitioning a randomly chosen subset of each mRNA species into each daughter cell. That extra level of realism is hard to introduce into an analytical model, but

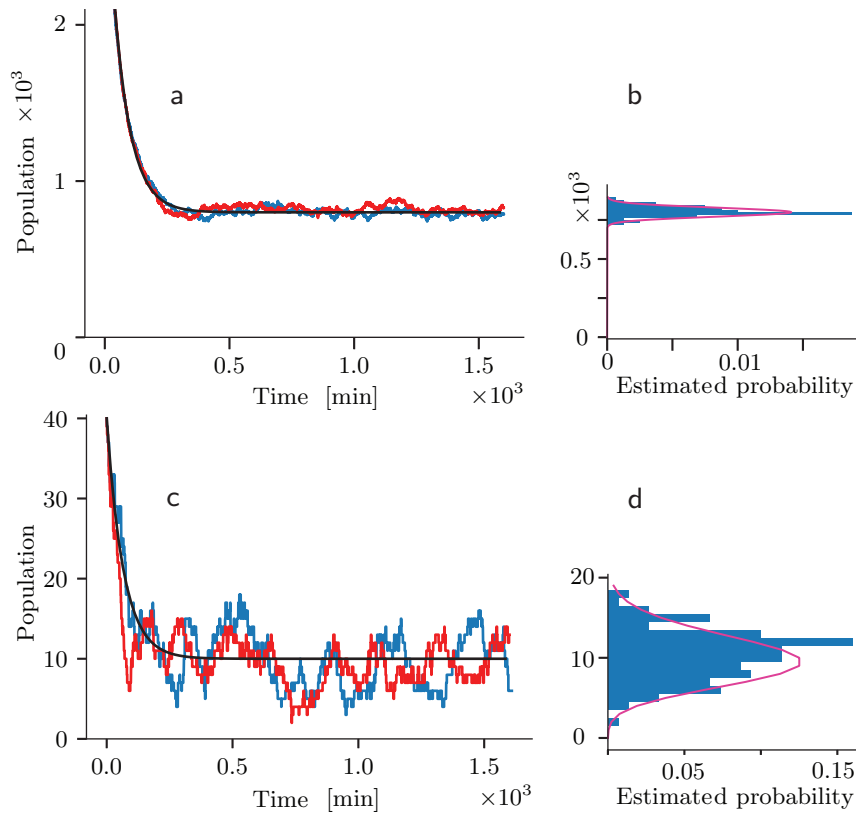


Figure 1: Behavior of a birth-death process. (a) The *bumpy traces* show two examples of simulated time series with $k_s = 12 \text{ min}^{-1}$, $k_d = 0.015 \text{ min}^{-1}$, and hence $n_* = 800$. The initial population was $n_0 = 3200$. The *smooth trace* shows the exponential relaxation predicted by the continuous, deterministic approximation (Eq. 1). (b) After the system comes to steady state, there is a tight distribution of n values across 100 runs of the simulation (*bars*). The *curve* shows the Poisson distribution with mean n_* for comparison. (c,d) The same but with $k_s = 0.15 \text{ min}^{-1}$ and $n_0 = 40$. Although individual instances (runs of the simulation) deviate strongly from the continuous, deterministic approximation, nevertheless the sample mean of the population $n(t)$ over 150 runs does follow that prediction. The distribution of steady-state fluctuations is again Poisson (*curve* in d). (See also [3].)

straightforward in a simulation. The results are similar to the ones just described, with a larger effective value of the clearance rate constant [3].

3.3 Upgrade to cover bursting processes

Bacteria are supposedly simple organisms. The birth-death process is simple, too, and it fits with the cartoons we see in textbooks. So it is interesting to follow the recent discovery that the model makes quantitative predictions for mRNA production that were experimentally disproven [11, 12, 3].

For example, recent advances in single-molecule imaging permit the direct measurement of $n(t)$ in individual cells, and disproved the model's prediction that the distribution of n in the "steady" state should be Poisson (Fig. 2c). Researchers found, however, that a simple modification of the birth-death model could accommodate these and other discrepant data. The required extension amounts to assuming that mRNA transcripts are generated in *bursts*, that the bursts themselves are initiated with a fixed probability per unit time, and that once initiated, a burst is also terminated with a fixed probability per unit time. Although this "bursting model" has two additional parameters compared to the original birth-death model, nevertheless it was overconstrained by the experimental data, so its success was a nontrivial test [12, 3]. Remarkably, detailed biochemical mechanisms for this behavior were found only some years after its indirect inference [13, 14, 15, 16], an important lesson for students to appreciate.

