



# Kent Academic Repository

**Chu, Dominique (2018) *Performance limits and trade-offs in entropy-driven biochemical computers*. Journal of Theoretical Biology, 443 . pp. 1-9. ISSN 0022-5193.**

## Downloaded from

<https://kar.kent.ac.uk/59768/> The University of Kent's Academic Repository KAR

## The version of record is available from

<https://doi.org/10.1016/j.jtbi.2018.01.022>

## This document version

Author's Accepted Manuscript

## DOI for this version

## Licence for this version

CC BY-NC-ND (Attribution-NonCommercial-NoDerivatives)

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal* , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

## Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

# Performance limits and trade-offs in entropy-driven biochemical computers

Dominique Chu

*School of Computing  
University of Kent  
CT2 7NF, Canterbury  
Tel.: +123-45-678910  
Fax: +123-45-678910*

---

## Abstract

It is now widely accepted that biochemical reaction networks can perform computations. Examples are kinetic proof reading, gene regulation, or signalling networks. For many of these systems it was found that their computational performance is limited by a trade-off between the metabolic cost, the speed and the accuracy of the computation. In order to gain insight into the origins of these trade-offs, we consider entropy-driven computers as a model of biochemical computation. Using tools from stochastic thermodynamics, we show that entropy-driven computation is subject to a trade-off between accuracy and metabolic cost, but does not involve time-trade-offs. Time trade-offs appear when it is taken into account that the result of the computation needs to be measured in order to be known. We argue that this measurement process, although usually ignored, is a major contributor to the cost of biochemical computation.

*Keywords:* biological computing, information thermodynamics, cost of computation, linear noise approximation

---

## 1. Introduction

Computing architectures based on biochemistry, rather than semi-conductor technologies, are attracting increasing interest as alternative models of computation [1]. Biochemistry can be used to engineer novel types of computers based on biological components. Examples include, DNA based computers [2, 3], robots controlled by slimemolds [4], or logic gates implemented in living cells [5, 6, 7, 8]. Beside this technological importance of biochemical computers, there is now also an increasing appreciation that information processing may be an important fitness contributing function for natural organisms [9, 10]. There are a number of biosystems that have been studied as *in vivo* special purpose computations. For example, kinetic proofreading [11, 12] greatly enhances the copying fidelity during translation and is often interpreted as an *in vivo* computation. A classical example of biochemical computation is bacterial sensing [13, 14, 15, 16], whereby cells measure molecular concentrations in their environment and modify internal pathways and gene expression levels in response. Chemotaxis [17], for instance, depends on organisms sensing a molecular concentration gradient by computing the difference between several measurements, either in time or across the cell volume. Most recently even bacterial growth dynamics has been interpreted as a computational process [18, 19].

Detailed case studies of biological computers often find performance limits to biochemical computations. For a simple gene-switch, Zabet and coworker found a trade-off between the cost, the accuracy and the speed of the computation [20, 21]. Similar trade-offs were established for other biological systems, including chemotaxis [22], regulation of nutrient uptake [18], and translation [23]; lower limits on the cost of sensing (not involving trade-offs) have also been found recently [24, 14].

Intuitively such trade-offs are to be expected. Bio-chemical networks are stochastic systems and as such subject to noise. Overcoming this noise requires energy input and time. Energy-time-accuracy trade-offs are also implied by the classical results on the physics of computation [25, 26, 27]. While there does not seem to

---

*Email address:* d.f.chu@kent.ac.uk (Dominique Chu)

be a lower limit for the energy used during a computation, Bennett [26] pointed out that in the zero energy limit the speed of computation goes to zero. Computations that complete within a finite time, therefore require finite energy resources.

The question is now whether one can go beyond both the individual case studies of biochemical computers and the intuitive arguments and establish a model which provides insights into the origins of the performance limits to biochemical computations. The task is a difficult one. For one, there is a wide variety of approaches to biochemical computation (only some of which are mentioned above). At the same time, there is no general definition of biological computation, i.e. it is not clear how to distinguish a reaction network that computes from one that does not.

For the purpose of this article, we will take a pragmatic approach with respect to the latter question and simply identify (in section 2.1) computation with out-of-equilibrium biochemical processes. According to this, every biochemical process that is not in equilibrium performs a computation. As far as the wide variety of biochemical computations are concerned, we will abstract away from specific models and define the concept of *entropy driven computers* (EDC) in section 2.2. This will capture many properties of *in vivo* computers as they appear in biological systems. EDCs are in many aspects different from real biological networks, but we will argue that they share important characteristics with a wide range of biochemical computers. It is perhaps best to think of EDCs as test-tube biochemistry, as it is frequently used in biological research to study reaction networks *in vitro*.

We model EDCs as continuous time Markov chain models of biochemical systems. We assume that each model is initialised in some state and then left to relax to equilibrium. We will then interpret this relaxation process as a computation. Throughout this article we will assume that the EDC is of mesoscopic scale. By this we mean that it is still affected by stochastic fluctuations, but that it is also within the range of validity of the linear noise approximation [28]. Simply put, this assumption states that the stochastic system behaves like the deterministic equivalent, plus some noise. The linear noise approximation is a very good approximation for mesoscopic systems and holds true for a wide range of biochemical systems, and hence for a wide range of biological “computers” such as gene regulatory networks, protein-protein interactions or intra and inter-cellular signalling systems, although clearly there are systems that will not be captured by this assumption.

For the purpose of the present contribution, we will identify the cost in energy of a computation with the entropy produced during the computation. While this does not quantify the actual metabolic cost of this computation, it is directly related to it. We will first show that the linear noise approximation implies that the entropy production scales linearly with the system size, while the time-scale to approach equilibrium (which we interpret as the computing time) remains invariant. This means that there is a trade-off between the cost of the computation and its accuracy, but there is no trade-off involving the speed of the computation. Contrary to previous work (or apparently so), this suggests that speed-energy trade-offs are not a fundamental property of biochemical computation *per se*

A trade-off involving time emerges only when it is taken into account that the result of the computation must be measured in order for the computation to have any impact in the world. Any measurement of the outcome of the computation in turn requires a measurement device. This device needs to be brought into contact with the computer to determine its state. Device and computer then form a joint system, which initially will be out of equilibrium but relaxes to an equilibrium. This relaxation constitutes the measurement process. Formally a measurement is thus also an entropy driven computation. As we will show below, restoring the computer to its original state, while at the same time leaving the measurement device in a state that indicates the result of the computation, requires both energy input and time. It also leads to a trade-off between the energy used and the speed with which the restoration can be completed with a given confidence. A second trade-off involving time arises from the stochastic nature of the computer. A single measurement only indicates the correct result with a certain probability. Repeated measurements are necessary in order to sample the state of the computer reliably, thus leading to a trade-off between accuracy and time.

