

# Engines of Parsimony: Part III

## Performance Trade-offs for Reversible Computers Sharing Resources

Hannah Earley\*

Department of Applied Mathematics and Theoretical Physics, University of Cambridge




Figure 1: This paper concerns interactions between reversible Brownian computational particles and populations of other particles such as species involved in chemical reactions or shared memory resources. Populations of particles are subject to well known statistical mechanical laws, and these laws can cause problems when computational particles attempt to interact with these populations.

### Lay Summary

In recent years, unconventional forms of computing ranging from molecular computers made out of DNA to quantum computers have started to be realised. Not only that, but they are becoming increasingly sophisticated and have a lot of potential to influence the future of computing. One interesting class of unconventional computers is that of reversible computers, which includes quantum computers. Reversible computing—wherein state transitions must be invertible, and therefore must conserve information—is largely neglected outside of quantum computing, but is a promising avenue for realising substantial gains in computational performance and energy efficiency.

Reversible computing is of great interest because, unlike traditional irreversible computing which is subject to a lower bound on energy cost per operation [1, 2], there is no lower bound to the energy usage of a reversible computer. It might be imagined that one could go as far as to get *processive*<sup>1</sup> computation for free using reversibility, such as with the scheme of Fredkin and Toffoli [3], but unfortunately mounting evidence [4, 5] suggested otherwise. Indeed, in Part I [6] of this series we proved that free processive computation is not possible, with instead the adiabatic regime being the best attainable. Still, for a given energy budget

\*h.earley@damtp.cam.ac.uk,  orcid.org/0000-0002-6628-2130

<sup>1</sup>That is, computation which proceeds forward at a finite, non-zero rate!

and spatial size these adiabatic reversible computers have asymptotically better performance scaling than irreversible computers subject to the same constraints.

This analysis neglected, however, to consider interactions between subunits of these computers. In Part II [7] we analysed how computational subunits interact and communicate with one another. One may also wish to consider how computational subunits interact with other, non-computational subsystems, which in general will be governed by the laws of thermodynamics and statistical mechanics. When the free energy available to the computational subunits is limiting, this becomes a challenging endeavour as perturbations to these subsystems may have a corresponding entropy cost for which our available free energy is insufficient.

Our Brownian computational subunits—themselves being small and particulate—will likely lack adequate internal memory for many programs of interest; as such, it can be expected that a realistic Brownian system will make heavy use of a shared pool of memory resources that can be dynamically distributed amongst the computational subunits according to their need. We therefore dedicate a section of this paper to analysing this particular case, and find that seemingly any possible scheme will not work unless additional free energy is supplied.

The same conclusion applies to any chemical reaction we wish to drive away from equilibrium. In order to make these interactions viable, we propose an abstract chemical scheme to dynamically infer how much free energy is needed, and to supply it by accumulating just enough from the (weak) ambient free energy. The scheme takes any reaction, arbitrarily far from equilibrium, and sequesters some of the reactants and products such that the extant reactants and products are in equilibrium. The scheme rapidly adjusts to changes in the equilibrium position, and makes use of molecular computation for its inference algorithm. Unfortunately there is a cost in applying this scheme: the reaction is slowed by a factor inversely proportional to the ‘computational bias’,  $b \ll 1$  (the net proportion of computational transitions that are successful). In fact, this overhead is comparable to that found for synchronisation interactions in Part II.

### Technical Abstract

This paper concludes a three-Part series on the limits the laws of physics place on the sustained performance of reversible computers. Part I concerned aggregate performance in terms of computational operations per unit time, but neglected to consider interactions among computational sub-units or between computational sub-units and shared resources such as memory or chemical species. Part II extended the analysis to consider the former set of interactions. In this Part we extend the analysis to consider the latter set, with a particular focus on resource distribution in the first half. It is found that most schemes imaginable fail to function effectively in the limit of vanishing ‘computational bias’  $b$ , which measures the net fraction of transitions which are successful, and falls as the system grows in size. Driving thermodynamically unfavourable reactions, such as resource distribution, is a very general problem for such systems and can be solved by supplying a sufficient excess of free energy. We propose a scheme to dynamically supply enough free energy for a given reaction, automatically and rapidly adapting to changes in the disequilibrium state of said reaction—including the case when the favourable reaction direction switches. The overhead of this scheme is no worse than the overhead found in Part II for communicating reversible computers under the same regime.

# 1. Introduction

This is the final Part of a three-Part series investigating the influence of the laws of physics on computation, particularly as they apply to large reversible Brownian computers (see Definition 1.1). It has long been known that reversible computation—computation which conserves information and whose every transition is invertible—allows for substantially better energy efficiency in principle. In particular, whilst irreversible computation is subject to a ‘Landauer limit’ bounding the energy cost of each computational operation from below [1], reversible computation has no such lower bound. Not only can reversible computers achieve better energy efficiency, but their total computational rate can be shown to scale asymptotically faster with spatial size than irreversible computers occupying the same region (see Frank [5] for some early results). In fact, we proved [6] that the adiabatic regime of reversible computation is the best possible in any quantum or classical computer, in which the total computational rate (in terms of primitive computational operations per unit time) scales as  $\sim \sqrt{AV} \sim R^{5/2}$  where  $A$  is the convex bounding surface area of the computer,  $V$  its volume, and  $R$  its radius, compared to  $\sim A$  for an irreversible computational system. Concretely, these are given by

$$\frac{4}{\sqrt{3}}\pi R^{5/2} \cdot \sqrt{\frac{\phi\lambda}{2k_B T}} \quad \text{and} \quad 4\pi R^2 \cdot \frac{\phi}{k_B T \Delta I},$$

respectively for a spherical region, where  $R$  is the radius of the sphere,  $\phi$  is the power supply<sup>2</sup> per unit surface area,  $k_B$  is Boltzmann’s constant,  $T$  is the surface temperature,  $\Delta I$  is the average amount of information erased per irreversible operation, and  $\lambda$  is a rate per unit volume of gross reversible computation (whether forwards or backwards).

As discussed in Part II [7], the net computational rate is not the only performance metric of interest. Therefore, in Part II we evaluated computational performance of reversible Brownian computers (and argued for applicability of the results to reversible computers in general) whose individual computational components were permitted to interact with one another, i.e. to communicate and to perform concurrent/parallel computation. A key property of these interactions is that computational ‘particles’ can be uniquely referenced, and therefore the traditional thermodynamic principles that apply to chemical systems were inapplicable. In this Part, we now consider the interface between traditional chemical species, of which there may be a variable number of identical units or particles, with the computational particles. For convenience, we shall refer to these as *klona* and *mona* respectively, following the formulation introduced in Part II. Recalling the definitions, we call particles which are uniquely distinguishable or addressable *mona* (sg. *monon*), after the Greek for unique, μοναδικός. Similarly, we call indistinct particles belonging to a species *klona* (sg. *klonon*). Klona, may be likened to traditional chemical species, such as glucose and ATP. Mona are less obvious; an example would be a 150 mL solution containing a 1 mmol dm<sup>−3</sup> concentration of random nucleic acids of length 80 nucleotides<sup>3</sup>. We can therefore

<sup>2</sup>The function of the power supply is to dissipate the entropy generated by the system, and thus the input power must have low entropic content in order to perform thermodynamic *work*. Therefore,  $\phi$  measures the rate of supply of *free energy*, which is generally lower than the gross energy supply rate.

<sup>3</sup>A nucleic acid, such as the DNA molecules that comprise our genome, is a polymer formed from an arbitrary sequence of four nucleotide monomers; random nucleic acids are easily obtained from DNA synthesis companies.

say that Part II studies mona-mona interactions, and Part III mona-klona interactions. We do not cover klona-klona interactions as these are well understood within the domains of chemistry and statistical physics.

Interactions between mona and klona are of great practical importance for Brownian computers; whilst mona by themselves may be capable of arbitrary computation, this computation only becomes useful when used to actuate some other system or when it otherwise produces some observable output. Examples may include the activation of a fluorophore upon reaching some logical condition, or directing the translation of a particular messenger RNA into a protein. Moreover, interaction with klona may be essential for computation: general computation requires access to an unbounded amount of memory, but by their very nature mona are limited in size and thus memory. Therefore it is important for mona to be able to recruit more memory as required, and release it for use by other mona when the need passes. Resources other than memory may also be considered, such as structural monomers for use in constructing and deconstructing molecular machines or computers.

**Definition 1.1** (Brownian Computers). A Brownian computer is a region of space (typically bounded or otherwise confined) containing particles. The definitions of particles and space are not important, and they may be complex macromolecules or simple elementary particles or even abstract entities in an abstract topological space. These particles are generally static in isolation, but when two or more particles collide they may alter their states during the interaction, in much the same way as a molecule of ATP and a molecule of  $\text{H}_2\text{O}$  may collide, react, and beget the products ADP,  $\text{P}_i$  and  $\text{H}^+$  (particle number need not be conserved). These interactions must be *reversible*, such that the products may collide and reform the reactants. To enable these interactions, particles must be endowed with velocities, and these velocities are assumed to follow a thermal distribution.