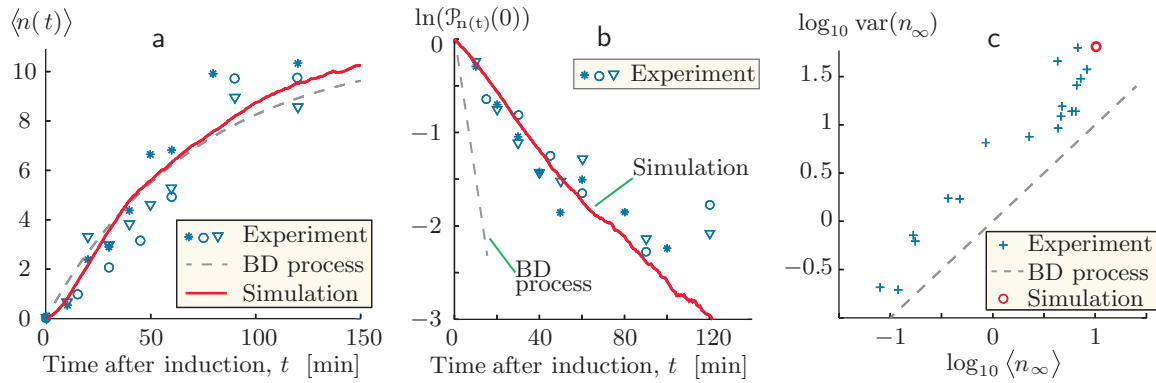


Figure 2: Indirect evidence for transcriptional bursting. (a) *Symbols*: Time course of the number of mRNA transcripts in a cell, $n(t)$, averaged over 50 or more cells in each of three separate, identical experiments. Data from the three trials are shown with three different symbols. All of the cells were induced to begin gene expression at time zero. The *dashed curve* shows a fit of the birth-death (BD) process (Eq. 1) to data, determining the apparent synthesis rate $k_s \approx 0.15/\text{min}$ and clearance rate constant $k_d \approx 0.014/\text{min}$. The *solid curve* shows the corresponding result from a computer simulation of the bursting model discussed in Sect. 3.3 (see also [3]). (b) Semilog plot of the fraction of observed cells that have zero copies of mRNA versus elapsed time. *Symbols* show data from the same experiments as in (a). *Dashed line*: The birth-death process predicts that initially $\mathcal{P}_{n(t)}(0)$ falls with time as $\exp(-k_s t)$, where k_s has the value found by fitting the data in (a). (Degradation was negligible in this experiment.) The experimental data instead yield initial slope $-0.028/\text{min}$. *Solid line*: Computer simulation of the bursting model. (c) Experiments were performed at each of many different levels of gene induction. For each level, a *cross* shows the variance of the late-time mRNA population n_{∞} versus its sample mean. This log-log plot of the data shows that they fall roughly on a line of slope 1, indicating that the Fano factor $(\text{var } n)/\langle n \rangle$ is roughly a constant. The simple birth-death process predicts that this constant is equal to 1 (*dashed line*), because mean equals variance for any Poisson distribution, but the data instead give the value ≈ 5 . The *circle* shows the result of the same simulation shown in (a,b). (Figure adapted from [3]; experimental data from [11].)

4 Kinetic proofreading

4.1 Word model

Ask a student, “What is the big secret of life?” and the answer will probably be “DNA,” or perhaps “evolution by natural selection.” Indeed, DNA’s high, but not perfect, degree of stability underlies life’s ability to replicate with occasional random modifications. But it is less well appreciated that the *stability* of a molecule of DNA does not guarantee the *accuracy* of its replication and transcription. There is another big secret here, just as essential to life as the well known ones. In fact, a wide range of molecular recognition events must have extremely high accuracy for cells and their organisms to function. Think of our immune cells, which must ignore the vast majority of antigens they encounter (from “self”), yet reliably attack a tiny subpopulation of foreign antigens differing only slightly from the self.

Translation of mRNA into proteins is an emblematic example of such a puzzle. It is true that artificial machines now exist that can read the sequence of mRNA. Then another artificial machine can take the resulting sequence of base triplets, decode it, and synthesize a corresponding polymer of amino acids (a polypeptide), which in some cases will then fold into a functional protein without further help. But the cells in our bodies, and even bacteria, do these jobs reliably *without* those huge and expensive machines, despite the incessant nanoscale thermal motion!

Merely intoning that a wonderful molecular machine called the ribosome accomplishes this feat doesn’t get us over the fundamental problem: At each step in translation, the triplet codon at the ribosome’s active site fits one of the many

available transfer RNA (tRNA) species *somewhat* better than it fits the other 19 options. But the binding energy difference, which quantifies “somewhat better,” only amounts to two or three hydrogen bonds. This translates into a fraction of time spent bound to the wrong tRNAs that is about 1/100 times as great as the corresponding quantity for the correct amino acid [17]. If the fraction of incorrect amino acids incorporated into a polypeptide chain were that high, then *every* protein copy longer than a few hundred amino acids would be defective!

In fact, the error rate of amino acid incorporation is more like 10^{-4} . The fact that this figure is so much smaller than the one seemingly demanded by thermodynamics remained puzzling for decades. After all, the ribosome is rather complicated, but it is still a nanoscale machine. *Which* of its features could confer this vast improvement in accuracy?

J. Hopfield and J. Ninio proposed an elegant physical mechanism, based on a known but seemingly pointless feature of the ribosome [18, 19, 17]. To explore it, we begin by paraphrasing a metaphor due to U. Alon [20]. Imagine that you run an art museum and wish to find a mechanism that picks out Picasso lovers from among all your museum’s visitors. You could open a door from the main hallway into a room with a Picasso painting. Visitors would wander in at random, but those who do not love Picasso would not remain as long as those who do. Thus, the concentration of Picasso lovers in the room would arrive at a steady value (with fluctuations, of course) that is enriched for the desired subpopulation.

To improve the enrichment factor further, you could hire an employee who occasionally closes the door to the main hallway, stopping the dilution of your

enriched group by random visitors. Then open a new exit doorway onto an empty corridor. Some of the trapped visitors will gratefully escape, but die-hard Picasso lovers will still remain, leading to a second level of enrichment. After an appropriate time has elapsed, you can then reward everyone still in the room with, say, tickets to visit the Picasso museum in Paris.

The original authors realized that in the ribosome, the initial, reversible binding of a tRNA was followed by a transformation analogous to closing the door in the preceding metaphor. This transformation involved hydrolysis of a GTP (guanosine triphosphate) molecule complexed with the tRNA, and hence it was nearly *irreversible*, due to the highly nonequilibrium concentration of GTP, compared to the hydrolysis products GDP and P_i (inorganic phosphate). Such hydrolysis reactions were well known to supply the free energy needed to drive otherwise unfavorable reactions in cells, but here their role is more subtle.

Hopfield and Ninio were aware that after the hydrolysis, incorporation of the amino acid was delayed and could still be preempted by unbinding of the tRNA complex. The existence of this pathway had previously seemed wasteful: An energy-rich GTP had been “spent” without anything “useful” (protein synthesis) being done. On the contrary, however, the authors argued that this second step implemented the mechanism in the art museum metaphor, giving the ribosome an independent second chance to dismiss a wrong tRNA that accidentally stayed bound long enough to progress to this stage. After all, spending some GTPs may be a modest price to pay compared to creating and then having to recycle an entire defective protein.

Hopfield coined the name “kinetic proofreading” for this mechanism, but we will refer to it as the “classic Hopfield–Ninio” (HN) mechanism because the original term is somewhat misleading. In chemical reaction contexts, a “kinetic” mechanism generally implies bias toward a product with lower activation barrier, even if it is less stable than another product with higher barrier. This preference is most pronounced at high, far-from-equilibrium catalytic rates [21]. In contrast, the classic HN proofreading model involves two sequential thermodynamic (quasiequilibrium) discriminations. Moreover, these discriminations take place prior to reading even the very next codon, in contrast to editorial proofreading, which generally happens after an entire manuscript is written. (Our choice of term also distinguishes the classic scheme from later models that are sometimes also called “kinetic proofreading.”)