## 2. Results

### 2.1. Computation by biochemical systems

The current modus operandi in the field of biochemical computing is to identify a biological system (such as sensing or proof-reading) as a computation when it implements a function that is naturally interpreted as a computation. This approach enables deep insights into specific examples, but is likely to miss most instantiations of biochemical computation. It would be much more useful to have a concept of biochemical computation that is independent of its function, just as in computer science computation is defined with respect to a number of specific mathematical models, not by reference to what is computed.

The best known model of computation is the *Turing machine*. This is a mathematical construct consisting of a “reading head” that is reading and writing a tape, while changing its internal states in the process, until it reaches a “halting state,” at which point the computation stops. It is believed that for every computable function there is a corresponding Turing machine that computes it. Based on this, one could be tempted to define a biochemical process as a computation if there is a Turing machine that simulates this process. This does not work however: The natural equivalent of a halting state in biochemical systems is the equilibrium state, i.e. the state of the biochemical system where reactions are in detailed balance. Unlike the halting state of a Turing machine, the equilibrium state is of a statistical nature. This means that on average there are no net-fluxes across the network of reactions [29, 30]. This does not mean, however, that reactions stop. Even in equilibrium there is an ongoing chemical activity. Crucially though, the sequence of reaction events is symmetric in time [28], such that an observer would not be able to tell apart an actual sequence of reactions from a (hypothetical) reversed sequence. Equilibrium is not time-directed. Computation, on the other hand, is necessarily time directed, mapping a particular input to a particular output. Equilibrium systems are therefore not able to compute. Sample paths of equilibrium biochemical systems can still be simulated and are thus computable by Turing machines, whether or not the system is in equilibrium. This demonstrates that not all processes that can be simulated by Turing machines are also themselves processing information.

For the purpose of this paper, we will adopt the simplest working hypothesis and postulate that the equilibrium state is the only halting state of biochemical computers. This implies that every biochemical system that is not in equilibrium is in the process of computing. By adopting this definition, we also accept that most biochemical computers will not do any useful calculations, just as almost all Turing machines do not compute anything of interest.

### 2.2. Entropy driven computation

In this section we define an EDC as a closed, stochastic, biochemical system, denoted by a fraktur  $S$ ,  $\mathfrak{S}$ . The system does not exchange particles with the environment. We conceptualise  $\mathfrak{S}$  as consisting of a (typically very large) number of discrete microstates  $\mathfrak{s}_0, \mathfrak{s}_1, \dots, \mathfrak{s}_m$  (see SI section 1 for more details). An EDC is initialised in a macrostate  $M_0^{\mathfrak{S}}$  characterised by a specified abundance for each of its constituent biochemical species at time  $t = 0$ ; see SI section 1 for a detailed explanation of what we mean by “macrostate.” After a transient period, the biochemical system approaches an equilibrium state  $M_{\infty}^{\mathfrak{S}}$  characterised by detailed balance. The approach to equilibrium is the computation. Strictly speaking, the transition to equilibrium takes an infinite amount of time. In practice, EDCs will be very close to equilibrium after a finite, possibly very short, time. We will model EDCs here as continuous time Markov chains. Then the time scale for the approach to equilibrium depends on the kinetic parameters appearing in the master-equation that defines  $\mathfrak{S}$ . We will take this time scale as the *speed* of the computation.

In contrast to EDCs, *in vivo* computations, i.e. cells, never approach equilibrium, but rather operate around non-equilibrium steady states. We acknowledge this, but still argue that for the present purpose equilibrium models are more convenient. They are also more revealing of underlying principles, for the following reasons: Firstly, steady state processes have an additional “maintenance” entropy production. This may be substantial, but is not related to the cost of the computation *per se*. Mathematically, it is possible to separate the contribution from the relaxation from the contribution to the entropy production that arises from the maintenance of the non-equilibrium state. By focussing on the simpler equilibrium case, we circumvent this mathematical complication that does not add anything to the present purpose. Secondly, as will become clear below, considering equilibrium processes naturally forces a conceptual separation between the computation *per se* and the reading of the result. Both are separate processes that limit computational

performance in qualitatively different ways. Finally, the equilibrium state is conceptually reminiscent of the halting state in Turing machines, which in turn is of fundamental importance for the theory of computation. Indeed, there is an interesting analogy in the relationship between on the one hand a non-equilibrium steady state, sustained by an organism and test-tube biochemical systems approaching equilibrium and on the other hand an operating system — which is not supposed to halt — and an algorithm — which must halt. Models of entropy driven computation therefore naturally link biology, physics and theoretical computer science. With this being said, the basic conclusions that we will reach depend primarily on the linear noise approximation, which in turn depends on system size, not on the distance from equilibrium.

### 2.3. Cost of the EDC proper

#### 2.3.1. The model

We assume that the stochastic system  $\mathfrak{S}$  is defined by the master equation, which is a differential equation for the probability to observe a particular combination of molecular abundances  $\mathbf{n}$  at time  $t$  [31].

$$\dot{P}(\mathbf{n}, t) = \sum_i (w_i(\mathbf{n} - \boldsymbol{\sigma}_i)P(\mathbf{n} - \boldsymbol{\sigma}_i, t) - w_i(\mathbf{n})P(\mathbf{n}, t)) \quad (1)$$

where  $w_i(\mathbf{n}) := k_i h_i(\mathbf{n})$  is the total rate of reaction  $i$  and  $h$  is the multiplier indicating how many combinations of molecules can realise this reaction;  $\mathbf{n}$  and  $\boldsymbol{\sigma}_i$  are the particle vector and the stoichiometric vector of reaction  $i$ . Altogether, the right hand side of the equation contains two terms. The first one expresses the possibility that the current state characterised by a particle composition  $\mathbf{n}$  was obtained from reaction  $i$  and the state before the reaction happened was  $\mathbf{n} - \boldsymbol{\sigma}_i$ . The second term formulates the possibility that before the last reaction the state of the system was  $\mathbf{n}$ , but reaction  $i$  took it away from this state (to a new state  $\mathbf{n} + \boldsymbol{\sigma}_i$ ).

Given such a stochastic model, we can scale the total number of particles in the initial conditions by a factor  $c$  and the reaction rate constants of bimolecular reactions by  $1/c$ . This amounts to scaling the volume  $V$  of the system while keeping the concentration fixed. In this way we construct an equivalence class  $\mathcal{S}$  of systems  $\mathfrak{S}$ . In general, the members in this class will behave differently. Most importantly, they show different amounts of fluctuations and approach equilibrium at different speeds. For mesoscopic volumes, however, the behaviour of the master equation is increasingly well approximated by the first order linear noise approximation [28] whereby the mean behaviour of the system is invariant to scaling. Scaling the system size in the regime of the linear noise approximation only affects the noise around the mean behaviour, which scales with  $V^{-1/2}$ . Importantly, scaling the volume does not affect the time-scale to reach equilibrium, which is determined by the mean behaviour. In the limiting case of an infinite volume the noise goes to zero and the system is described by a differential equation whose trajectory corresponds (for monostable systems) to the mean of the linear noise approximation.