A Brownian computer is characterised by being able to interpret the current state of the particles as a computational state: perhaps some particles are ‘computers’ in their own right, and thus the evolution of their state over time is isomorphic to some model of computation, or perhaps the joint state of a number of particles must be considered. The former, wherein the computational particles are mona, is more powerful, whereas the latter approach, representing computational state with klona, typically results in the entire reaction volume being dedicated to a single program and thus limits the computational power. Multiple programs can be run if their components are orthogonal in the sense that they don’t interact with components of a different program, but designing orthogonal klona is challenging and limited. An alternative mitigation approach is to use multiple bounded sub-volumes containing just the klona for that particular program; these sub-volumes may then be considered as abstract ‘particles’ or mona, as can the combination of orthogonal sub-klona (forming a virtual sub-volume, if you will). Computational mona can be driven forward in computational phase space by a *computational bias* (Definition 1.2, below).

An example would be that of molecular computers such as can be found in the field of DNA

---

As such, there are  $4^{80} \approx 1.5 \times 10^{48}$  possible distinct 80-mers (nucleic acid of length 80). The birthday paradox says that we need  $1.4 \times 10^{24}$  of these random molecules before it is likely that there are duplicates present. Our 150 mL,  $10 \text{ mmol dm}^{-3}$  solution contains just shy of  $10^{20}$  such molecules, and it can be calculated that the probability that there are any duplicates in the solution is less than 1 in 350 million. Therefore, it is almost certain that every nucleic acid in said solution is a monon.

computing [8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. These consist of molecules of DNA dissolved in water, wherein the ability for nucleic acids to bind to other nucleic acids in a sequence-specific manner<sup>4</sup> is exploited in various different ways to program cascades of reactions encoding arbitrary computation. Those DNA computers using DNA-strand displacement as their mechanism have a klonal representation of computational state, whereas those using the Tile Assembly Model have a monal representation.

**Definition 1.2** (Computational Bias). Recall from Part I [6] that in a Brownian computer, transitions are mediated by systems of klona representing ‘positive’ and ‘negative’ bias. Physically, such klona act as a source of free energy and a real world example is given by ATP and ADP + P<sub>i</sub> in biochemical systems. In the simplest case, there are  $\oplus$  and  $\ominus$  klona such that  $\forall n. |n\rangle + \oplus \rightleftharpoons |n+1\rangle + \ominus$  where  $|n\rangle$  is the  $n^{\text{th}}$  state of some monon, and this is positively biased when the concentration of  $\oplus$  exceeds that of  $\ominus$ . We define the bias in this case to be  $b = ([\oplus] - [\ominus])/([\oplus] + [\ominus])$ .

In this paper we shall be most interested in the case of resource distribution, and will consider various such schemes in Section 2. One approach is a centralised store which mona can access, either by travelling there themselves or by sending a ‘courier’ monon on their behalf to fetch and retrieve (or carry and return) a resource. This approach is undesirable. If the couriers float freely in solution, then reaching the central store and returning will both take a time  $\sim \mathcal{O}(V)$ ; if on the other hand we introduce a lattice as in Part II [7], then the average distance will be proportional to the radius  $R$  and, given  $b \sim R^{-1/2}$ , the average time will be  $\sim \mathcal{O}(b^{-3})$ . As the time for a single net computational step is  $\sim \mathcal{O}(b^{-1})$ , both of these are untenable. Instead we shall seek a decentralised approach by making use of klona; intuitively it must become harder to access a resource as its scarcity increases, but we would like the timescale to do so to not be substantially more costly than  $\sim (b[X])^{-1}$  where  $[X]$  is the available concentration of the resource.


We will unfortunately find that the reactions involved in all possible schemes are unfavourable except in a very limited range of resource availability. In Section 3 we shall return to considering general mona-klona interactions, and demonstrate how mona can dynamically drive any unfavourable reaction with klona by inferring the number of bias tokens required to balance the entropic cost of the reaction. Alas, there will turn out to be an overhead in achieving this even in the case when the driven reaction is at equilibrium and hence cost-free.

## 2. Resource Distribution Schemes


A resource distribution scheme for a resource  $X$  is a set of klona from which  $X$  can be extracted and into which it can be deposited. Formally, resources are conserved discrete quantities analogous to charges in physics. Letting the set of resource species be indexed by  $\mathcal{R}$  and the set of resource klona by  $\Sigma$ , then to each klonon  $\sigma \in \Sigma$  is associated a resource charge  $\vec{q}_\sigma \in \mathbb{N}^{\mathcal{R}}$ : a vector over

---

<sup>4</sup>Nucleic acids are polymers of arbitrary sequences of four monomers. In the case of DNA, these monomers are adenine, guanine, cytosine and thymine or A, G, C and T. The nucleotide monomers of nucleic acids are special in that they are subject to base-pairing rules: A preferentially binds to T and vice-versa, and G preferentially binds to C and vice-versa. The consequence is that a DNA sequence 5’-CTGCGATCGC-3’ will generally only bind stably to its (reverse) complementary sequence 5’-GCGATGCAAG-3’.




(a) A possible state diagram for a monon interacting with a resource pool; the particular resource distribution scheme is left abstract.




(b) An illustration of the statistical states and reactions for the Free Klona resource distribution scheme.

Figure 2: An assortment of illustrative examples of some of the resource distribution schemes for a Brownian amorphous computer as considered in this paper. (continued)



(c) An illustration of the statistical states and reactions for the Bounded Carrier resource distribution scheme.



(d) An illustration of the statistical states and reactions for the Isobaric Compartment resource distribution scheme. The question marks serve to indicate that the apparent isergonicity of the resource reactions is in question, and as we find, is unfortunately not the case.

Figure 2: (continued) An assortment of illustrative examples of some of the resource distribution schemes for a Brownian amorphous computer as considered in this paper.

the field of naturals and indexed by the resource species. In any reaction, the net resource charge must be conserved; for example, for the reaction  $|n\rangle + \sum_{\sigma \in \Sigma} \nu_{\sigma} \sigma \rightleftharpoons |n+1\rangle + \sum_{\sigma \in \Sigma} \nu'_{\sigma} \sigma$  involving the transition of a monon between states  $|n\rangle$  and  $|n+1\rangle$  and the interaction of said monon with resource klona according to the stoichiometries  $\vec{\nu} \rightleftharpoons \vec{\nu}'$ , this conservation law can be expressed as  $\vec{q}_{|n+1\rangle} - \vec{q}_{|n\rangle} = \sum_{\sigma \in \Sigma} (\nu_{\sigma} - \nu'_{\sigma}) \vec{q}_{\sigma}$ .

**Free Klona** The simplest possible scheme is that in which the resource is freely present in solution, as shown in Figure 2b. That is, a system with  $\mathcal{R} = \{X\}$ ,  $\Sigma = \{X\}$  and  $(\vec{q}_X)_X = 1$ . This system will be a viable resource pool so long as the free energy change involved in both extraction and deposition of  $X$  is less than the available computational free energy,  $\log \frac{p}{q} = 2 \operatorname{arctanh} b$ . Note that our notion of ‘free energy’ uses a different convention than is normally taken, namely we are giving the change in information entropy of the universe  $\Delta h$ , whereas normally the free energy change is expressed as  $\Delta G = -kT\Delta h$ . The free energy change vanishes when the system is in equilibrium, which for  $X$  an ideal gas can be calculated (via the Sackur-Tetrode equation) to occur at a concentration

$$[X]_{(\text{eq.})} = e^{5/2} \left( \frac{4\pi m u}{3h_P^2} \right)^{3/2},$$

where  $h_P$  is Planck’s constant,  $m$  is the mass of  $X$ , and  $u$  is its average internal energy. As we interact with the system, the concentration will be drawn away from this equilibrium position and as a result the absolute free energy difference will increase. For this scheme, the free energy difference corresponds to that due to adding or removing a particle of  $X$  and can therefore be identified with the *Chemical Potential*,  $\mu_X$ . In units of information, this can be expressed as  $\mu_X = \log \frac{[X]\gamma_X}{[X]_{(\text{eq.})}\gamma_X}$  where  $\gamma_X$  is the activity coefficient of  $X$  that quantifies the deviation of the properties of  $X$  from an ‘ideal’ substance. If  $Q$  is the concentration of resource  $X$  (here trivially equal to  $[X]$ ) then we can now determine how far from equilibrium the resource system can be brought, and therefore how much of the resource  $Q$  is accessible to the computational mona. For a small change  $\delta Q$ , we find

$$\left| \frac{\delta Q}{Q} \right| < 2b \left| 1 + \frac{\partial \log \gamma_X}{\partial \log Q} \right|^{-1} + \mathcal{O}(b^2).$$

We don’t expect the activity coefficient to depend strongly on  $Q$ , and certainly not to the extent of  $\frac{\partial \log \gamma}{\partial \log Q} = -1$ , which would imply a range of isentropic equilibria, and therefore we can only usefully access a proportion  $\mathcal{O}(b) \ll 1$  of the available resource  $X$  under this scheme.