The qualitative word-model given earlier in this section may seem promising. But the corresponding kinetic equations make for difficult reading and understanding. Better intuition could emerge from a presentation that stays closer to the concrete ideas of discrete actors randomly arriving, binding, unbinding, and so on, visibly implementing the ideas behind the “museum” metaphor. The following sections will argue that stochastic simulation can realize that goal. In a nutshell,

An effectively irreversible step, or at least a step far from equilibrium, gives rise to enhanced accuracy. The free energy of GTP hydrolysis is the price paid for this accuracy.

It would also be valuable to confirm a key result of the analytic approach, which predicts that the enhancement of accuracy depends on GTP, GDP, and P_i being

held far from chemical equilibrium, so that the hydrolysis step is nearly irreversible (the “door shuts tightly” in the museum metaphor). In fact, the model predicts *no enhancement* of accuracy when this chemical driving force is low [18]. Far from equilibrium, however, the predicted error fraction can be as low as the *square* of the equilibrium value (or even a higher power if multiple rounds of sequential testing are employed).

4.2 A single ribosome in a bath of precursors

This section’s goal is to formulate the word-model of Sect. 4.1 in the context of mRNA translation, then set up a stochastic simulation (see also [22]). Later sections will show how students can explore the expectations raised at the end of the preceding section.

We will assume that a single ribosome is complexed with a single mRNA and has arrived at a particular codon. This complex sits in an infinite bath containing several free, dissolved species at fixed concentrations (Fig. 3a):

- C denotes correct tRNA (that is, the species that matches the codon currently being read), loaded with the corresponding amino acid. We will neglect the possibility of a tRNA being incorrectly loaded; accurate loading is the concern of another proofreading mechanism that we are not studying now [24, 25].
- W is similar to C, but refers to the wrong tRNA for the codon under study.
- Some reactions form complexes of tRNA with guanosine phosphates: C·GTP, C·GDP, W·GTP, and W·GDP. (For simplicity, we suppress any mention of

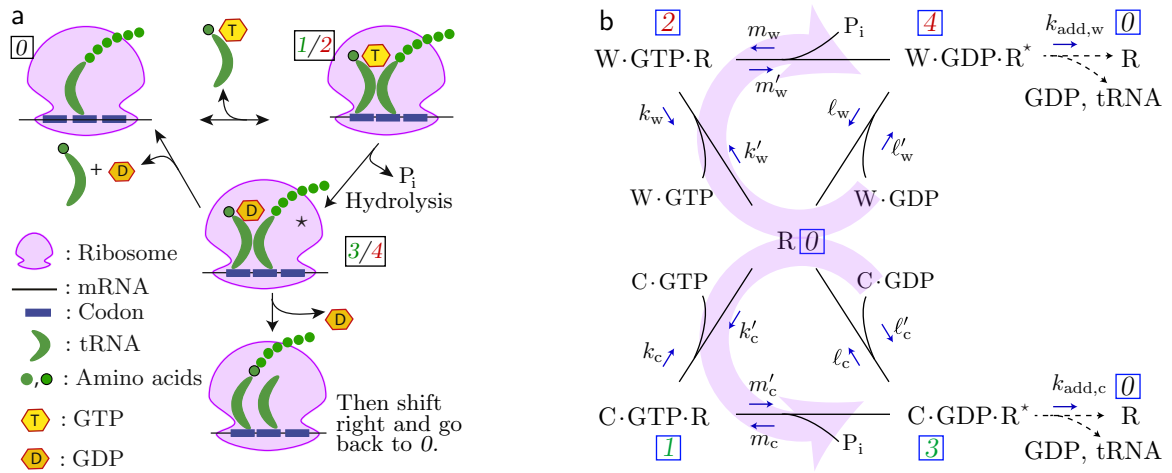


Figure 3: Two representations of the classic Hopfield–Ninio mechanism. (a) Traditional cartoon expressing the catalytic cycle of the ribosome (after [23]). (b) Corresponding kinetic diagram. The large pale arrows indicate the net circulation in each cycle under cellular conditions, where GTP is held out of equilibrium with GDP and P_i . The symbol R denotes a ribosome complexed with mRNA; R^* is the corresponding complex “activated” by GTP hydrolysis. At far right, R indicates the ribosome with one additional amino acid added to the nascent polypeptide chain. The classic Hopfield–Ninio proofreading model assumes that unbinding rates k_c and ℓ_c are smaller than their mismatched counterparts k_w and ℓ_w , but that other constants are all equal for correct and wrong tRNA.

elongation factors, one of which, “EF-Tu,” is also included in these complexes but is only implicit in the classic HN mechanism.)

Fig. 3b denotes the ribosome-mRNA complex by R . In state 0, this complex is not bound to any tRNA. (More precisely, no tRNA is bound at the “A” site of the ribosome; a previously bound tRNA, together with the nascent polypeptide chain, is bound at another site (Fig. 3a), which we do not explicitly note.) Surrounding this state, Fig. 3b shows four other states 1–4 in which the ribosome is bound to the complexes introduced earlier. The upper part of the figure describes wrong tRNA binding and possible incorporation; the lower part corresponds to the correct tRNA. Horizontal arrows at the top and bottom denote hydrolysis of GTP, which is coupled to a transformation of the ribosome into an activated state, R^* .

Although any chemical reaction is fundamentally reversible, under cellular conditions the concentration ratio $[P_i][C\cdot GDP]/[C\cdot GTP]$ is far below the equilibrium value, so that the reactions in Fig. 3b are predominantly in the direction shown by the pale arrows. This was one of the conditions in Hopfield’s original proposal. (Sect. 4.6 will explore relaxing it.)

Again, we are assuming that a *single* ribosome bounces around this state diagram in the presence of fixed concentrations of feedstocks either imposed in vitro by the experimenter or supplied by a cellular milieu. There are two ways to “exit the museum exhibit by the second door”: After hydrolysis, the ribosome can reject its tRNA-GDP complex with probability per unit time ℓ . Or, with probability per unit time k_{add} it can add its amino acid to the nascent polypeptide, translocate the tRNA to the second binding site, and eject any tRNA already bound there. Either way, the

main binding site becomes vacant and, for the purposes of this state diagram, the ribosome returns to state 0.

Supplementary computer code 3 implements a Gillespie simulation on the five states of the ribosome (Fig. 3b).