#### 2.3.2. Entropy production

In this sub-section, we will show that the operation of an EDC is subject to a trade-off between the noise and the cost. An EDC starts with the stochastic system  $\mathfrak{S}$  initialised in some macrostate  $M_0$ . The system then relaxes into an equilibrium state  $M_\infty$ , performing a computation in the process. We will identify the cost of the computation with the amount of entropy generated during the relaxation. There are two components to the entropy. Firstly, the exported entropy or heat dissipated to the environment which accrues whenever the system makes the transition from state  $\mathfrak{s}_i$  to state  $\mathfrak{s}_j$ . This is given by the ratio of the forward rate of the transition  $w_{ji}$  and the corresponding backwards rate  $w_{ij}$  [32].

$$\Delta S_Q = k_B \ln \frac{w_{ji}}{w_{ij}} \quad (2)$$

Using this definition the heat dissipated to the environment is positive, whereas heat extracted from the reservoir is negative. No heat is generated by a system that transitions from a state  $\mathfrak{s}_i$  to a state  $\mathfrak{s}_j$  and then back again. More generally the heat generated by a system that transitions from an initial state  $\mathfrak{s}_0$  to some final state  $\mathfrak{s}_j$  via a number of intermediate states is independent of the particular path, and depends only on the initial and final state. Taking into account that the heat generated in each transition is additive, and applying the detailed balance condition that relates the steady state probabilities  $\pi_i, \pi_j$  of states  $\mathfrak{s}_i$  and  $\mathfrak{s}_j$  of a stochastic system with the transition probabilities, i.e.

$$\pi_i w_{ij} = \pi_j w_{ji} \quad (\text{detailed balance condition})$$




Figure 1: Numerical calculations of the exported entropy (**left**) and the Shannon entropy (**right**) for the system  $A + B \rightleftharpoons C + D \rightleftharpoons E$ . The system size is the number of  $A$  and  $B$  at time  $t = 0$ . The concentration was kept constant as the particle size was increased. The entropy is given in units where  $k_B = 1$ . The forward rate constant was set to 0.3 and the backwards rate constants from  $E$  and  $C + D$  were 0.05 and 0.1 respectively. Prism [33] was used to solve the numerics; for the model file see supplementary information.

we obtain:

$$\Delta S_Q(\mathfrak{s}_0 \rightarrow \mathfrak{s}_j) = k_B \ln \frac{\pi_j}{\pi_0} \quad (3)$$

In addition to the exported entropy, the computation also produces Shannon entropy. This Shannon entropy of a microstate  $\mathfrak{s}_i$  can be written as  $\Delta S_{\text{int}} = -k_B \ln p_i$ , where  $p_i$  is the probability to observe the system  $\mathfrak{S}$  in microstate  $\mathfrak{s}_i$ . On average this gives a contribution  $\langle S_{\text{int}} \rangle = -k_B \sum_i p_i \ln(p_i)$ . The Shannon entropy produced during the computation is then given by  $\langle S_{\text{int}}(\infty) \rangle - \langle S_{\text{int}}(0) \rangle$ .

For realistically sized systems it is usually not possible to solve the master eq. 1 exactly, even in steady state. Analytic expressions for the entropy production are therefore difficult to obtain. However, general scaling arguments can be given in the limit of the linear noise approximation. In this regime, one expects the total entropy during the approach to equilibrium to scale linearly with the system size. Consider, for example, a biochemical system with a mono-modal steady state distribution and discrete, steady state probabilities  $\{\pi_0, \pi_1, \dots, \pi_k, \dots, \pi_m\}$ , where  $m + 1 \gg 1$  is the total number of states of the system. Assume now that the initial state is  $\mathfrak{s}_0$ , then the heat dissipation is given by averaging eq. 3 over all possible end states:

$$\langle S_Q \rangle = k_B \sum_{i=0}^m \pi_i \ln \left( \frac{\pi_i}{\pi_0} \right) = k_B \sum_{i=0}^m \pi_i \ln(\pi_i) - k_B \ln \pi_0. \quad (4)$$

The first term on the right hand side equals the Shannon entropy (up to the sign). The total entropy produced  $\langle S \rangle = \langle S_Q \rangle + \langle S_{\text{int}} \rangle$  is therefore given by

$$\langle S \rangle = -k_B \ln \pi_0. \quad (5)$$

In order to understand the scaling of the entropy production, we approximate the discrete steady state distribution by a continuous Gaussian distribution with mean  $\mu$  and variance  $\sigma^2$ . According to the linear noise approximation, the variance and the mean scale linearly with  $V$ , such that we obtain:

$$\langle S \rangle = -k_B \ln \pi_0 \sim \frac{x_0^2 - 2x_0\mu + \mu^2}{2\sigma^2} \sim V. \quad (6)$$

We stress that this scaling argument is only valid for mesoscopic systems when the linear noise approximation is good. In particular, for very small systems there can be a non-linear relationship between the system size and the heat dissipation. In those cases the master equation needs to be solved to determine the entropy production. This is usually problematic. See SI section 5 for an explicit calculation of the heat dissipation for a minimal example and figure 1 for a graphical representation.

In summary: The ability to determine the result of the EDC is limited by the noise of the system in equilibrium, which scales like  $V^{-\frac{1}{2}}$  in the linear noise approximation. The cost of the computation scales

linearly with the size of the system. Hence there is a trade-off between entropy production (and thus energy cost) and the accuracy of the EDC. No time-trade-offs arise here. In the linear noise approximation, the time evolution of the mean does not depend on the system size.

#### 2.4. Trade-offs arising from the measurement

The equilibration of the EDC is only one part of the computation. In order, for the outcome of the computation to be known, a measurement must be performed to determine the macrostate  $M_\infty^{\mathfrak{S}}$  and the result needs to be recorded. This comes at a cost [34, 35]. We assume that the measurement is done by means of the measurement device  $\tilde{\mathfrak{S}}$ , which is itself an EDC; see SI section 1 for details on the concept of measurement.

We will first discuss how  $\tilde{\mathfrak{S}}$  can be used for binary measurements on observables of  $\mathfrak{S}$ , i.e. in order to determine whether a particular biochemical species  $L$  of  $\mathfrak{S}$  is above or below a threshold abundance. We will assume that  $\tilde{\mathfrak{S}}$  is a bistable system, i.e. in equilibrium, it can be in one of two (transient) macrostates, indicating the result of the computation.