**Bounded Carriers** Perhaps the primary issue with free klona is that there is a change in particle number, implying a substantial entropy change when resources are traded with the pool. A possible improvement is suggested in ensuring total particle number remains constant by binding the resource to carrier klona, as shown in Figure 2c. This may manifest, for example, as  $\Sigma = \{X^0, X^1\}$  with  $(\vec{q}_{X^0})_X = 0$  and  $(\vec{q}_{X^1})_X = 1$ . More generally, we have polymeric klona  $\Sigma = \{X^k : k \in [0, n] \cap \mathbb{N}\}$  with  $(\vec{q}_{X^k})_X = k$  for some integer  $n > 1$ . Once again, we shall consider a perturbation to the equilibrium distribution of the pool; it will be convenient to introduce



a partition function  $\mathcal{Z}(\beta)$  where the thermodynamic variable  $\beta$  correlates with the amount of resource stored by the pool. The partition function will take the form  $\mathcal{Z} = \sum_{k=0}^n e^{-\varepsilon_k - \beta k}$  where  $\varepsilon_k$  contains all the thermodynamically pertinent information about  $X^k$  klonas, such as its energy and internal degrees of freedom. In fact, we will take  $\varepsilon_k = 0$  as any  $k$ -dependence of the  $\varepsilon_k$  would only serve to reduce the effective value of  $n$  by narrowing the distribution<sup>5</sup>.

The free energy change for a reaction is given by the logarithm of the ratios of reaction rates (see, e.g., Section 4 of Earley [6] for a derivation). Therefore in this case it is given by  $\Delta h = \log \sum_{k=1}^n e^{-\beta k} - \log \sum_{k=0}^{n-1} e^{-\beta k}$ , and so at equilibrium (when  $\Delta h = 0$ ) we find  $\beta = 0$ . For a perturbation from equilibrium, this reduces (exactly) to  $\delta \Delta h = -\delta \beta$ . Now, the average X-charge of each klonon is given by  $\bar{k} = -\partial_\beta \log \mathcal{Z}$ , and the charge density by  $Q = \sum_{k=0}^n k[X^k] \equiv \bar{k}[X]$ . The proportional resource accessibility is then given by

$$\left| \frac{\delta Q}{Q} \right| = \left| \frac{-\delta \beta \text{var } k}{\bar{k}} \right| < \frac{1}{3} b(n+2)$$

using  $\bar{k} = \frac{1}{2}n$  and  $\text{var } k = \frac{1}{12}n(n+2)$ . This suggests taking  $n \gtrsim b^{-1}$  to ensure full resource availability. Unfortunately, doing so would mean the total system charge scales as  $V/b \sim V^{7/6}$  which clearly cannot be sustained in the limit of large systems. One potential resolution is to reduce the carrier concentration to  $[X] \sim b^{1/2}$  and then take  $n \sim b^{-1/2}$  so that the total charge scales with  $V$ ; whilst this would slow down resource reactions commensurately, mona could use this additional time to accumulate  $\sim b^{-1/2}$  tokens in order to raise their free energy budget to  $\Delta h \lesssim b^{1/2}$ , thereby obtaining a resource accessibility of  $|\delta Q/Q| \sim 1$  at a time penalty of  $\sim b^{-3/2}$ .

**Unbounded Carriers** It is also possible to consider polymeric klonas of unbounded size, i.e. the limit  $n \rightarrow \infty$ , in which case the partition function is given by  $\mathcal{Z} = (1 - e^{-\beta})^{-1}$ . It should be noted that in this limit there is no equilibrium distribution of the resources, as  $\mathcal{Z}$  diverges as  $\beta \rightarrow 0$ . In fact, distributions only exist for  $\beta > 0$  whilst the free energy constraint bounds it from above and hence  $0 < \beta < 2b$ . This range corresponds to  $\bar{k} > (2b)^{-2} - \mathcal{O}(1)$  and therefore a total enclosed charge  $\gtrsim V^{4/3}$ ; for lower densities, the only favourable reaction is resource deposition. Therefore, whilst conceptually simpler, unbounded carriers are to be avoided.

**Isobaric Compartments** The problem inherent to seemingly all schemes is that a resource reaction must necessarily reduce the quantity of reactants and increase the quantity of products, and this will in turn bias the resource reactions in the opposite direction per *Le Chatelier's* principle. As our last attempt to circumvent these issues, we introduce a scheme exploiting varying partial volumes in order to maintain equal concentrations of reactants and products. This ‘Isobaric Compartment’ scheme is illustrated in Figure 2d and consists of dividing the system into two compartments by means of a massless and infinitely flexible membrane, impermeable to the carrier species  $X^0$  and  $X^1$ . Referring to the two compartments by  $\Gamma_\oplus$  and  $\Gamma_\ominus$ , we then establish the invariants that  $\Gamma_\oplus$  only contains  $X^1$  and  $\Gamma_\ominus$  only contains  $X^0$ . A monon wishing to exchange

<sup>5</sup>Assuming that the equilibrium distribution is unimodal, this corresponds to a reduction in variance. If the distribution is multimodal then instead there is a reduced ‘local’ variance about each mode.

resources with the pool must first bind to the membrane, and then to maintain the invariant it will translocate the carrier species between the two compartments. At equilibrium, the concentrations of carrier species in each compartment must be equal and thus it would appear that we have been successful in eliminating the free energy cost for this mona-klona interaction.

Unfortunately, whilst the carrier species concentrations are equal, concentration is not the correct quantity for determining the reaction rates and thence the free energy difference. From a collision theory perspective, the reaction rates are given by the frequency with which mona and klona collide (in the correct orientation and conformation). Consider the case where all but one carrier reside within  $\Gamma_{\oplus}$ : if the total number of carriers is  $N$  then the mean volume of  $\Gamma_{\ominus}$  will be  $V/N$  and the mean concentration  $1/(V/N) = N/V$ . This would suggest that a monon bound to the interfacial membrane will encounter the  $X^0$  klona as often as it does an  $X^1$  klona, but this clearly cannot be the case. If we imagine the  $X^0$  klona as residing within a small bubble of volume  $V/N$ , then whilst its local concentration is indeed the same as for the  $X^1$  klona, the bubble itself is free to migrate anywhere along the membrane. Consequently, in a system in which particles can bind to the boundary of the system in addition to exploring their volume, the effective concentration becomes  $N/(\alpha + \gamma)$  where  $\alpha$  is the effective ‘volume’ of the boundary (proportional to its area) and  $\gamma$  is the enclosed conventional volume. For our isobaric system,  $\gamma_{\oplus} = N_{\oplus}\gamma_0$  (and similarly for  $\gamma_{\ominus}$ ) where the constant  $\gamma_0 = V/(N_{\oplus} + N_{\ominus})$  is the volume available to each particle. We can therefore write  $[X^0]_{\text{eff.}} = [X^0]/(\alpha/V + [X^0]\gamma_0)$  where  $[X^0] = N(X^0)/V$  is the ‘true’ concentration; redefining terms, we let  $[X^0]_{\text{eff.}} = [X^0]/(\alpha + [X^0]\gamma)$ .

The system can be generalised to carriers for  $n > 1$  by letting the membrane be semi-permeable to  $\{X^k : 0 < k < n\}$ , as these klona can participate as both reactants and products in either reaction. We can then write  $\mathcal{Z}_{\ominus} = \sum_{k=0}^{n-1} e^{-\beta_{\ominus}k}$  for the partition function of  $\Gamma_{\ominus}$  and similarly for  $\Gamma_{\oplus}$ . The effective concentrations of  $X^k$  in each compartment must be equal for  $0 < k < n$ , i.e.  $(e^{-\beta_{\ominus}k}/\mathcal{Z}_{\ominus})[X]_{\ominus\text{eff.}} = (e^{-\beta_{\oplus}k}/\mathcal{Z}_{\oplus})[X]_{\oplus\text{eff.}}$ . As this must apply for each  $0 < k < n$ ,  $(\beta_{\oplus} - \beta_{\ominus})k$  must be constant with respect to  $k$  and therefore  $\beta_{\oplus} = \beta_{\ominus} = \beta$ . Moreover,  $[X]_{\oplus\text{eff.}}/[X]_{\ominus\text{eff.}} = \mathcal{Z}_{\oplus}/\mathcal{Z}_{\ominus} = e^{-\beta}$ . Equilibrium occurs once again for  $\beta = 0$  with corresponding ‘true’ concentrations  $[X]_{\oplus} = [X]_{\ominus} = \frac{1}{2}[X]_{\text{total}}$ . The free energy difference is then given by

$$\delta\Delta h = -\delta\beta = \frac{\delta[X]_{\oplus}}{\frac{1}{2}[X]} \frac{2\alpha}{\alpha + \frac{1}{2}\gamma[X]}.$$

Finally, the resource charge is  $Q = \bar{k}_{\oplus}[X]_{\oplus} + \bar{k}_{\ominus}[X]_{\ominus}$  and hence the proportional availability is given by

$$\begin{aligned} \left| \frac{\delta Q}{Q} \right| &= \left| \frac{-\delta\beta \frac{1}{2}[X](\text{var}_{\oplus} k + \text{var}_{\ominus} k) + \delta[X]_{\oplus}}{\frac{1}{2}n[X]} \right| \\ &< 2b \left( \frac{\frac{1}{6}(n^2 - 1)}{n} + \frac{\alpha + \frac{1}{2}\gamma[X]}{2\alpha n} \right) \\ &< b \left( \frac{1}{3}n + \frac{1}{6n} + \frac{1}{4\alpha n} \right), \end{aligned}$$

recalling that  $\gamma \equiv 1/[X]$ .