4.3 Visualization of the simulation results

To keep the project modular, we constructed a simulation code that writes its state trajectory to a file. A second code then reads that file and creates a visual output. The first of these codes operates similarly to Sect. 3.1, but with a four-way choice of what transition to make after each waiting interval. The second code can be almost as simple as the one described in Sect. 2.2. However, students with more time (perhaps in a capstone project) can make a more informative display with a reasonable additional effort, as follows.

The supplementary videos not only show the state that is current at the end of each video frame; they also animate the pending arrivals of new complexes that are about to bind and the departures of old ones that have unbound without incorporation. By this means, the videos give a rough sense of the “narrative” in the trajectory being shown. These improvements are not difficult to add once the basic code is working. Alternatively, students can construct the basic version, then be shown these videos.

The exponential distribution of waiting times implies that there will be episodes with several events happening rapidly, interspersed with long pauses. For this reason, it is useful to view the simulation in two ways: Once with a shorter time

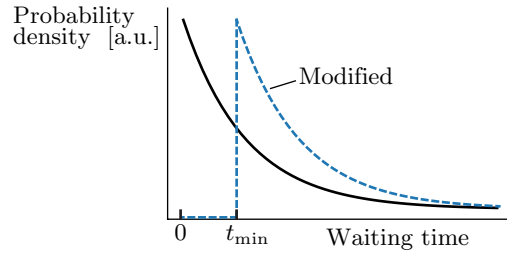


Figure 4: Modified waiting times. *Solid line*: Exponential distribution of waiting times. *Dashed line*: Shifted exponential distribution obtained by adding the constant t_{\min} to each sample.

step that resolves most individual events but covers only a limited time interval (supplementary video 3), and then with a coarser time step to see the entire synthesis trajectory (supplementary video 4).

We also found it useful (solely for visualization purposes) to alter the distribution of waiting times in a simple way that relieves visual congestion without, we think, too much damage to the realism of the simulation. Our modification, shown in the supplementary videos, was to add a small fixed delay, for example one half of one video frame, to every transition waiting time (Fig. 4).

4.4 Classic Hopfield–Ninio model

Following Hopfield, we initially assume that the rate constant for incorporation, k_{add} , is the same regardless of whether the tRNA is correct or incorrect. We also suppose that the binding rates $k'_c = k'_w$ and $\ell'_c = \ell'_w$ also have this property; for example, all of them may be diffusion-limited [17]. Only the *unbinding* rates differ

in the classic HN model:

$$k_w = \phi_{-1}k_c \text{ and } \ell_w = \phi_3\ell_c.$$

Here $\phi_{-1} = \phi_3 \approx 100$ is the preference factor for unbinding the wrong tRNA (relative to the correct one). Again following Hopfield, we will also take the hydrolysis rate constants to be equal: $m'_w = m'_c$ (and $m_w = m_c$). To visualize wrong incorporations within a reasonable time frame, we raised the probability of incorrect choices: The preference ratios ϕ_{-1} and ϕ_3 were lowered from their realistic value of 100 to just 5. Other values in column 4 of Table I were loosely inspired by rate constants estimated from experimental data in a simplified form (see Appendix A.1). (Sect. 4.5 will follow the experimental values more closely.) These effective rate constants are either a constant probability per unit time (unbinding and hydrolysis) or else a probability per unit time with the substrate concentration already lumped in (binding and condensation). The values we chose were appropriate for the concentrations of reactants present in the experiment.

Supplementary videos 3–4 show the resulting behavior. Perhaps the most important impression we get from viewing these animations is that *the cell is a busy place*. The riot of activity, the constant binding events that end with no “progress” (and often not even GTP hydrolysis), are hallmarks of chemical dynamics that are hard to appreciate in textbook discussions, yet vividly apparent in the simulation. This is especially apparent in supplementary video 4, which shows a typical run of 25 amino acid incorporations. Because there are many unproductive binding and unbinding events in the simulation, not every event is shown in detail in video 4. However, focusing on the GDP-tRNA rejections shows that more correct tRNAs

<i>Description</i>	<i>Symbol</i>	<i>Name in code</i>	<i>Classic HN</i>	<i>Realistic model</i>	<i>Equilibrium</i>
binding GTP complex, s^{-1}	k'_c	<code>kc_on</code>	40	40	40
unbinding GTP complex, s^{-1}	k_c	<code>kc_off</code>	0.50	0.5	0.5
binding GDP complex, s^{-1}	ℓ'_c	<code>lc_on</code>	0.001	0.001	0.26
unbinding GDP complex, s^{-1}	ℓ_c	<code>lc_off</code>	0.085	0.085	0.085
hydrolysis and P_i release, s^{-1}	m'_c	<code>mhc</code>	0.01	25	0.01
condensation/ P_i binding, s^{-1}	m_c	<code>msc</code>	0.001	0.001	0.26
binding GTP	$k'_w = \phi_1 k'_c$	<code>kw_on</code>	$\phi_1 = 1$	0.68	1
unbinding GTP	$k_w = \phi_{-1} k_c$	<code>kw_off</code>	$\phi_{-1} = 5$	94	5
binding GDP	$\ell'_w = \phi_{-3} \ell'_c$	<code>lw_on</code>	$\phi_{-3} = 1$	0.0027	1
unbinding GDP	$\ell_w = \phi_3 \ell_c$	<code>lw_off</code>	$\phi_3 = 5$	7.9	5
hydrolysis	$m'_w = \phi_2 m'_c$	<code>mhw</code>	$\phi_2 = 1$	0.048	1
condensation	$m_w = \phi_{-2} m_c$	<code>msw</code>	$\phi_{-2} = 1$	1	1
incorporation, s^{-1}	$k_{add,c}$	<code>kaddC</code>	0.01	4.14	0.01
incorporation	$k_{add,w} = \phi_{add} k_{add,c}$	<code>kaddW</code>	$\phi_{add} = 1$	0.017	1

Table 1: Illustrative values for the rates shown in Fig. 3b. The third column gives variable names used in supplementary computer code 3. See the Appendix for discussion of the numerical values. The fifth column follows [23, 26], who refer to our ϕ_i as f_i . The last column uses rates from the HN model, but with hydrolysis and incorporation modified to satisfy equilibrium.

than incorrect tRNAs make it past GTP hydrolysis, and that the few incorrect tRNAs that do make it past are quickly rejected in the second proofreading step. In the instance shown, only one incorrect amino acid was incorporated out of 25 incorporations, much lower than the $1/5$ error rate expected from single step equilibrium binding. Supplementary video 3 provides a more detailed look at this process. The videos also show clearly the jerky, nonuniform progress of synthesis, with some amino acid incorporations happening after much longer delays than others. That feature is by now well documented by single-molecule experiments.