Measurement is only possible if  $\mathfrak{S}$  and  $\tilde{\mathfrak{S}}$  are temporarily brought into contact. We will assume that the contact remains limited to specific interfaces. The biochemical systems  $\mathfrak{S}$  and  $\tilde{\mathfrak{S}}$  should not be allowed to mix, as it would be difficult to separate the systems again after the measurement. Instead, we consider here a protocol whereby  $\mathfrak{S}$  and  $\tilde{\mathfrak{S}}$  remain separated by a wall. We model the interface between system and measuring device as a single “trans-membrane” receptor placed in the surface of  $\tilde{\mathfrak{S}}$ . The external part of the sensor contains a number of binding sites for molecules of type  $L$  of  $\mathfrak{S}$ . The inside of the receptor can interact with  $\tilde{\mathfrak{S}}$ , as will be described in the next section. We will assume a Monod-Wyman-Changeux (MWC) [36] receptor (but other models are possible too). When its external sites are bound then the receptor is heavily biased towards the “active” state. Otherwise it is heavily biased towards being deactivated. When there are several binding sites, then the receptor can support ultra-sensitivity. This means that for low abundances of  $L$  the binding sites are almost never bound; above a threshold abundance the sites are almost always bound. In this way, MWC receptors can be used for binary measurements of external concentrations.

Upon contact, the system and the measurement device form a new system  $\mathfrak{S}\tilde{\mathfrak{S}}$ . The joint system will therefore evolve from an initial macrostate  $M_t^{\mathfrak{S}\tilde{\mathfrak{S}}}$  at time  $t$  of the measurement, to a macrostate  $M_{t+T_{\text{meas}}}^{\mathfrak{S}\tilde{\mathfrak{S}}}$ , where  $T_{\text{meas}}$  is the time required for the measurement. Upon separation of  $\mathfrak{S}$  and  $\tilde{\mathfrak{S}}$  the new macrostate of the measurement device should be a (transient) record of the state of  $\mathfrak{S}$ . Once the measurement is completed it is necessary to restore  $\tilde{\mathfrak{S}}$  as an independent system, i.e. to separate  $\mathfrak{S}$  and  $\tilde{\mathfrak{S}}$ .

Conceptually, there are thus two aspects to the measurement process that are relevant for the question we consider here. Firstly, the process of measurement itself and secondly, the separation of the measurement device from the system  $\mathfrak{S}$ . We find that the binary measurement — while it cannot be done for free — does not imply any trade-offs. In contrast, the separation step leads to a trade-off between energy, cost and time.

##### 2.4.1. The measurement proper

We choose  $\tilde{\mathfrak{S}}$  to be bistable, i.e. there are two macrostates  $M^m$  and  $M^g$  characterised by a high amount of  $m$  or  $g$  respectively. Bistability can occur in equilibrium biochemical systems. Stochastic bistable systems will switch spontaneously, but possibly rarely, between the two macrostates. The expected time between switching events will depend on the system size and the “well depth.” The latter is essentially determined by the kinetic parameters of the system and indicates the difficulty of escaping a metastable state. Larger systems are more stable but even relatively small systems can be sufficient to virtually guarantee stability over any practically relevant time-scale.

In order to switch a bistable system at a *particular time*, it is necessary to take it out of equilibrium by introducing a net probability flux to the desired state of the system. This implies breaking detailed balance and hence dissipates heat. A controlled switch of the bistable system can therefore only be achieved at a certain expense of work. Formally, it would be sufficient to connect a source of free energy (a “chemical battery”) to  $\tilde{\mathfrak{S}}$  during the switch only and disconnect it afterwards. In practice, due to the nature of the biochemical systems, it is difficult to remove such a battery without extra work input. It is therefore much better to operate the device permanently from a battery away from equilibrium.






Figure 2: The trade-off between dissipation and stability. For networks of different size  $N$ , defined as the total number of particles in the system, the energy dissipation after 600 time units was recorded against the stability of the memory. Small systems are performing better than large systems. Stability is defined as the number of  $m$  molecules divided by the number of  $m$  and  $g$  molecules combined after 600 time units. Each point in the graph represents the average over 5000 individual simulations. For each value of  $N$  we sampled initial amounts of ATP from 5 to 300 in steps of 5. The dissipation is the difference between the number of ATP at time 0 and at time 600 averaged over 5000 simulations.

#### 2.4.2. A possible design of a measurement device


We will now construct a measurement device  $\bar{\mathfrak{S}}$  consisting of the molecular species  $\{c, d, e, g, m\}$ . It is bistable in  $m$  and  $g$ , i.e. at any one time only one of those two species is present in a high concentration. We will identify the two metastable states by the macrostate  $M^g$  and  $M^m$  respectively. These states will indicate whether the concentration of species  $L$  of  $\bar{\mathfrak{S}}$  is above or below some threshold. There are many ways to achieve bi-stability in biochemical systems. We choose a mechanism based on two competing autocatalytic reactions: Molecules of type  $c$  are converted into  $d$ . This conversion is catalysed by  $g$ . Molecules of type  $d$  are reversibly converted into  $g$ . Altogether,  $g$  is thus autocatalytic. To achieve bi-stability, we add symmetrically a second autocatalytic circuit consisting of  $e$  which is produced from  $c$  catalysed by  $m$ . Molecules of  $e$  are also reversibly converted into  $m$ . The reactions can be summarised as follows:



To create bi-stability, this system needs to be extended by an antagonistic interaction between  $g$  and  $m$ .



Finally, we also posit that  $g$  and  $m$  inter-convert reversibly, albeit at a very low rate. Forward and backwards rates are equal.




If initialised with a sufficient number of  $c$  and one  $g$  and  $m$  each, then the system will, after a transient period, evolve to a macrostate  $M_\infty^g$  where  $g > m$  or  $M_\infty^m$  characterised by  $m > g$ . The probabilities of either of these states will be equal if the parameters are symmetric between the autocatalytic pathways of  $m$  and  $g$ .

#### 2.4.3. Switching

In order to switch the state of  $\bar{\mathfrak{S}}$ , we couple the chemistry of the device to the internal part of the MWC receptor, which can be in an active state  $r^*$  or an inactive state  $r$ . Here we assume that the active receptor




catalyses the inter-conversion of  $m$  and  $g$ , adding the following reactions to the measurement device:



In the absence of a driving force, the active receptor accelerates the approach to  $m = g$ . Given the right range of parameters it will thus temporarily disrupt bi-stability and bring  $\bar{\mathfrak{S}}$  into a mono-stable regime. Upon removal of the catalysing reaction, i.e. when the internal receptor reverts to the inactive state,  $\bar{\mathfrak{S}}$  will relax quickly to  $M^g$  or  $M^m$ , both with equal probability. No directed switch is possible.

A directed switch is only possible if the catalytic reaction is driven preferentially into one direction. One possibility is to couple it to a biochemical battery, i.e. a reaction that is not in equilibrium. A modified reaction scheme could then be the following:



If there is a large excess of ATP over ADP then this would result in a net drive towards  $g$  whenever the receptor is activated. Assuming the excess of  $g$  is sufficient, the activation of the receptor leads to a switch of the macrostate to  $M^g$ .

#### 2.4.4. Measurement and recording

The measurement procedure is as follows: (i) Reset  $\bar{\mathfrak{S}}$ . (ii) Initiate the measurement proper on  $\mathfrak{S}$  by bringing  $\mathfrak{S}$  and  $\bar{\mathfrak{S}}$  into contact mediated by the MWC receptor. (iii) Wait for a fixed amount of time  $T_{\text{meas}}$ . (iv) Separate  $\mathfrak{S}$  and  $\bar{\mathfrak{S}}$ .