As with the Bounded Carrier scheme we can let  $n \sim 1/b$  to ensure full resource availability, but we can also choose to let  $\alpha \sim b$  with  $n = 1$ . The constant  $\alpha$  is proportional to the ratio of

the membrane area to the system volume. A naive implementation would of course have  $\alpha \sim b^2$  which, whilst ensuring full resource availability, would limit the number of mona able to interact with the resource scheme proportional to  $A$ . Instead, consider replacing the system's volume with a space-filling tubule (or perhaps a tubule-lattice) and partitioning the tubule along its length; this system is still an instance of the Isobaric Compartment scheme, but allows an  $\alpha$  as large as  $\sim 1$  (if we wanted to allow all mona to interact with the resource system simultaneously). By making the tubule diameter  $\sim b^{-1}$ , a value of  $\alpha \sim b$  can be achieved in which case the proportion of mona able to simultaneously interact with the resource system will be  $\sim V^{5/6}$ . This is, however, a time penalty of  $\sim b^{-2}$  and hence is strictly worse than the Bounded Carrier scheme.

### 3. Driving Unfavourable Reactions

We conjecture that the results in the previous section are general; that is, no resource scheme can achieve better resource availability under the same constraints. In order to make use of a greater proportion of the resource pool, more free energy must be supplied and, ideally, the amount of free energy supplied should be the minimum necessary to avoid wasting the supply of negentropy. As alluded to in the introduction, we shall now consider general mona-klona interactions of arbitrary free energy cost.


**Reactions of Known or Bounded Cost** Consider an arbitrary reaction  $\sum_i \nu_i X_i \rightleftharpoons \sum_i \nu'_i X_i$  with forward reaction rate  $\alpha$  and backward rate  $\beta$ . Recalling that the free energy change for the forward reaction is given by  $\Delta h = \log \frac{\alpha}{\beta}$ , it can be seen that  $n > -\Delta h/2 \operatorname{arctanh} b$  computational bias tokens are needed to drive the reaction forward in the case  $\Delta h < 0$ ; similarly, if  $\Delta h > 0$  then  $n > \Delta h/2 \operatorname{arctanh} b$  tokens are needed to drive the reaction backward. In addition, if we wish to couple such an unfavourable reaction to our computational mona then we should ensure by some means that the reaction does not occur in uncoupled isolation, such as raising the activation energy of the reaction or by employing a reaction mechanism that requires interaction with the mona. This kind of coupling abounds in nature, wherein a variety of bias systems are exploited. In fact, the  $\text{ATP} : \text{ADP} + \text{P}_i$  bias system is itself generated from a number of other bias sources. One particular example involves the conversion of the free energy of a transmembrane proton gradient for which  $n \gtrsim 4$ ; the molecular machine responsible for this coupling, ATP synthase, is a molecular motor which translocates 11 protons<sup>6</sup> across the membrane per revolution, in the process converting 3 molecules of ADP to ATP, an entropically unfavourable process.

When  $n$  is known and fixed (or, at least, bounded from above) an optimal approach is provided by the simple scheme that performs  $n$  transitions between each unfavourable reaction. Indexing

---

<sup>6</sup>The exact number varies between species, and even between different organelles of the same species, but is typically between 10 and 14.

the mona states,  $M_k$ , modulo  $n + 1$ , the scheme can be written thus



The reactions between  $M_k$  and  $M_{k+1}$  serve to increase the ratio  $[M_n]/[M_0]$  such that it counteracts the ratio  $\alpha'/\beta' = \alpha/\beta < 1$ . To solve for the steady state dynamics of this system, we write the currents

$$S(M_k \mapsto M_{k+1}) = p\lambda[M_k] - q\lambda[M_{k+1}]$$

$$S(M_n \mapsto M_0) = \alpha'[M_n] - \beta'[M_0]$$

and use the fact that, at steady state, the currents must be equal. Moreover, the current will correspond to the net rate at which the target reaction is driven, and therefore we can find the optimal choice of  $n > |\Delta h/2b|$  to maximise this rate (as for excessively large  $n$ , the chain of intermediate monon states will dilute the proportion of transitions performing the desired reaction).

We first show that the system makes optimal use of the consumed bias tokens by calculating the value of  $n$  for which the current vanishes, and hence for which the reaction is exactly balanced. For vanishing current, the principle of detailed balance applies and the steady state distribution is obtained simply by  $[M_k]/[M_0] = (\frac{p}{q})^k$  and  $[M_0]/[M_n] = \alpha/\beta$ . This yields  $n = \log \frac{\beta}{\alpha} / \log \frac{p}{q} = -\Delta h/2 \operatorname{arctanh} b$ , precisely the threshold value of  $n$ . The attentive reader may point out that this value of  $n$  is not necessarily integral, and therefore a cyclic chain of length  $n + 1$  may not exist; this is true, but we shall be seeking larger values of  $n$  to increase the rate of the reaction and for which picking the nearest integral value shall be acceptable. Alternatively, one could consider the superposition of two such chains for  $\lfloor n \rfloor$  and  $\lceil n \rceil$  in the appropriate proportions.

To solve for non-vanishing current, we perform two telescopic sums

$$nS = \sum_{k=0}^{n-1} S(M_k \mapsto M_{k+1}) \quad \frac{1}{p} S \sum_{k=0}^{n-1} t^{-k} = \sum_{k=0}^{n-1} (t^{-k} \lambda[M_k] - t^{-k-1} \lambda[M_{k+1}])$$

$$= b\lambda[M] + q\lambda[M_0] - p\lambda[M_n], \quad S \frac{p^n - q^n}{p - q} = p^n \lambda[M_0] - q^n \lambda[M_n],$$

where  $t \equiv \frac{p}{q}$ , and then substitute  $S(M_n \mapsto M_0)$  for the current:

$$b\lambda[M] = (p\lambda + n\alpha')[M_n] - (q\lambda + n\beta')[M_0],$$

$$\left( q^n \lambda + \frac{p^n - q^n}{p - q} \alpha' \right) [M_n] = \left( p^n \lambda + \frac{p^n - q^n}{p - q} \beta' \right) [M_0].$$

Hence, solving for  $[M_0]$  and  $[M_n]$ , we can obtain the steady state current thus:

$$S = b\lambda[M] \frac{q^n \beta' - p^n \alpha'}{(q^{n+1} - p^{n+1})\lambda + n(q^n \beta' - p^n \alpha') + (q\alpha' - p\beta') \frac{1}{b}(p^n - q^n)}.$$

After some rearrangement, we find that the current is maximised when

$$(t^{n/2} - t^{-n/2}\gamma)^2 = \frac{q \log t}{b}(\gamma - 1)(t\gamma - 1) + \frac{q\lambda \log t}{\alpha}(t\gamma + 1)$$


where  $\gamma \equiv \beta/\alpha$ . For small bias, this reduces to  $n = 2n_0 + \mathcal{O}(b)$  where  $n_0 = |\Delta h|/2 \operatorname{arctanh} b$  is the threshold value. When  $n$  is large, this gives a current  $S \rightarrow [M] \min(2b^2\lambda\frac{\alpha}{\beta}, 2b\alpha)$ . This assumes that the intermediate monon transitions are not useful; if the intermediate transitions in fact perform useful work then this current can be multiplied by  $n$  and, in the limit of large  $n$ , the current tends to  $b\lambda$ : the net transition rate for uncoupled mona.

**Reactions of Unknown or Unbounded Cost** The preceding scheme suffers from a number of shortcomings due to the fact that  $n$  is fixed; when  $n_0$  is lower than that designed for, the system expends more bias tokens than necessary, whereas when it is greater the system will stall. Moreover the system is asymmetric with respect to direction of the driven reaction: if  $\beta > \alpha$  then the forward reaction should be driven with  $n$  tokens whilst the reverse reaction needs no additional bias, but this complicates the act of running a subroutine in reverse as each bias chain for the forward reaction must be removed, and a bias chain must be introduced for each instance of the reverse reaction. This is further complicated if the intermediate transitions along the chain perform useful computation. Whilst these problems are not insurmountable, it is possible to solve them all with a revised scheme in which  $n$  is permitted to vary.

A scheme with variable  $n$  is incompatible with the specific implementation of that for fixed  $n$ , and so we first migrate the mona chains into a distinct system of klona consisting of the species  $\{X_k^{(n)} : n, k \in \mathbb{N} \wedge 0 \leq k \leq n\} \cup \{X_{-k}^{(-n)} : n, k \in \mathbb{N} \wedge 0 \leq k \leq n\}$  and subject to reactions  $X_k^{(n)} \xrightleftharpoons[\frac{q}{p}]{} X_{k+1}^{(n)}$ , with steady state  $[X_k^{(n)}]/[X_0^{(n)}] = (\frac{p}{q})^k$ . To demonstrate practicability, an abstract molecular realisation is provided in Appendix B. The idea is to sequester some of the reactants and products in an inactive state, with those remaining in an active state being at equilibrium (for an isenthalpic reaction, this manifests as equal concentrations). Representing the reactants by A and the products by B, this effective sequestration behaviour is achieved by the reaction  $M_i + A:X_{n_0}^{(n_0)} \xrightleftharpoons[\frac{q}{p}]{} M_{i+1} + B:X_0^{(n_0)}$  with  $n_0 = -\Delta h/2 \operatorname{arctanh} b$ . For fixed  $n_0$ , this scheme already solves all the aforementioned problems as the complexed system  $A:X_{n_0}^{(n_0)} \rightleftharpoons B:X_0^{(n_0)}$  is an equilibrated abstraction over the underlying  $A \rightleftharpoons B$  reaction, and hence can be coupled directly to the computational mona system.