A typical run created a chain of 100 amino acids, of which 6 were wrong. This error rate of $\approx (6/100) = 0.06$ is far smaller than the naive expectation of $1/\phi_3 = 0.20$. This is the essence of the classic HN mechanism; we see it taking shape in the animation, as many wrong tRNA complexes bind but are rejected, either prior to or after GTP hydrolysis. We also see many *correct* complexes bind and get rejected, before or after GTP hydrolysis. This is the price paid for accuracy in the classic HN proofreading model. The error rate in the simulation is slightly larger than $1/(\phi_3)^2 = 1/5^2 = 0.04$, however. The discrepancy is expected, because the limiting value $1/(\phi_3)^2$ is only achieved in the limit as the incorporation and hydrolysis catalytic rates are sent to zero [26, 18].

4.5 More realistic model

Much has been learned about ribosome dynamics after Hopfield's and Ninio's original insights [27, 28]. We now know that each step in our model consists of substeps. For example, GTP hydrolysis is subdivided into GTPase activation

followed by actual hydrolysis, the latter step probably depends on a rearrangement of “monitoring bases” in the ribosomal RNA, and so on [29].

The model studied in Sect. 4.4 was designed to show the HN mechanism in its “classic,” or pure, form, and how it can enhance fidelity even without help from the effects just described. For example, we assumed that the only dependence on right versus wrong tRNA was via unbinding rates. Indeed, such dependence was later seen at the single-molecule level [30]. But it now appears that some of the forward rates also depend on the identity of the tRNA [31, 32, 33], an effect sometimes called “internal discrimination.” In the limit that the ribosome uses only internal discrimination (activation barrier heights of correct and incorrect tRNA binding differ and the equilibrium constants are the same), minimum error is obtained at fast catalytic rates [34]. This is in contrast to the HN scheme, which achieves minimum error as catalytic rate tends to zero.

In our stochastic simulation model, it is straightforward to add internal discrimination effects by altering rate constants (Table I column 5). See Appendix A.2 for discussion of the values.

Supplemental video 5 shows that the ribosome with experimentally measured rates can be faster and more efficient than a ribosome with only classic HN proof-reading. We see both a bias for correct tRNA binding/hydrolysis and a bias for rejection of wrong tRNAs before GTP hydrolysis. Of the 26 correct tRNA binding events in this run, 25 resulted in successful incorporation. This is compared to the fraction $24/10245 = 0.002$ of productive correct tRNA binding events in a typical run of the classic HN ribosome simulation. In addition, of the 30 incorrect tRNA

binding events on the realistic ribosome, all 30 resulted in rejection. The error fraction of the ribosome with realistic rates is also more accurate than the classic HN ribosome. Of 10 000 amino acids simulated, 18 wrong amino acids were incorporated (error fraction of 0.0018), compared to 6/100 for the classic HN ribosome. This simulated error fraction of 0.0018 is consistent with the analytic prediction of 0.0017 from first-passage times [26].

However, the simulated ribosome with in-vitro measured rates is still not as fast or accurate as the real *E. coli* ribosome in vivo, which translates at 15–20 amino acids per second with an error rate of 1/10 000 [35]. Thus, the realistic ribosome likely evolved to combine HN proofreading (quasiequilibrium, energetic proofreading) with internal discrimination (unequal forward rates) to optimize speed, efficiency, and accuracy [36]. Despite this, simulating the classic HN model of proofreading is still a valuable exercise for students. By visualizing discrimination via only a difference in unbinding rates, students see the minimal components necessary to attain high accuracy in a broad class of biological reactions. Also, the classic HN mechanism illustrates an essential part of biological proofreading which fundamentally relies on non-equilibrium physics.

There is also recent evidence pointing to *two* kinetic proofreading steps, that is, two sequential, nearly irreversible steps each of which can be followed by unbinding of tRNA [37, 38]. Our simulation could be extended to include such effects, whereas analytic methods would quickly become intractable. Finally, additional interesting steps arise during “translocation,” in which the previous tRNA, from which the nascent peptide chain has been released, and the current tRNA, now carrying that

chain with an additional amino acid, are both shifted one step inside the ribosome, freeing the binding site so that the entire cycle can begin again (Fig. 3a). Because this step is not related to accuracy, we have simplified by omitting it from our model.

4.6 Role of thermodynamic driving force

For comparison, we return to the classic Hopfield–Ninio model, this time operating at nearly equilibrium concentrations of GTP, GDP, and P_i to demonstrate the importance of the “one-way door” (the GTP hydrolysis step). Table I column 6 summarizes the rates for this undriven model (see Appendix A.3).

With these rates, the reaction still creates a chain, because we assumed a fixed probability per time to irreversibly add an amino acid whenever the ribosome visits its activated state. But this time a typical run gave 17 errors in a chain of length 100, illustrating the significance of the thermodynamic driving force in reducing the error rate. This error rate of about 0.17 is consistent with the Michaelis-Menten predicted error of $1/\phi_3 = 0.20$, which is the lowest the error can be for a classic HN model in equilibrium conditions [18].

Analysis of the events in the simulation showed that of the 17 wrong amino acids incorporated, 10 were through direct binding of GDP·tRNA. Thus, for many amino acids, the first discrimination step was bypassed, resulting in the high error rate observed in the simulation.

To gain more insight into the role of the irreversible GTP hydrolysis step, some students may wish to rerun the simulation with different incorporation and hydrolysis rates. For example, a simulation with $m'_c = 25$ and $k_{\text{add}} = 4.14$ results in a

simulation with many tRNAs flipping between GDP and GTP states, another way in which the two discrimination steps become coupled into one.

5 Conclusion

The models described here show fairly elementary physical principles that lie at the heart of cell biology. Specifically, gene expression and kinetic proofreading are two important, fundamental topics that are well within reach of undergraduates.