The reset step is necessary so as to ensure that the measurement device is in a known state prior to starting the measurement. Without this step,  $\bar{\mathfrak{S}}$  is in state  $\bar{M}^m$  with probability 1/2. This is problematic, because the measurement of  $\mathfrak{S}$  will only effect a switch in  $\bar{\mathfrak{S}}$  if the concentration of  $L$  is indeed above the threshold. The reset can be implemented as a measurement of some species  $I$  of a reference volume that is in a known state. The measurement should be mediated by an auxiliary receptor different from the one that measures  $\mathfrak{S}$ . The active state of this auxiliary receptor should effect a switch to  $M^g$  when the concentration of  $I$  in the reference system is high (which should always be the case).

Once the reset is completed, the measurement of  $\mathfrak{S}$  can be performed. Following this step the macrostate of  $\bar{\mathfrak{S}}$  is a record of the state of  $\mathfrak{S}$ . The metastable states of  $\bar{\mathfrak{S}}$  are only stable over a finite (although possibly very long) time, because the system may transition spontaneously to a different state.

In principle it is possible to remove the battery from the bistable system post-measurement. In practice this will be difficult because it entails extracting the ATP and ADP molecules from  $\bar{\mathfrak{S}}$  and requires biochemical work. Therefore, it is better not to remove the battery. In this case, however, the macrostate of  $\bar{\mathfrak{S}}$  remains only stable for as long as the battery is sufficiently charged, i.e. as long as there is an excess of ATP over ADP. Once the battery has run out  $\bar{\mathfrak{S}}$  will no longer be bistable (see SI fig. 1).

The biochemical battery is discharged by two processes. Firstly, during switching the receptor is active with high probability and the battery will drive the conversion from  $g$  to  $m$  while using ATP. When the external binding sites of the MWC receptor are occupied the receptor will be active most of the time, resulting in a high rate of discharge. Secondly, there is spontaneous activation of the receptor even if no ligands  $L$  are binding to the outside receptor. The rate of spontaneous activation of the receptor determines the base-rate of ATP usage/battery discharge.

Note that the ability of  $\bar{\mathfrak{S}}$  to measure and record the state of  $\mathfrak{S}$  is not subject to a trade-off between system size and accuracy. In the particular model we present here, small measurement devices  $\bar{\mathfrak{S}}$  perform better than those with a large number of molecules, while also dissipating substantially less energy (see SI fig. 1). Similarly, there are no trade-offs between the speed of the measurement and its accuracy. The speed is related to the scale of the reaction rates inside of  $\bar{\mathfrak{S}}$  which does not affect the heat dissipation of the system.

There are, however, other sources of trade-offs that are closely connected to the measurement process. The accuracy of the binary measurement corresponds to the ability of  $\bar{\mathfrak{S}}$  to indicate whether or not the concentration of the system is below or above a threshold. The parameter that determines this accuracy is the number of binding sites for the ligand  $L$ . The receptor number does not influence the cost of switching, but it does increase the cost and time of separating  $\mathfrak{S}$  and  $\bar{\mathfrak{S}}$ , as will be discussed in the next section.

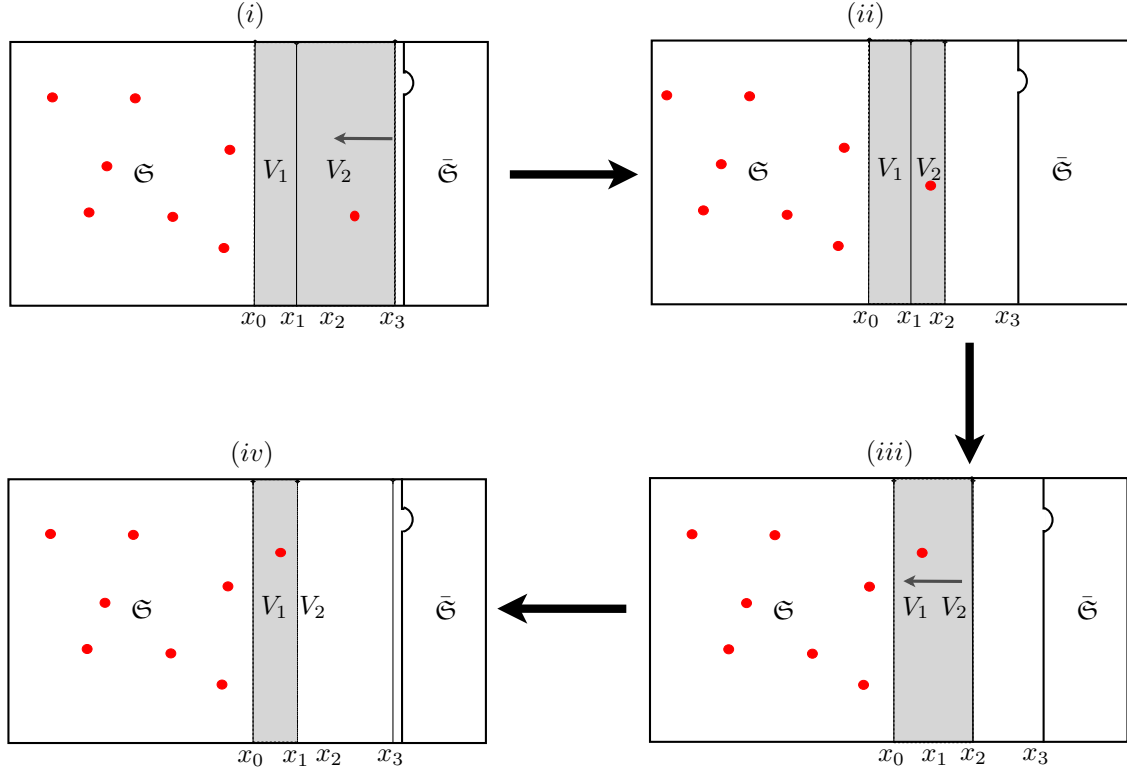


Figure 3: A schematic representation of the resetting procedure. The steps (i) to (iv) are illustrated from top left to bottom right. In this particular case, the particle was successfully captured in the first step and is thus successively transported into the volume  $V_1$ .

#### 2.4.5. The separation step

Once the measurement of the system has been completed the coupled system  $\mathfrak{S}\bar{\mathfrak{S}}$  needs to be separated, while restoring  $\mathfrak{S}$  to a state that is statistically equivalent to the state before the measurement. This simply reflects the condition that the measurement should not alter the system. In our specific case the restoration of the system requires the separation of  $\mathfrak{S}$  and  $\bar{\mathfrak{S}}$  and that the ligand  $L$  is returned to the system. We will find that this separation step gives rise to a trade-off involving time. To do this, we use an abstract model of separation which, using ideas from thermodynamics, estimates minimal work requirements. While the model is not biological realistic, we will nonetheless find that it leads to biologically relevant trade-offs in terms of binding rates that also relate to known limitations of biochemical sensors.