In the course of system operation, the value of  $n_0$  will almost certainly vary. Introducing the sequestration klona for all values of  $n \in \mathbb{Z}$ , and letting  $\alpha' \equiv \alpha \sum [X_n^{(n)}]$  and  $\beta' \equiv \beta \sum [X_0^{(n)}]$ , we see that the desired equilibrium condition is given by  $\alpha' = \beta'$ . Therefore, if the values of  $\alpha'$  and  $\beta'$  deviate from equilibrium then we should counteract this deviation by perturbing the  $n$ -distribution of the sequestration klona. Ideal kinetics are attained if this  $n$ -distribution converges to  $\alpha' = \beta'$  exponentially fast, i.e.  $\dot{n} \propto \beta' - \alpha'$ . This target kinetics arises because  $[X_n^{(n)}] = (\frac{p}{q})^n [X_0^{(n)}]$  and  $\frac{p}{q} > 1$ , hence increasing  $n$  increases  $\alpha'/\beta'$ . A first attempt at implementation might be given by  $\forall n \in \mathbb{Z}. M_i + A:X_n^{(n)} \xrightleftharpoons[\frac{q}{p}]{} M_{i+1} + B:X_0^{(n-1)}$ , but this doesn't quite replicate the desired kinetics. Instead we separate the concerns of the computational bias and the reaction coupling

thus:



Assuming that at steady state  $[M_{i+\frac{1}{3}}] = [M_{i+\frac{2}{3}}] \equiv \mu$ , the coupled reaction then replicates the kinetics  $\dot{n} = k\mu(\beta' - \alpha')$ . For symmetry, we have introduced two intermediate mona with their inter-mona transitions used for coupling to the computational bias. In order to ensure that only a single token's worth of bias is consumed in the reaction, we also introduce two leak reactions such that the total free energy change is precisely that for a single token:  $\Delta h = \log \frac{p+p}{q+p} + \log \frac{p+q}{q+q} = 2 \operatorname{arctanh} b$ .

To verify that this system has the desired properties, we now solve for the steady state distribution. Employing detailed balance, we find for each  $n$  that  $[X_n^{(n)}]/[X_n^{(n)}]$

$$\frac{[X_n^{(n)}]}{[X_0^{(n-1)}]} = \frac{\beta}{\alpha} \equiv \left(\frac{p}{q}\right)^{n_0} \implies \frac{[X_0^{(n)}]}{[X_0^{(0)}]} = \left(\frac{p}{q}\right)^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}-n_0)^2},$$

where we have used the intra- $X_k^{(n)}$  steady state distribution,  $[X_k^{(n)}]/[X_0^{(n)}] = (\frac{p}{q})^k$ . This steady state distribution resembles a discrete Gaussian in  $n$ , and hence is relatively tightly centered about  $n_0$ ; this is important to ensure good control over the distribution and that perturbing its mean value is practicable. As a sanity check, we can show that this distribution indeed gives  $\beta' = \alpha'$ :

$$\begin{aligned} \frac{\beta'}{\alpha'} &= \left(\frac{p}{q}\right)^{n_0} \frac{\sum_{n \in \mathbb{Z}} \left(\frac{p}{q}\right)^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}-n_0)^2}}{\sum_{n \in \mathbb{Z}} \left(\frac{p}{q}\right)^n \left(\frac{p}{q}\right)^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}-n_0)^2}} \\ &= \frac{\sum_{n \in \mathbb{Z}} \left(\frac{p}{q}\right)^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}+n_0)^2}}{\sum_{n \in \mathbb{Z}} \left(\frac{p}{q}\right)^{-\frac{1}{2}(n-\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}+n_0)^2}} \\ &= 1. \end{aligned}$$

Finally we seek to characterise the rates for the coupled reaction; first, we shall require the weight for each of the  $\{X_k^{(n)} : k\}$  subsystems:

$$\sum_{k=0}^n [X_k^{(n)}] = \frac{[X_0^{(n)}]}{t-1} \begin{cases} t^{n+1} - 1 & n \geq 0 \\ t - t^n & n \leq 0 \end{cases}$$

where  $t = \frac{p}{q}$ . Next, we use  $t^n t^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}-n_0)^2} = t^{-\frac{1}{2}(n-\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}+n_0)^2}$  to find the normalisation,

$$\begin{aligned} [X_0^{(0)}]^{-1} &= 1 + \frac{1}{t-1} [t \sum_{n < 0} - \sum_{n > 0}] t^{-\frac{1}{2}(n+\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}-n_0)^2} \\ &\quad + \frac{1}{t-1} [t \sum_{n > 0} - \sum_{n < 0}] t^{-\frac{1}{2}(n-\frac{1}{2}-n_0)^2 + \frac{1}{2}(\frac{1}{2}+n_0)^2}. \end{aligned}$$

This is non-trivial for general  $n_0$ , so we restrict to the two limiting cases of  $n_0 = 0$  and  $|n_0| \gg 1$ . In the former case, the normalisation simplifies to  $[X_0^{(0)}] = \frac{b}{1+b}$  and so we have  $\beta' = \alpha' = \frac{\alpha b}{1+b} \sum_{n \in \mathbb{Z}} t^{\frac{1}{8} - \frac{1}{2}(n - \frac{1}{2})^2}$ . The sum is bounded from below by the integral  $2 \int_0^\infty t^{-\frac{1}{2}(s + \frac{1}{2})^2} ds$ , and from above by  $2 \int_0^\infty t^{-\frac{1}{2}(s - \frac{1}{2})^2} ds$ , from which we obtain  $\alpha' = \alpha \sqrt{\pi b} + \mathcal{O}(b)$ . In the latter case, suppose  $n_0 \gg 1$  is positive (i.e.  $\beta \gg \alpha$ ). The normalisation will then be dominated by the sums over  $n > 0$  and so has asymptotically exact approximation

$$\begin{aligned} [X_0^{(0)}]^{-1} &= \frac{t}{t-1} \sum_{n \in \mathbb{Z}} t^{-\frac{1}{2}(n - \frac{1}{2} - n_0)^2 + \frac{1}{2}(\frac{1}{2} + n_0)^2} - \frac{1}{t-1} \sum_{n \in \mathbb{Z}} t^{-\frac{1}{2}(n + \frac{1}{2} - n_0)^2 + \frac{1}{2}(\frac{1}{2} - n_0)^2} \\ &= \frac{1}{b} [pt^{\frac{1}{2}(\frac{1}{2} + n_0)^2} - qt^{\frac{1}{2}(\frac{1}{2} - n_0)^2}] \sum_{n \in \mathbb{Z}} t^{-\frac{1}{2}(n + \frac{1}{2} - n_0)^2}, \end{aligned}$$

hence  $\alpha'$  reduces to  $2b\alpha + \mathcal{O}(b^2)$  in the case of  $\beta \gg \alpha$ , or  $2b\beta + \mathcal{O}(b^2)$  in the case of  $\alpha \gg \beta$ . In the ideal case, in which only the excess B (when  $\beta > \alpha$ ) is sequestered, we should have  $\alpha' = \beta' = \alpha$ . The consequence of the factor  $(\sqrt{\pi b}, 2b)$  is to further suppress the current through the coupled transition. As the uncoupled monon transitions have current  $\propto b$ , this results in an overhead of  $\sim b^{-3/2}$  in the best case and  $\sim b^{-2}$  in the worst. This overhead is comparable to the time penalties we found for communication and synchronisation interactions in Part II [7], as for the coupled reaction to proceed both the monon and the complexed klonon must be in receptive states. Drawing further from the analogy, one may wonder whether reactions involving multiple klona are subject to even greater overheads; fortunately, this is not the case: provided only one sequestration klonon participates, by choosing a single representative from each of the reactant klona and product klona to engage in complexation with the sequestration klona, then the results stand.

## 4. Conclusion

In this paper we studied the interaction of mona, such as molecular computers, with klona, shared resources, in a system with a limiting supply of free energy. As in Part II [7], we found this to be very expensive in comparison to independent computation. In general the cost, in terms of transition time, to couple mona to an arbitrary system of klona is of order  $\mathcal{O}(b^{-2})$  (an overhead of  $\mathcal{O}(b^{-1})$ ) but, in contrast to the mona-mona interactions of Part II, lighter overheads of order  $\mathcal{O}(b^{-1/2})$  or better can be achieved under certain conditions. For example, whilst we found that all resource distribution schemes we could devise (some particular examples of which were illustrated in Section 2) did not admit a fractional resource availability above  $\mathcal{O}(b)$  if time overheads were to be avoided, by allowing a time overhead of  $\mathcal{O}(b^{-1/2})$  (thence a time penalty of  $\sim \mathcal{O}(b^{-3/2})$ ) the full resource pool was rendered accessible. We conjectured that these results are in fact general bounds on the capabilities of any such system.