A module that introduces stochastic simulation need not dominate a semester course: One class week is enough for the first exposure. Indeed, the entire simulation plus visualization in supplementary computer code 1 consists of just *seven short lines* of code, and yet it creates a valuable educational experience not available in a static textbook. Moreover, the opening material is not specifically biological in character; it can serve as a stepping stone to more complex simulations relevant for a variety of courses.

Acknowledgments

We are grateful to Ned Wingreen and Anatoly Kolomeisky for correspondence and to Bruce Sherwood for help with software. This work was partially supported by the United States National Science Foundation under Grants PHY-1601894 and MCB-1715823. Some of the work was done at the Aspen Center for Physics, which is supported by NSF grant PHY-1607611, and at the Physical Biology of the Cell summer school (Marine Biological Laboratory).

A Choice of illustrative parameters

A.1 Classic HN proofreading (Sect. 4.4)

Here we outline our choice of rate parameters in Table I, column 4. Rather than pursue these biochemical details, students can simply be told that these are interesting values.

Zaher and Green measured $k_{\text{add},c} \approx 4.14 \text{ s}^{-1}$ and $m'_c \approx 25 \text{ s}^{-1}$ [31]. Interestingly, with those forward rates and the classic HN model's stipulation that only the unbinding rates differ, we found in the stochastic simulation that protein synthesis had a high error rate. For example, using the parameter values shown Table I, column 4 but with the two forward rates above, a typical run gave 49 errors out of 1000 amino acids simulated.

In fact, this breakdown of kinetic proofreading was already predicted in Hopfield's original work, which pointed out that the error rate will only approach $1/\phi_3^2$ if additional conditions are met:

$$m'_c \lesssim k_c \text{ and } k_{\text{add},c} \lesssim \ell_c.$$

This condition gives the two binding/unbinding steps, which are both quasi-equilibrium, time to reject a wrong aa-tRNA before hydrolysis or incorporation; it was not satisfied for the parameter values just mentioned. Thus, if the wrong aminoacyl-tRNA (aa-tRNA) bound to the ribosome, there was a low probability it would unbind before either GTP hydrolysis or incorporation.

Thus, for the purpose of illustrating a pure HN model of proofreading, we modified the experimental values of $k_{\text{add},c}$ and m'_c to the ones shown in Table I.

We also required illustrative values of m_c and ℓ'_c , which were not measured experimentally by Zaher and Green. However, these values are constrained by thermodynamic consistency [39], which requires that even if the reaction is run far from equilibrium, nevertheless the reactions must be *capable* of creating an equilibrium state. When a reaction graph contains a closed loop (cycle), as ours does, this condition requires that

$$k'_{c,\text{eq}}m'_{c,\text{eq}}\ell_{c,\text{eq}} = k_{c,\text{eq}}m_{c,\text{eq}}\ell'_{c,\text{eq}} \quad \text{equilibrium} \quad (2)$$

Assuming for simplicity that substrate-dependent reactions are first order gives

$$\ln \frac{k'_c m'_c \ell_c}{k_c m_c \ell'_c} = \ln \frac{[\text{C}\cdot\text{GDP}]_{\text{eq}}}{[\text{P}_i]_{\text{eq}}[\text{C}\cdot\text{GTP}]_{\text{eq}}} + \ln \frac{[\text{C}\cdot\text{GTP}]}{[\text{P}_i][\text{C}\cdot\text{GDP}]} \quad (3)$$

The corresponding equation for wrong tRNAs is of the same form, except with wrong tRNA rate constants used. Dividing Eq. 3 for the wrong tRNA by the equation for the right tRNA gives the consistency condition that $\phi_1\phi_2\phi_3 = \phi_{-1}\phi_{-2}\phi_{-3}$. The values in Table I column 4 satisfy this condition.

Adamczyk and Warshel calculated the free energy change for GTP hydrolysis on free EF-Tu via a molecular dynamics simulation as $\Delta G'^0 \approx -18$ kcal/mol [40], which we used to calculate the first term on the right. For the second term on the right side of Eq. 3, $[\text{C}\cdot\text{GTP}]$ was taken from Zaher and Green since the simulation's rate constants were based on those experiments. In Zaher and Green's experiments, EF-Tu·GTP complexes were incubated with $0.5 \mu\text{M}$ of aa-tRNA for 15 minutes before injection into a stopped-flow instrument. The $[\text{C}\cdot\text{GDP}]$ and $[\text{P}_i]$ can then be estimated by using

$$[\text{C}\cdot\text{GDP}] = [\text{P}_i] = [\text{C}\cdot\text{GTP}]_{\text{ini}}(1 - e^{-k_{\text{cat}}t}) \approx [\text{C}\cdot\text{GTP}]_{\text{ini}}k_{\text{cat}}t, \quad (4)$$

where t is the incubation time of 15 minutes and k_{cat} is the rate of GTP hydrolysis on free EF-Tu-aa-tRNA complexes. Fasano et al. calculated $k_{\text{cat}} \approx 5.56 \cdot 10^{-5} \text{ s}^{-1}$ [41]. Thus, $[\text{C} \cdot \text{GDP}] = [\text{P}_i] = 0.025 \mu\text{M}$ and $[\text{C} \cdot \text{GTP}] = 0.475 \mu\text{M}$.

Using these non-equilibrium concentrations and the assumption that $m_c = \ell'_c$, Eq. 3 yields the common value of $2.36 \cdot 10^{-12} \text{ s}^{-1}$. A similar argument gives the same values for m_w and ℓ'_w . For the purposes of our simulation, however, these values are essentially zero; over the limited duration of our simulation the corresponding transitions don't occur. We replaced them all by another small value, 0.001 s.

As a further confirmation that the rates chosen are appropriate for a pure Hopfield scheme, we also checked the following conditions described by Hopfield [18]:

- $m'_c \lesssim k_c$.
- $k_{\text{add}} < \ell_c$.
- $m'_c k'_w / (m'_c + k'_w) > \ell'_w$.
- $m_w \lesssim \ell_w, k_{\text{add}} \lesssim \ell_w$.