We stipulate that contact between  $\mathfrak{S}$  and  $\bar{\mathfrak{S}}$  is mediated by a straight wall. The system  $\mathfrak{S}$  is fixed in space, but  $\bar{\mathfrak{S}}$  can be moved to the right, which opens up a volume  $V_0$  between the two systems (see fig. 3). We also assume a number of removable walls that can be inserted at arbitrary (but fixed) points to further sub-partition  $V_0$ . We now describe a protocol to restore  $\mathfrak{S}$ . This model involves 3 walls. A simpler — but energetically less favourable — model is possible and described in the SI section 4.

The idea is to insert a wall (Wall 1) close to the membrane of  $\bar{\mathfrak{S}}$  at  $x_0$ ; see fig. 3. It separates  $\mathfrak{S}$  from  $\bar{\mathfrak{S}}$ . The measurement device is then moved to the right, opening up the volume  $V_0$ . Wall 2 is then inserted at some distance to the right of wall 1 at position  $x_1$  so as to form the volume  $V_1$  between it and wall 1. The separation procedure is as follows. (i) After waiting  $T$  time units insert wall 3 at  $x_3$ , which is immediately to the left of the membrane of  $\bar{\mathfrak{S}}$ . Wall 3 separates the receptor (and possibly ligands bound to it) from the volume  $V_0$ . (ii) Slide wall 3 to the point  $x_2$ . Wall 3 now forms a volume  $V_2$  between itself and wall 2 and a volume  $V_0 - V_1 - V_2$  between itself and the membrane of  $\bar{\mathfrak{S}}$ . (iii) Remove wall 2; this opens the volume  $V_1 + V_2$  between  $\mathfrak{S}$  and wall 3. (iv) Slide wall 3 from position  $x_2$  to position  $x_1$ . At this point the cycle is started again at (i) by inserting wall 2 at  $x_3$ .

The minimal work required to separate the system is given by (see SI section 3 for the details of the

calculation):

$$\begin{aligned}
W_{\text{tot}} &= \ln\left(\frac{x^{-1}}{\kappa + x^{-1}}\right) + \\
&+ \ln\left(\frac{V_2(x^{-1} + \kappa)}{(V_0 - x^{-1} - \kappa)(x^{-1} + \kappa + V_2)}\right) - 1
\end{aligned} \tag{7}$$

Here  $0 \leq \kappa \leq V_0 - V_2$  is a constant expressing the size of  $V_1$ ; see SI section 3 for details. This shows that if the concentration  $x \gg 1$  of particles in  $\tilde{\mathfrak{S}}$  is high, then the work scales as the logarithm of the concentration. More importantly though, within any particular equivalence class  $\mathcal{S}$  of systems, where the concentration does not change, the separation work is independent of the system size and not affected by scaling of the volume  $V$ .

### 2.5. Trade-offs involving time

The separation step of the two systems gives rise to two different trade-offs involving time. Firstly, the work estimates calculated above are valid for quasi-static processes. Finite time processes will always require more work [37]. A second trade-off emerges from the waiting time for the ligand-receptor bond to equilibrate during step (i) of the separation procedure. The time scale for equilibration is  $\tau \sim k_u + k_b/V_0$ , where  $k_b$  and  $k_u$  are the binding and binding/dissociation rates to/from the receptor. The equilibrium probability to find the particle unbound is  $p_u = V_0/(V_0 + \eta)$  where  $\eta = k_b/k_u$ . From this, we obtain the number of iterations  $\langle n_\epsilon \rangle$  of the resetting procedure necessary to ensure that the average number of bound ligands is  $\epsilon$ :

$$\langle n_\epsilon \rangle = \left\lceil \frac{\ln\left(\frac{\epsilon}{l}\right)}{\ln(\eta) + \ln V_0 + \eta} \right\rceil, \tag{8}$$

where  $\lceil y \rceil$  denotes the smallest integer  $> y$  and  $l$  is the number of ligand binding sites at the external receptor. Multiplying this by the time scale  $\tau$  gives a time-scale indicator for the restoration of  $\mathfrak{S}$  (see fig. 4b). A trade-off between the measurement time and the work required arises here, because both quantities depend on the volume  $V_0$ . A larger volume makes it more likely for the unbinding receptor to be captured, but equally increases the amount of work that needs to be done (see fig. 4a).

Further trade-offs arise when a concentration needs to be determined with an accuracy higher than binary. To see this, assume that  $d_{\text{max}}$  is the maximal abundance of  $L$  in  $\mathfrak{S}$ . Given a set of perfect binary measurement devices  $\mathfrak{S}_i$  one could use half-interval search to determine the abundance of the particles with an accuracy of  $d_{\text{max}}/N$  by using  $\lceil \log(N) \rceil$  sequential binary measurements. Each binary measurement comes at a fixed cost and takes a finite amount of time. Altogether, this means that a higher accuracy can only be achieved at the expense of a longer waiting time and a higher energy cost. Instead of performing half-interval search, one could alternatively perform a number of concurrent measurements. In this case, there would be no additional time penalty, but  $N$  measurements would be required in order to achieve an accuracy of  $d_{\text{max}}/N$ , which is much worse than the logarithmic scaling of the half-interval search.

## 3. Discussion

Computational processes in biological systems typically show a trade-off between cost, accuracy and speed. Intuitively such trade-offs are expected, but their precise origin remained unclear. Our model shows that the origins of time trade-offs and trade-offs involving accuracy are quite distinct. The key-parameter controlling the accuracy of the computation of an EDC is the system size, which also controls the cost. Yet, in mesoscopic systems the speed of the computation is independent of the system size, but depends only on the kinetic parameters. Trade-offs involving time arise during the process of measuring the outcome of the computation.

One could take the point of view that the measurement process is not an integral part of the computation, and that it is therefore unreasonable to charge this cost to the computation account. Indeed, it is true that there are many stochastic processes that happen in the world where the state of the system is never determined. However, information processing is only useful to the cell when the result is somehow utilised