To drive a general out-of-equilibrium system of klona, one needs to supply sufficient free energy to overcome the entropy cost of the reaction. When this cost is known or bounded, such as in biological systems, the definition of ‘sufficient’ can be precomputed and built in to the specifications of the system. In general, however, the cost may be unknown or unbounded, or the favourable direction of the reaction may switch over time. We proposed a scheme in Section 3 to dynamically infer this cost, as well as automatically applying the correct amount of free energy

tokens to pay for it, thereby exposing an equilibrated interface to the disequilibrium reaction. A more detailed specification of the (abstract) molecular realisation of such a scheme is presented in Appendix B. This equilibrated interface can then be readily coupled to our monon transitions, however the cost must still be paid in terms of transition time overhead and indeed it is, as the concentration of the interface is commensurately reduced. Unfortunately our proposed scheme has a minimum overhead of  $\mathcal{O}(b^{-1/2})$  even in the case where the underlying reaction is at or near equilibrium (where it is theoretically possible to achieve zero overhead). It is unknown whether a better scheme exists, whilst still retaining the ability to drive reactions that are arbitrarily far from equilibrium.

As any non-zero overhead to a class of interactions will lead to this class ‘freezing out’ in very large reversible computers, their frequency must be minimised—the proposed scheme notwithstanding. Once again referencing the conclusions of Part II, these interactions thus minimised can make use of a subpopulation bias klona of greater free energy to drive these unfavourable interactions more effectively.

## A. Acknowledgements

The author would like to acknowledge the invaluable help and support of his supervisor, Gos Micklem. This work was supported by the Engineering and Physical Sciences Research Council, project reference 1781682.

## B. Molecular Counter Implementation

Here we present example implementations of the counter klona  $X_k^{(n)}$  used in Section 3. These implementations are abstract, but point towards a viable molecular realisation.

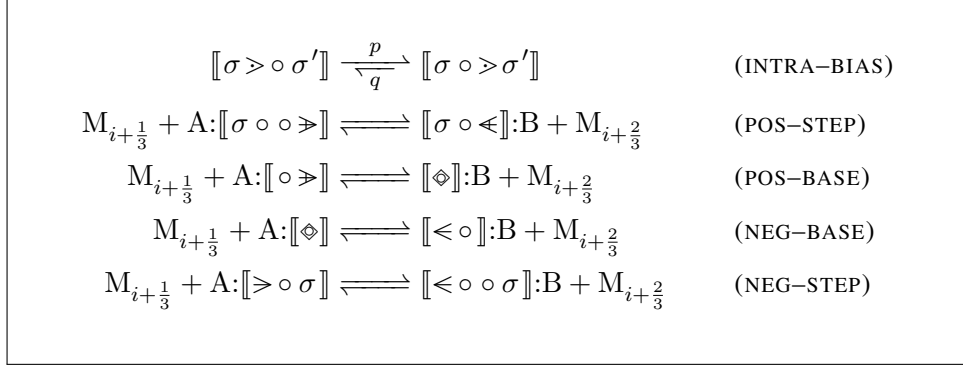
### B.1. Unary Representation

The simplest implementation of counters is given by representing numbers in unary, i.e. a Peano axiomatic approach. For a natural integer  $n \in \mathbb{N}$ , we propose a polymeric representation consisting of  $n$  repeated  $\circ$  monomers and represented thus,

$$\underbrace{[\circ \circ \cdots \circ \circ]}_n.$$

We can then introduce a subcounter  $0 \leq k \leq n$  by incorporating a marker monomer,  $\succ$ , placed between the monomers and which can migrate between the  $n + 1$  possible positions. Moreover, we can handle non-positive values of  $n$  and  $k$  via the alternative marker  $\succ$ . Whilst not strictly necessary, we also differentiate the case  $n = k = 0$  with a third marker monomer,  $\diamond$ . Example





Listing 1: The abstract molecular reaction scheme implementing  $X_k^{(n)} \rightleftharpoons X_{k+1}^{(n)}$  and  $X_n^{(n)} \rightleftharpoons X_0^{(n-1)}$  for the unary representation case, with  $\triangleright$  representing any marker monomer and  $\sigma, \sigma'$  representing arbitrary strings of  $\circ$  monomers.

polymers are given for  $n = -4, 0, 4$ :

$$\begin{array}{lll}
X_0^{(4)} = \llbracket \triangleright \circ \circ \circ \rrbracket, & & X_0^{(-4)} = \llbracket \circ \circ \circ \triangleright \rrbracket, \\
X_1^{(4)} = \llbracket \circ \triangleright \circ \circ \rrbracket, & & X_{-1}^{(-4)} = \llbracket \circ \circ \circ \triangleright \rrbracket, \\
X_2^{(4)} = \llbracket \circ \circ \triangleright \circ \rrbracket, & X_0^{(0)} = \llbracket \diamond \rrbracket, & X_{-2}^{(-4)} = \llbracket \circ \circ \triangleright \circ \rrbracket, \\
X_3^{(4)} = \llbracket \circ \circ \circ \triangleright \rrbracket, & & X_{-3}^{(-4)} = \llbracket \circ \triangleright \circ \circ \rrbracket, \\
X_4^{(4)} = \llbracket \circ \circ \circ \circ \triangleright \rrbracket, & & X_{-4}^{(-4)} = \llbracket \triangleright \circ \circ \circ \rrbracket.
\end{array}$$

Notice the mirrored interpretation of marker positions as values of  $k$  for positive and negative values of  $n$ ; this is intentional and serves to simplify the implementation of the intra- $n$  biasing reactions. It should also be noted that these polymers are achiral and so their mirror images, e.g.  $X_1^{(4)} = \llbracket \circ \circ \circ \triangleleft \rrbracket$ , are equivalent representations of the same polymer. This property greatly simplifies the implementation of the  $X_n^{(n)} \rightleftharpoons X_0^{(n-1)}$  reactions. Altogether, this implementation can be achieved with five reactions as listed in Listing 1. In fact, this scheme can be implemented with just four reactions by replacing  $\diamond$  with  $\triangleright$  and combining reactions (POS-STEP) and (POS-BASE), but we prefer the given implementation for its symmetry. More concretely, these reactions engender the following bi-infinite Markov chain:

$$\begin{array}{c}
\cdots \rightleftharpoons \llbracket \triangleright \circ \circ \rrbracket \rightleftharpoons \llbracket \circ \triangleright \circ \rrbracket \rightleftharpoons \llbracket \circ \circ \triangleright \rrbracket \rightleftharpoons \llbracket \triangleright \circ \rrbracket \rightleftharpoons \llbracket \circ \triangleright \rrbracket \rightleftharpoons \llbracket \diamond \rrbracket \\
\llbracket \diamond \rrbracket \rightleftharpoons \llbracket \triangleleft \circ \rrbracket \rightleftharpoons \llbracket \circ \triangleleft \rrbracket \rightleftharpoons \llbracket \triangleleft \circ \circ \rrbracket \rightleftharpoons \llbracket \circ \triangleleft \circ \rrbracket \rightleftharpoons \llbracket \circ \circ \triangleleft \rrbracket \rightleftharpoons \cdots
\end{array}$$

## B.2. Binary Representation

Whilst particularly simple, the unary representation is very space inefficient; worse, incrementing and decrementing  $n$ , for the purpose of modulating access to a sparse resource/disequilibrium




Figure 3: The state diagram for a Turing Machine [18] implementing incrementation of a natural number in binary positional notation. The alphabet is  $\{\emptyset, 0, 1\}$ , with  $\emptyset$  indicating a blank square; numbers may be provided with any number of leading 0s, and the initial position should be the least significant digit.