A.2 More realistic model (Sect. 4.5)

To model a ribosome as it might work in the cell, in this simulation we more closely followed the rates measured by Zaher and Green (Table I, column 5). The rate constants m'_c and $k_{\text{add},c}$ were set to their experimentally measured values of 25 s^{-1} and 4.14 s^{-1} , respectively. The ϕ_i values were also modified to match experimental measurements [26, 23]. ϕ_2 and ϕ_3 were not measured in Zaher and Green, so we follow Banerjee et al., who chose $\phi_2 = 1$ and got ϕ_3 by the thermodynamic

consistency condition discussed in Appendix A.1. Banerjee et al. also chose $\phi_{-2} = 1$, then got ϕ_{-3} by imposing the consistency condition: $\phi_1\phi_2\phi_3 = \phi_{-1}\phi_{-2}\phi_{-3}$.

A.3 Model with no chemical driving force (Sect. 4.6)

To illustrate the importance of the GTP hydrolysis step in the HN model, rate values were chosen to simulate protein synthesis without GTP, GDP, and P_i concentrations out of equilibrium. The rates chosen for the equilibrium proofreading model were similar to those in the classic HN model, except that the equilibrium consistency condition, Eq. 2, was used to calculate ℓ'_c and m_c . In addition, the ϕ s were kept the same as in Appendix A.1, that is, satisfying Hopfield's condition that only the unbinding rates differs between wrong and right tRNAs. Using the same assumption that $m_c = \ell'_c$ for simplicity yielded $m_c = \ell'_c = 0.26 \text{ s}^{-1}$, higher than in the classic HN model.

Supplementary online material

Videos:

1. ChenVideo1-BeadJump.mov at **TEMPORARY LINK**
<http://www.physics.upenn.edu/biophys/PMLS/Media/brownian/BeadJump.mov>:
Video micrograph of a micrometer-scale bead in thermal motion, hopping between the minima of a symmetric double-well potential field created by an optical trap. Courtesy Adam Simon (see [2]).
2. ChenVideo2-Flipper.mov at **TEMPORARY LINK** <https://vimeo.com/269210861>:

Hopping in a symmetric, two-state well. Animation generated by supplementary computer code 1, reminiscent of the behavior seen in video 1.

3. ChenVideo3-HNslow.mp4 at **TEMPORARY LINK** <https://vimeo.com/283759767> :

Proofreading in the classic Hopfield–Ninio model. Animation generated by supplementary computer code 4, from a simulated trajectory generated by supplementary computer code 3 with parameters discussed in Sect. 4.4. In this and the following videos, the left green number is the frame number and the right green number is the state’s sequence number. Green objects represent correct transfer-RNA/amino acid complexes; red objects represent incorrect choices. Cylindrical objects represent complexes containing GTP, arriving from solution. A sphere in the center position represents a complex containing GDP. The growing chain of spheres on the right represent amino acids incorporated into the nascent polypeptide. The simulation was done assuming cellular conditions, that is, GTP far from equilibrium with GDP and inorganic phosphate. Only a short initial time interval is shown (the first 3% of full simulation), so that tRNA binding, unbinding, hydrolysis, and incorporation events can be seen in detail. The frame rate was chosen to be 15 frames/second, and the total duration was chosen to be 10 000 s (150 000 frames). The actual simulation time is 26 692 seconds (41 185 states). The value of t_{\min} (see Sect. 4.3) was chosen to be $(0.5)(1/15) \text{ s} = 0.033 \text{ s}$.

4. ChenVideo4-HNfast.mp4 at **TEMPORARY LINK** <https://vimeo.com/284065985> :

The same as in video 3, but speeded up and covering the incorporation of 25

amino acids. The error rate for binding was deliberately taken to be unrealistically large, $1/\phi_3 = 0.20$, in order to generate some errors in a short simulation. The frame rate was chosen to be 30 frames/second, and the total duration was chosen to be 600s (18 000 frames). The actual simulation time is 26 692 seconds (41 185 states). The value of t_{\min} (see Sect. 4.3) was chosen to be $(0.42)(1/30)\text{ s} = 0.014\text{ s}$.

5. `ChenVideo5-InternalDisc.mp4` at **TEMPORARY LINK** <https://vimeo.com/283760592>

:

Proofreading in a model with internal discrimination. The same as in video 3, but with parameters discussed in appendix A.2. The same chain length of 25 was used and no wrong amino acids are incorporated because the realistic ribosome has an error rate much lower than $1/25$ (Sect. 4.5). The frame rate was chosen to be 15 frames/second, and the total duration was chosen to be 150 s (2250 frames). The actual simulation time is 8.44 seconds (138 states). The value of t_{\min} (see Sect. 4.3) was chosen to be $(0.5)(1/15)\text{ s} = 0.033\text{ s}$.

Computer codes: These codes are also available from https://github.com/NelsonUpenn/PNelson_code.

They were written in Python version 3. The `.ipynb` files run in the Jupyter Notebook environment and use the VPython 7 package [4]. The `.py` files may be run in any Python implementation. One way to obtain VPython and Jupyter is from the free Anaconda distribution (<http://anaconda.com>): After regular installation, issue the command

```
$ conda install -c vpython vpython
```


to your operating system's command shell. For more details see [42].

1. `ChenCode1-flipper.ipynb`: Simulate two-state transitions.
2. `ChenCode2-transcrip.py`: Simulate birth-death process.
3. `ChenCode3-riboProof.py`: Simulate tRNA selection in a ribosome.
4. `ChenCode4-kproofBackend.ipynb`: Display the result of Code 3 as an animation.

References

- [1] R W Chabay and B A Sherwood. *Matter and interactions*. Wiley, New York, 4th edition, 2015.
- [2] A Simon and A Libchaber. Escape and synchronization of a Brownian particle. *Phys. Rev. Lett.*, 68(23):3375–3378, 1992.
- [3] P Nelson. *Physical models of living systems*. W. H. Freeman and Co., New York, 2015.
- [4] Glowscript VPython and VPython 7. <http://www.glowscript.org/docs/VPythonDocs/index.html>, accessed 5/2018.
- [5] B Alberts, D Bray, K Hopkin, A Johnson, J Lewis, M Raff, K Roberts, and P Walter. *Essential cell biology*. Garland Science, New York, 4th edition, 2014.
- [6] D T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.*, 22(4):403–434, 1976.