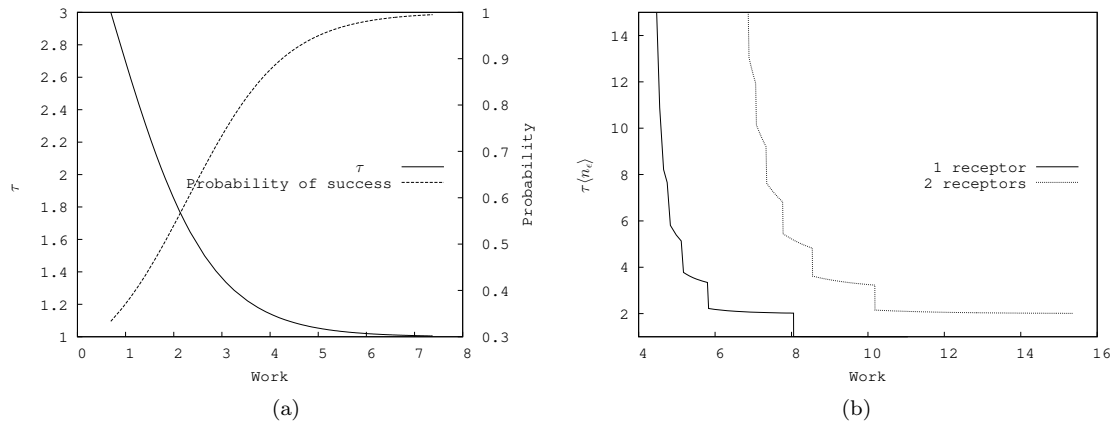


Figure 4: (a) Trade-off between the time and the work spent on the restoration. Here we chose  $k = k_u = 1$  and  $V_1 = V_2 = 0.25$ . The reaction volume  $V_0$  was varied from 0.5 to 200. The line labelled "Probability" denotes the probability that the particle bound to the receptor has been "captured," i.e. pushed to  $V_1$ . (b) Trade-off between the time spent on computation and the work, for 1 and 2 receptors. Here we chose  $\epsilon = 0.01$ ,  $k/k_u = 1$ ,  $N = 200$ ,  $V_{\mathcal{S}} = 10$  and  $V_2 = 1$ .

by the cell. For instance, during chemotaxis a bacterium determines the external concentration of some molecular species and then alters the motion of its flagella. More generally, if any computation is to have an effect on the world at all, then its result must be known. A computation that just happens without the result ever being communicated is thus not a meaningful computation at all. Measurement, therefore, is crucial to computation and cannot be separated.

Within our model, there are two parts to the measurement process. (i) The cost of recording the result of the computation and (ii) the separation/restoration cost. As expected, the cost of recording the measurement does not scale with the size of the measurement system. There are no fundamental reasons to assume that large systems are necessary in order to store information. Indeed, every electronic computer uses microscopic magnetic systems in order to store bits on disc drives. At a first glance more surprising is that the time-cost stems from the time required to separate the system from the measurement device.

The model of separation we used is, of course, not biologically realistic. Nonetheless, it reflects a real biochemical effect. The separation cost is directly related to the work required to break the bond between the external protein and the receptor, which features prominently in various biological incarnations of what is essentially the same effect. A well known example is the classical insight that the unbinding rate of the ligand from the receptors at the cell surface fundamentally limits the ability of cells to measure molecular concentrations in their environment. This is the ultimate origin of the Berg-Purcell limit [15, 38, 24]. Also relevant in this context is [39], where it was shown how in a biological control system trade-offs arise from the dynamics of binding and unbinding of external ligands.

Further trade-offs, beyond the ones described here, arise because the system  $\mathcal{S}$  is subject to noise. This means that a single measurement of the system will only result in a probabilistic answer. In order to increase the confidence that the correct answer was obtained, repeated measurements need to be performed (see for example [20]). In terms of our model, each of these sampling events comes at a measurement cost. Altogether, the accuracy of the biochemical computation can be increased in two ways: Either the EDC is made bigger, and thus noise reduced, while the cost is increased; or, alternatively, the EDC remains constant but its result is sampled more frequently. This increases the time and the cost.

## References

## References

- [1] M. Amos, Cellular Computing, Oxford University Press, 2004.

- [2] G. Seelig, D. Soloveichik, D. Y. Zhang, E. Winfree, Enzyme-free nucleic acid logic circuits, *Science* 314 (5805) (2006) 1585–1588. doi:10.1126/science.1132493.  
URL <http://dx.doi.org/10.1126/science.1132493>
- [3] M. Lakin, A. Phillips, Modelling, simulating and verifying Turing-powerful strand displacement systems, in: *Proceedings of the 17th International Conference on DNA Computing and Molecular Programming, DNA'11*, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 130–144.  
URL <http://dl.acm.org/citation.cfm?id=2042033.2042047>
- [4] A. S. Tsuda S., Zauner K.P., *The Phi-Bot: A Robot Controlled by a Slime Mould*, Springer-Verlag London, 2009, Ch. 10, pp. 213–232.
- [5] A. Friedland, T. Lu, X. Wang, D. Shi, G. Church, J. Collins, Synthetic gene networks that count., *Science* 324 (5931) (2009) 1199–1202. doi:10.1126/science.1172005.  
URL <http://dx.doi.org/10.1126/science.1172005>
- [6] R. Sole, J. Macia, Expanding the landscape of biological computation with synthetic multicellular consortia, *Natural Computing* (2013) 1–13doi:10.1007/s11047-013-9380-y.  
URL <http://dx.doi.org/10.1007/s11047-013-9380-y>
- [7] R. Silva-Rocha, V. de Lorenzo, Implementing an OR-NOT (ORN) logic gate with components of the SOS regulatory network, *Molecular Biosystems* 7 (8) (2011) 2389–2396. doi:10.1039/c1mb05094j.  
URL <http://dx.doi.org/10.1039/c1mb05094j>
- [8] M. Amos, I. Axmann, N. Blüthgen, F. de la Cruz, A. Jaramillo, A. Rodriguez-Paton, F. Simmel, Bacterial computing with engineered populations, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373 (2046) (2015) 20140218. doi:10.1098/rsta.2014.0218.  
URL <https://doi.org/10.1098/rsta.2014.0218>
- [9] S. Walker, H. Kim, P. Davies, The informational architecture of the cell, *Phil. Trans. R. Soc. A* 374 (2063) (2016) 20150057. doi:10.1098/rsta.2015.0057.  
URL <http://dx.doi.org/10.1098/rsta.2015.0057>
- [10] P. Davies, S. Walker, The hidden simplicity of biology, *Reports on Progress in Physics* 79 (10) (2016) 102601. doi:10.1088/0034-4885/79/10/102601.  
URL <http://dx.doi.org/10.1088/0034-4885/79/10/102601>
- [11] J. Ninio, Kinetic amplification of enzyme discrimination., *Biochimie* 57 (5) (1975) 587–595.
- [12] A. Fluit, E. Pienaar, H. Viljoen, Ribosome kinetics and aa-trna competition determine rate and fidelity of peptide synthesis, *Computational Biology and Chemistry* 31 (5-6) (2007) 335–346. doi:10.1016/j.compbiolchem.2007.07.003.  
URL <http://dx.doi.org/10.1016/j.compbiolchem.2007.07.003>
- [13] T. Gregor, D. Tank, E. Wieschaus, W. Bialek, Probing the limits to positional information., *Cell* 130 (1) (2007) 153–164. doi:10.1016/j.cell.2007.05.025.  
URL <http://dx.doi.org/10.1016/j.cell.2007.05.025>
- [14] C. Govern, P. ten Wolde, Optimal resource allocation in cellular sensing systems., *Proceedings of the National Academy of Science USA* 111 (49) (2014) 17486–17491. doi:10.1073/pnas.1411524111.  
URL <http://dx.doi.org/10.1073/pnas.1411524111>
- [15] H. C. Berg, E. M. Purcell, Physics of chemoreception., *Biophys J* 20 (2) (1977) 193–219. doi:10.1016/S0006-3495(77)85544-6.  
URL [http://dx.doi.org/10.1016/S0006-3495\(77\)85544-6](http://dx.doi.org/10.1016/S0006-3495(77)85544-6)
- [16] P. Mehta, D. Schwab, Energetic costs of cellular computation., *Proceedings of the National Academy of Science USA* 109 (44) (2012) 17978–17982. doi:10.1073/pnas.1207814109.  
URL <http://dx.doi.org/10.1073/pnas.1207814109>