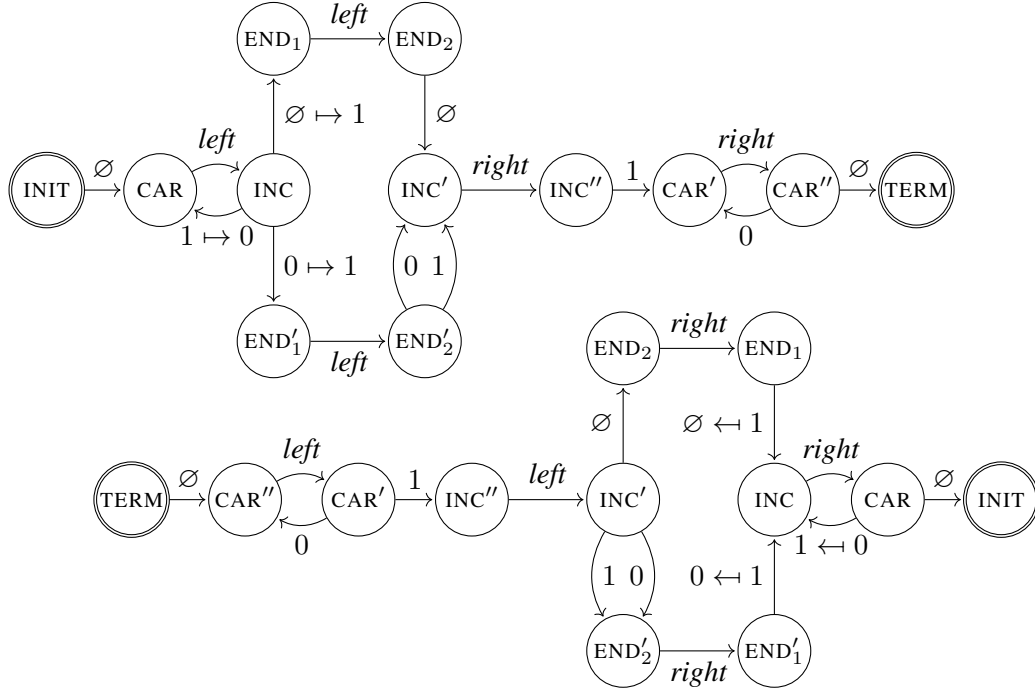


Figure 4: The state diagram for a Reversible Turing Machine [19] (and its reverse) implementing incrementation (resp. decrementation) of a natural number in binary positional representation. The alphabet is  $\{\emptyset, 0, 1\}$ , with  $\emptyset$  indicating a blank square; numbers must be provided in the described prefix-free form, and the initial position should be the (empty) square to the immediate right of the least significant digit. The box-shaped ‘subroutine’ is used to handle the case when a new digit must be prepended; as this is a reversible branch, we must ensure that the converged state  $INC'$  is able to uniquely determine the branch from whence it came.

```

[] \INC\ ['1];
['0 x · b] \INC\ ['1 x · b];
['1 · b] \INC\ ['0 · b']:
b \INC\ b'.

```

Listing 2: A recursive N/alethe [20] program implementing incrementation of a natural number in binary positional representation. Numbers are provided in the described prefix-free form as a list of digits  $\{0, 1\}$  with the least-significant bit first. The second case pattern matches on two bits in order to ensure its output pattern is distinct from that of the first case, and this is where a program will stall should a number be provided which is not in prefix-free form.


reaction, will itself require interaction with a resource pool in order to acquire and release  $\circ$  monomers. Fortunately an exponentially more compact representation exists in the form of positional notations such as binary, with spatial complexity logarithmic in  $|n|$  and hence needing far fewer interactions with its respective monomer pool.

Unlike with the unary case, we cannot exploit the structure of the representation of  $n$  to represent  $k$ . Instead we shall need to store both explicitly. We will also require the ability to recognise the states  $k = 0$  and  $k = n$ , both for determining whether the coupled reaction  $A \rightleftharpoons B$  should proceed and for ensuring the reactions that generate the intra- $n$  distribution do not produce illegal values of  $k$ . These two states correspond respectively to all the digits of  $k$  being 0 or being identical to those of  $n$ . Since  $|k| < |n|$ , its minimal representation will be no larger than that of  $n$  and hence a convenient structure for encoding the pair  $(n, k)$  with the ability to easily check the condition  $n = k$  is given by a ‘zipped’ double-stranded polymer which we denote by  $\llbracket \begin{smallmatrix} n \\ k \end{smallmatrix} \rrbracket$ ; for example, in base 2 the pair  $(11, 6)$  would be given by

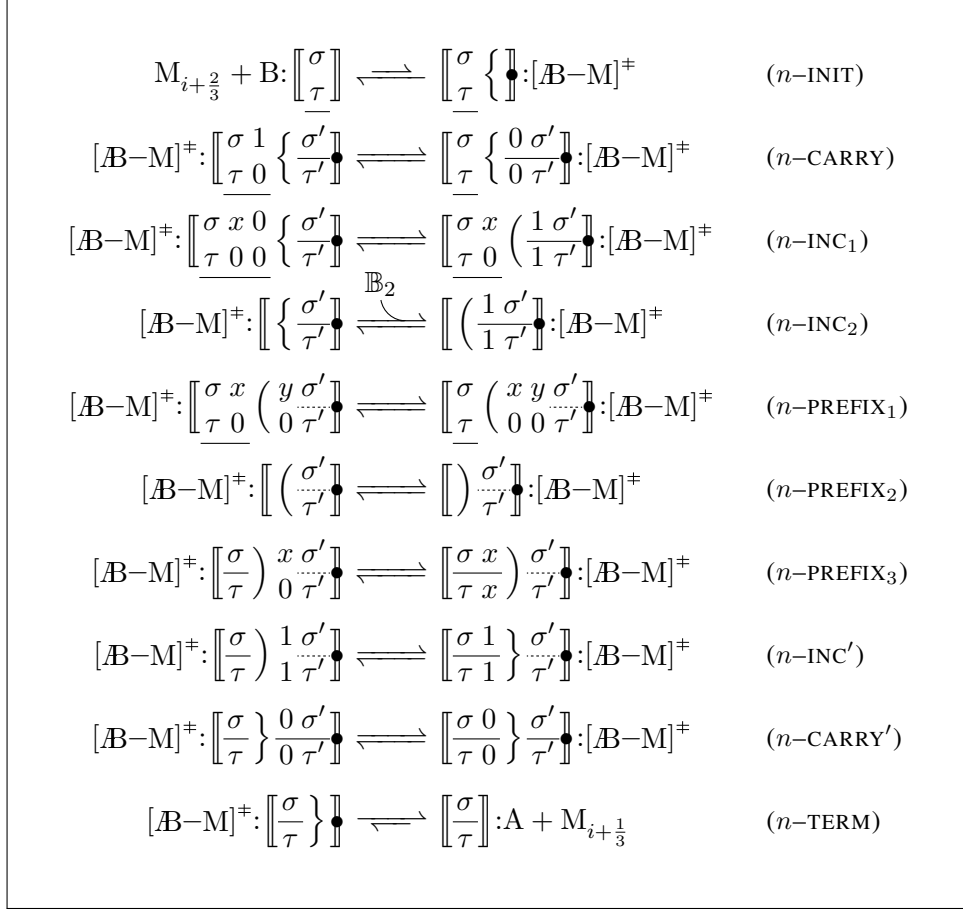
$$\llbracket \begin{smallmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \end{smallmatrix} \rrbracket.$$

In contrast to the unary case, this representation *does* have an intrinsic polarity, with the leftmost digits being the most significant. To ensure logical reversibility of the forthcoming reaction scheme, the polymer representation must be unique for each  $(n, k)$  pair, but positional representations admit a countably infinite equivalence class for each integer (e.g. 3 has base-10 representations  $\{3, 03, 003, \dots\}$ ). We resolve this ambiguity by trimming all leading zeros of  $n$  and trimming  $k$  to the same length. This has the consequence that  $(0, 0)$  has the ‘empty’ representation  $\llbracket \rrbracket$ ; it is certainly possible to introduce a special case for  $(0, 0)$  of  $\llbracket \begin{smallmatrix} 0 \\ 0 \end{smallmatrix} \rrbracket$  for æsthetic reasons, but this would significantly complicate the reaction scheme for no other benefit.

Thus far, we have only considered the representation for non-negative  $k, n \in \mathbb{N}$ ; there are a number of approaches one could take to extend the representation, from simply incorporating a *sign* trit,  $\{-, 0, +\}$ , similar to the approach for the unary case, to a two’s complement approach.



Listing 3: The abstract molecular reaction scheme implementing  $X_k^{(n)} \rightleftharpoons X_{k+1}^{(n)}$  for the binary representation case. To only use a single bias token's worth of free energy, one can make use of a similar scheme to that used in Equation (1), and perhaps making some of the reactions in this scheme unbiased. Alternatively one can accept the use of many tokens, which has the advantage of increasing the rate of the coupled reaction due to the effectively higher bias, but has the disadvantage of slower convergence of the sequestration klona to steady state—potentially interfering with correct operation of the coupled reaction during periods of high flux. In this scheme, the  $\langle$  and  $\rangle$  symbols represent klona that mark the current state and progress of the overall reaction. It is of note that, in the  $(k\text{-CARRY})$  rule, there is no provision for  $Q$  to be set on the bit to the left of  $\frac{x}{1}$ , as such a state would imply we are trying to increment  $k \geq n$  and is thus illegal.



Listing 4: The abstract molecular reaction scheme implementing  $X_0^{(n)} \rightleftharpoons X_{n+1}^{(n+1)}$  for the binary representation case. This overall reaction is unbiased, which raises the concern that the sequestration klona could spend an undesirable amount of time in the intermediate states above. To limit this possibility, the intermediate states should be made more unfavourable and this is achieved by the complementary biases in the rules ( $n$ -INIT) and ( $n$ -TERM). These biases need not be provided by the bias tokens, and in fact that is unideal as the bias is vanishingly small; instead it can be implemented by introducing an enthalpy change, such that there is a chance of thermal activation of the reaction, whilst the klona still prefer to exist in the non-intermediate states. As in the scheme for the intra- $n$  reaction, klona are introduced to mark the state and progress of the overall reaction; in this case,  $\{$ ,  $\}$ ,  $($ , and  $)$ , with the  $\{$  and  $\}$  klona corresponding primarily to incrementation of  $n$  and the  $($  and  $)$  klona to copying of the prefix of  $n$  into  $k$ .