- [7] D T Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2361, 1977.
- [8] P J Choi, L Cai, K Frieda, and X S Xie. A stochastic single-molecule event triggers phenotype switching of a bacterial cell. *Science*, 322(5900):442–446, 2008.
- [9] A Eldar and M B Elowitz. Functional roles for noise in genetic circuits. *Nature*, 467(7312):167–173, 2010.
- [10] M E Lidstrom and M C Konopka. The role of physiological heterogeneity in microbial population behavior. *Nat. Chem. Biol.*, 6(10):705–712, 2010.
- [11] I Golding, J Paulsson, S M Zawilski, and E C Cox. Real-time kinetics of gene activity in individual bacteria. *Cell*, 123(6):1025–1036, 2005.
- [12] I Golding and E C Cox. Chapter 8: Spatiotemporal dynamics in bacterial cells: Real-time studies with single-event resolution. *Meth. Cell Biol.*, 89:223–251, 2008.
- [13] S Chong, C Chen, H Ge, and X S Xie. Mechanism of transcriptional bursting in bacteria. *Cell*, 158(2):314–326, 2014.
- [14] S A Sevier, D A Kessler, and H Levine. Mechanical bounds to transcriptional noise. *Proc. Natl. Acad. Sci. USA*, 113(49):13983–13988, 2016.
- [15] S A Sevier and H Levine. Properties of gene expression and chromatin structure with mechanically regulated elongation. *Nucleic Acids Res*, 46(12):5924–5934, 2018.

- [16] A Klindziuk and A B Kolomeisky, “Theoretical Investigation of the Transcriptional Bursting Mechanism: a Multi-State Approach.” *J. Phys. Chem.*, submitted.
- [17] R Phillips, J Kondev, J Theriot, and H Garcia. *Physical biology of the cell*. Garland Science, New York, 2d edition, 2012.
- [18] J J Hopfield. Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity. *Proc. Natl. Acad. Sci. USA*, 71(10):4135–4139, 1974.
- [19] J Ninio. Kinetic amplification of enzyme discrimination. *Biochimie*, 57(5):587–595, 1975.
- [20] U Alon. *An introduction to systems biology: Design principles of biological circuits*. Chapman and Hall/CRC, Boca Raton FL, 2006.
- [21] P Sartori and S Pigolotti. Kinetic versus energetic discrimination in biological copying. *Phys. Rev. Lett.*, 110(18):188101, 2013.
- [22] D M Zuckerman. Physical lens on the cell: Active (“kinetic”) proofreading. <http://physicallensonthecell.org/cell-biology-phenomena/active-kinetic-proofreading>, accessed 5/2018.
- [23] K Banerjee, A B Kolomeisky, and O A Igoshin. Elucidating interplay of speed and accuracy in biological error correction. *Proc. Natl. Acad. Sci. USA*, 114(20):5183–5188, 2017.

- [24] J J Hopfield, T Yamane, V Yue, and S M Coutts. Direct experimental evidence for kinetic proofreading in amino acylation of tRNA^{Ile}. *Proc. Natl. Acad. Sci. USA*, 73(4):1164–1168, 1976.
- [25] T Yamane and J J Hopfield. Experimental evidence for kinetic proofreading in the aminoacylation of tRNA by synthetase. *Proc. Natl. Acad. Sci. USA*, 74(6):2246–2250, 1977.
- [26] K Banerjee, A B Kolomeisky, and O A Igoshin. Accuracy of substrate selection by enzymes Is controlled by kinetic discrimination. *J. Phys. Chem. Lett.*, 8(7):1552–1556, 2017.
- [27] M V Rodnina. Mechanisms of decoding and peptide bond formation. In M V Rodnina, W Wintermyer, and R Green, editors, *Ribosomes: Structure, function, and dynamics*, chapter 16, pages 199–212. Springer, New York, 2011.
- [28] I Bahar, R L Jernigan, and K A Dill. *Protein actions: Principles and modeling*. Garland Science, New York, 2017.
- [29] P Satpati, J Sund, and J Aqvist. Structure-based energetics of mRNA decoding on the ribosome. *Biochemistry*, 53(10):1714–1722, 2014.
- [30] S C Blanchard, R L Gonzalez, H D Kim, S Chu, and J D Puglisi. tRNA selection and kinetic proofreading in translation. *Nat. Struct. Mol. Biol.*, 11(10):1008–1014, 2004.
- [31] H S Zaher and R Green. Hyperaccurate and error-prone ribosomes exploit distinct mechanisms during tRNA selection. *Mol. Cell*, 39(1):110–120, 2010.

- [32] A Prabhakar, J Choi, J Wang, A Petrov, and J D Puglisi. Dynamic basis of fidelity and speed in translation: Coordinated multistep mechanisms of elongation and termination. *Protein Science*, 26(7, SI):1352–1362, 2017.
- [33] M V Rodnina, N Fischer, C Maracci, and H Stark. Ribosome dynamics during decoding. *Phil. Trans. R. Soc. Lond. B, Biol. Sci.*, 372(1716):20160182, 2017.
- [34] C H Bennett. Dissipation-error tradeoff in proofreading. *BioSystems*, 11(2-3):85–91, 1979.
- [35] R Milo and R Phillips. *Cell Biology by the numbers*. Garland Science, New York, 2016.
- [36] R Rao and L Peliti. Thermodynamics of accuracy in kinetic proofreading: dissipation and efficiency trade-offs. *J. Stat. Mech.*, article P06001, 2015.
- [37] K-W Jeong, U Uzun, M Selmer, and M Ehrenberg. Two proofreading steps amplify the accuracy of genetic code translation. *Proc. Natl. Acad. Sci. USA*, 113(48):13744–13749, 2016.
- [38] J Chen, J Choi, S E O’Leary, A Prabhakar, A Petrov, R Grosely, E V Puglisi, and J D Puglisi. The molecular choreography of protein synthesis: translational control, regulation, and pathways. *Q. Rev. Biophys.*, 49:e11, 2016.
- [39] T L Hill. *Free energy transduction and biochemical cycle kinetics*. Springer, New York, 1989.

- [40] A J Adamczyk and A Warshel. Converting structural information into an allosteric-energy-based picture for elongation factor Tu activation by the ribosome. *Proc. Natl. Acad. Sci. USA*, 108(24):9827–9832, 2011.
- [41] O Fasano, E De Vendittis, and A Parmeggiani. Hydrolysis of GTP by elongation factor Tu can be induced by monovalent cations in the absence of other effectors. *J. Biol. Chem.*, 257(6):3145–3150, 1982.
- [42] J M Kinder and P Nelson. *A student's guide to Python for physical modeling*. Princeton Univ. Press, Princeton NJ, updated edition, 2018.