- [17] U. Alon, M. Surette, N. Barkai, S. Leibler, Robustness in Bacterial Chemotaxis, *Nature* 397 (1999) 168–171.
- [18] D. Chu, D. Barnes, The lag-phase during diauxic growth is a trade-off between fast adaptation and high growth rate., *Scientific Reports* 6 (2016) 25191. doi:10.1038/srep25191.  
URL <http://dx.doi.org/10.1038/srep25191>
- [19] D. Chu, In silico evolution of diauxic growth., *BMC Evolutionary Biology* 15 (2015) 211. doi:10.1186/s12862-015-0492-0.  
URL <http://dx.doi.org/10.1186/s12862-015-0492-0>
- [20] N. Zabet, D. Chu, Computational limits to binary genes., *Journal of the Royal Society Interface* 7 (47) (2010) 945–954. doi:10.1098/rsif.2009.0474.  
URL <http://dx.doi.org/10.1098/rsif.2009.0474>
- [21] D. Chu, N. Zabet, A. Hone, Optimal parameter settings for information processing in gene regulatory networks, *BioSystems* 104 (2011) 99–108. doi:10.1016/j.biosystems.2011.01.006.  
URL <http://www.cs.kent.ac.uk/pubs/2011/3081>
- [22] G. Lan, P. Sartori, S. Neumann, V. Sourjik, Y. Tu, The energy-speed-accuracy trade-off in sensory adaptation, *Nature Physics* 8 (5) (2012) 422–428.
- [23] M. Johansson, J. Zhang, M. Ehrenberg, Genetic code translation displays a linear trade-off between efficiency and accuracy, *Proceedings of the National Academy of Science USA* 109 (1) (2012) 131–136. doi:10.1073/pnas.1116480109.  
URL <http://dx.doi.org/10.1073/pnas.1116480109>
- [24] K. Kaizu, W. de Ronde, J. Paijmans, K. Takahashi, F. Tostevin, P. ten Wolde, The Berg-Purcell limit revisited., *Biophysical Journal* 106 (4) (2014) 976–985. doi:10.1016/j.bpj.2013.12.030.  
URL <http://dx.doi.org/10.1016/j.bpj.2013.12.030>
- [25] R. Landauer, C. Bennett, R. Laing, W. Zurek, Maxwells Demon, Information Erasure, and Computing, Princeton University Press, 1990, pp. 187–288.  
URL <http://www.jstor.org/stable/j.ctt7zts1p.8>
- [26] C. Bennett, The thermodynamics of computation. A review, *International Journal of Theoretical Physics* 21 (12) (1982) 905–940. doi:10.1007/bf02084158.  
URL <http://dx.doi.org/10.1007/bf02084158>
- [27] R. P. Feynman, A. Hey, Feynman Lectures On Computation, Westview Press, 2000.
- [28] N. van Kampen, Stochastic Processes in Physics and Chemistry, Elsevier, Amsterdam, 2007, third edition.
- [29] H. Qian, D. Beard, Thermodynamics of stoichiometric biochemical networks in living systems far from equilibrium, *Biophysical Chemistry* 114 (2-3) (2005) 213220. doi:10.1016/j.bpc.2004.12.001.  
URL <http://dx.doi.org/10.1016/j.bpc.2004.12.001>
- [30] D. Beard, S. Liang, H. Qian, Energy balance for analysis of complex metabolic networks., *Biophys J* 83 (1) (2002) 79–86. doi:10.1016/S0006-3495(02)75150-3.  
URL [http://dx.doi.org/10.1016/S0006-3495\(02\)75150-3](http://dx.doi.org/10.1016/S0006-3495(02)75150-3)
- [31] D. Gillespie, A rigorous derivation of the chemical master equation, *Physica A: Statistical Mechanics and its Applications* 188 (1-3) (1992) 404–425. doi:10.1016/0378-4371(92)90283-v.  
URL [http://dx.doi.org/10.1016/0378-4371\(92\)90283-v](http://dx.doi.org/10.1016/0378-4371(92)90283-v)
- [32] U. Seifert, Entropy Production along a Stochastic Trajectory and an Integral Fluctuation Theorem, *Physical Review Letters* 95 (4) (2005) 040602. doi:10.1103/PhysRevLett.95.040602.

- [33] M. Kwiatkowska, G. Norman, D. Parker, PRISM: Probabilistic symbolic model checker, in: P. Kemper (Ed.), Proc. Tools Session of Aachen 2001 International Multiconference on Measurement, Modelling and Evaluation of Computer-Communication Systems, 2001, pp. 7–12, available as Technical Report 760/2001, University of Dortmund.
- [34] T. Ouldridge, C. Govern, P. ten Wolde, Thermodynamics of computational copying in biochemical systems, *Physical Review X* 7 (2). doi:10.1103/physrevx.7.021004.  
URL <https://doi.org/10.1103/physrevx.7.021004>
- [35] T. Ouldridge, P. ten Wolde, Fundamental costs in the production and destruction of persistent polymer copies, *Physical Review Letters* 118 (15). doi:10.1103/physrevlett.118.158103.  
URL <https://doi.org/10.1103/physrevlett.118.158103>
- [36] S. Marzen, H. Garcia, R. Phillips, Statistical mechanics of Monod–Wyman–Changeux (MWC) models, *Journal of Molecular Biology* 425 (9) (2013) 1433–1460. doi:10.1016/j.jmb.2013.03.013.  
URL <http://dx.doi.org/10.1016/j.jmb.2013.03.013>
- [37] I. Bena, C. V. den Broeck, R. Kawai, Jarzynski equality for the Jepsen gas, *Europhysics Letters (EPL)* 71 (6) (2005) 879–885. doi:10.1209/epl/i2005-10177-0.  
URL <http://dx.doi.org/10.1209/epl/i2005-10177-0>
- [38] W. Bialek, S. Setayeshgar, Cooperativity, sensitivity, and noise in biochemical signaling., *Physical Review Letters* 100 (25) (2008) 258101.
- [39] D. Chu, Limited by sensing - a minimal stochastic model of the lag-phase during diauxic growth, *Journal of Theoretical Biology* 414 (2017) 137–146. doi:10.1016/j.jtbi.2016.10.019.  
URL <https://doi.org/10.1016%2Fj.jtbi.2016.10.019>