The two’s complement approach is particularly attractive as it would require only minimal adjustments to the non-negative reaction scheme. The two’s complement representation of negative numbers is realised by taking the equivalence class for the integer, e.g.  $\{11, 011, 0011, \dots\}$  for 3, and complementing each bit to obtain, e.g.,  $\{00, 100, 1100, \dots\}$ . That is, each negative integer has an infinite prefix of 1s instead of 0s as for the non-negative case. This approach is used in most conventional CPUs because no special-casing for negative numbers is necessary: adding a negative number  $m$ , with two’s complement  $m'$ , to a positive number  $n$  is equivalent to adding the positive numbers  $m'$  and  $n$ . As a simple justification for this, consider decrementing  $1000 \dots 000$ . We will need to ‘borrow’ a 1 from each digit to the left, eventually obtaining  $111 \dots 111$ . In the limit of infinitely many 0s, this results in infinitely many 1s. For brevity we shall restrict our attention to the non-negative case, but extending to the full domain is a straightforward—if tedious—exercise.

Whilst recognising the cases  $k = 0$  and  $k = n$  is not difficult in our representation, comparing  $\sim \log n$  digits each time we wish to execute the coupled reaction is not ideal. By a slight increase in the complexity of the  $X_k^{(n)} \rightleftharpoons X_{k+1}^{(n)}$  and  $X_n^{(n)} \rightleftharpoons X_0^{(n-1)}$  algorithms, we can render these checks trivial. In particular, we augment the polymer with two bits of state for each monomer pair which we call  $Q$ , for eQuality, and  $Z$ , for Zero. As biochemical precedent, compare with phosphorylation sites. The function of these states is best understood by example; for  $n = 11$ , the  $Q$ - $Z$  states for each of  $k = 0, 1, 8, 11$  are given by

$$\begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix}.$$

Namely, a prefix of zeroes is marked by the  $Z$  state (bottom lines) and a prefix matching  $n$  is marked by the  $Q$  state (middle lines), and these are easily achieved by implementing two local invariants for each digit pair  $\begin{smallmatrix} x \\ y \end{smallmatrix}$ :

- ( $Q$ ) The  $Q$  state is *on* if and only if  $y = x$  and the  $Q$  state of the digit pair to its left (if it exists) is also *on*.
- ( $Z$ ) The  $Z$  state is *on* if and only if  $y = 0$  and the  $Q$  state of the digit pair to its left (if it exists) is also *on*.

It can be seen that these two states are orthogonal as the leading bit of  $n$  must be 1, with the possible exception of  $(0, 0)$  where it makes sense to define  $Q$  and  $Z$  to be both *on* despite the lack of lack of monomers on which to mark said states.

Before implementing a reaction scheme, we first require an algorithmic understanding of reversibly incrementing/decrementing integers in positional notation. As it is possible (and indeed trivial) to do so for integers in unary notation as demonstrated by Listing 1 we should expect it to also be possible for positional notation, and indeed it is. The irreversible algorithm for incrementing a binary number is very simple (Figure 3): if there are any trailing 1 bits then we carry the 1 (flipping the intermediate 1 bits to 0) until we reach the first 0, which we flip to 1. If we reach the end of the number, we prepend a new 1 bit. Making this algorithm reversible is conceptually simple, but expressing it in the form of a Reversible Turing Machine (RTM, as introduced by Bennett [19]) is somewhat involved. An example implementation is provided in

Figure 4, and can be broken down into three stages: (1) carrying; (2) incrementing the least significant 0 bit, or prepending a fresh 1 bit, followed by the logic to reversibly join these two branches; (3) ‘reverse’ carrying, in which we return to the starting position and use the fact that the suffix we are processing takes the form  $100 \cdots 000$  to ensure we can logically reverse this stage. It can be readily checked that this scheme as presented is reversible, and indeed we provide the mechanical reversal of the RTM which may be seen to implement the operation of reversibly decrementing a positive natural number. To demonstrate that the algorithm is in fact conceptually simple, an equivalent recursive program written in a higher level reversible language is shown in Listing 2.

Finally, the reaction schemes for  $X_k^{(n)} \rightleftharpoons X_{k+1}^{(n)}$  and  $X_n^{(n)} \rightleftharpoons X_0^{(n-1)}$  are made manifest in Listings 3 and 4 respectively. In addition to implementing the reversible incrementation procedure in an abstract molecular sense, these schemes must also incorporate the relevant logic for ensuring  $k$  is appropriately bounded and that the  $Q$  and  $Z$  invariants are preserved (although these may be transiently broken within the intermediate states). Furthermore the logic for the second reaction (which we choose to implement in reverse, i.e.  $X_0^{(n)} \rightleftharpoons X_{n+1}^{(n+1)}$ , per the freedom reversibility affords us) must reversibly copy the value of  $n$  into  $k$ , and thus the second scheme must read the entire length of the polymer. As for the case of  $n \equiv 2^p - 1$  in  $n \mapsto n + 1$  (resp.  $n + 1 \mapsto n$ ), the scheme recruits (resp. releases) a pair of binary monomers ( $\mathbb{B}_2$ ) which will be provided by some resource pool mediated by its own sequestration klona. To prevent the molecular algorithms running multiple times at once, the intermediate states are distinguished via an additional state in the form  $\llbracket \cdot \rrbracket$ .

## References

- [1] Rolf Landauer. ‘Irreversibility and heat generation in the computing process’. In: *IBM J. Res. Dev.* 5.3 (1961), pp. 183–191.
- [2] Leo Szilard. ‘Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen’. In: *Zeitschrift für Physik* 53.11-12 (1929), pp. 840–856. English translation: Leo Szilard. ‘On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings’. In: *Behavioral Science* 9.4 (1964), pp. 301–310.
- [3] Edward Fredkin and Tommaso Toffoli. ‘Conservative Logic’. en. In: *Collision-Based Computing*. Ed. by Andrew Adamatzky. Springer London, 1981, pp. 47–81.
- [4] Charles H Bennett. ‘The thermodynamics of computation—a review’. en. In: *Int. J. Theor. Phys.* 21.12 (Dec. 1982), pp. 905–940.
- [5] Michael Patrick Frank. ‘Reversibility for efficient computing’. PhD thesis. Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1999.
- [6] Hannah Earley. ‘Engines of Parsimony: Part I. Limits on Computational Rates in Physical Systems’. In: (2020). arXiv: [2007.03605](https://arxiv.org/abs/2007.03605) [[cond-mat.stat-mech](#)].
- [7] Hannah Earley. ‘Engines of Parsimony: Part II. Performance Trade-offs for Communicating Reversible Computers’. In: (2020). arXiv: [2011.04054](https://arxiv.org/abs/2011.04054) [[cond-mat.stat-mech](#)].
- [8] Paul Wilhelm Karl Rothmund. ‘A DNA and restriction enzyme implementation of Turing Machines’. In: *DNA based computers* 27 (1995), pp. 75–119.
- [9] Erik Winfree. *Simulations of Computing by Self-Assembly*. Tech. rep. California Institute of Technology, Department of Computer Science, 1998.

- [10] Lulu Qian and Erik Winfree. ‘A simple DNA gate motif for synthesizing large-scale circuits’. en. In: *J. R. Soc. Interface* 8.62 (Feb. 2011), pp. 1281–1297.
- [11] Lulu Qian and Erik Winfree. ‘Scaling up digital circuit computation with DNA strand displacement cascades’. en. In: *Science* 332.6034 (June 2011), pp. 1196–1201. ISSN: 0036-8075, 1095-9203.
- [12] David Soloveichik et al. ‘Computation with finite stochastic chemical reaction networks’. en. In: *Nat. Comput.* 7.4 (Dec. 2008), pp. 615–633. ISSN: 1567-7818, 1572-9796.
- [13] Georg Seelig et al. ‘Enzyme-free nucleic acid logic circuits’. en. In: *Science* 314.5805 (Dec. 2006), pp. 1585–1588. ISSN: 0036-8075, 1095-9203.
- [14] Luca Cardelli. ‘Two-Domain DNA Strand Displacement’. In: (June 2010). arXiv: [1006.2993 \[cs.LG\]](#).
- [15] Lulu Qian, Erik Winfree and Jehoshua Bruck. ‘Neural network computation with DNA strand displacement cascades’. en. In: *Nature* 475.7356 (July 2011), pp. 368–372. ISSN: 0028-0836, 1476-4687.
- [16] David Soloveichik, Georg Seelig and Erik Winfree. ‘DNA as a universal substrate for chemical kinetics’. en. In: *Proc. Natl. Acad. Sci. U. S. A.* 107.12 (Mar. 2010), pp. 5393–5398.
- [17] Kevin Montagne et al. ‘Programming an in vitro DNA oscillator using a molecular networking strategy’. In: *Mol. Syst. Biol.* 7.1 (2011), p. 466.
- [18] Alan Mathison Turing. ‘On computable numbers, with an application to the Entscheidungsproblem’. In: *J. of Math* 58.345-363 (1936), p. 5.
- [19] Charles H Bennett. ‘Logical Reversibility of Computation’. In: *IBM J. Res. Dev.* 17.6 (Nov. 1973), pp. 525–532.
- [20] Hannah Earley. ‘The  $\lambda$  Calculus. A declarative model of reversible programming with relevance to Brownian computers’. In: (2020). arXiv: [2011.14989 \[cs.PL\]](#).