

Fast Decisions with Biophysically Constrained Gene Promoter Architectures

Tarek Tohme , Massimo Vergassola , and Thierry Mora 

*Laboratoire de physique de l'École normale supérieure, CNRS, PSL University, Sorbonne Université,
and Université de Paris, 75005 Paris, France*

Aleksandra M. Walczak 

*Laboratoire de physique de l'École normale supérieure, CNRS, PSL University, Sorbonne Université,
and Université de Paris, 75005 Paris, France
and Applied Physics Research Group, Vrije Universiteit Brussel, Brussels, Belgium*



(Received 12 July 2025; accepted 4 December 2025; published 21 January 2026)

Cells integrate signals and make decisions about their future state in short amounts of time. A lot of theoretical effort has gone into asking how to best design gene regulatory circuits that fulfill a given function, yet much less is known about the constraints that performing that function in a small amount of time imposes on circuit architectures. Using an optimization framework, we explore the properties of a class of promoter architectures that distinguish small differences in transcription factor concentrations under time constraints. We show that the full temporal trajectory of gene activity allows for faster decisions than its integrated activity represented by the total number of transcribed mRNA. The topology of promoter architectures that allow for rapidly distinguishing low transcription factor concentrations results in a low, shallow, and noncooperative response, while at high concentrations, the response is high and cooperative. In the presence of noncognate ligands, networks with fast and accurate decision times need not be optimally selective, especially if discrimination is difficult. While optimal networks are generically out of equilibrium, the energy associated with that irreversibility is only modest and negligible at small concentrations. Instead, our results highlight the crucial role of rate-limiting steps imposed by biophysical constraints.

DOI: [10.1103/sxnt-sbx9](https://doi.org/10.1103/sxnt-sbx9)

I. INTRODUCTION

Cellular decisions about fate and phenotypes are encoded in the states of regulatory biochemical networks. These networks sense the state of the environment, often encoded by concentrations of signaling molecules, such as morphogen gradients in development [1–4], antigens in an immune response [5–9], or sugars and aminoacids in chemotaxis [10–14], and respond by expressing subsets of genes. The ability of cells to precisely and rapidly estimate and respond to these concentrations is tied to the architecture of the readout circuits: promoters and signaling cascades. This readout occurs in an environment with high intrinsic noise [15,16]. A lot is known about the molecular details of specific promoter architectures and optimal circuits for precision, robustness, evolvability, and discrimination potential have been explored to understand the constraints that these functions bring [16–22]. However, less work has been concerned with how the ability to reliably readout concentration can be constrained by limited time.

Time matters for development, immune response, and chemotaxis. Cells commit to different transcriptional programs through a set of decisions about which genes to up and down regulate. However, they do not have a lot of time to measure the environmental cues that provide them with information needed to make these decisions. The design of the readout algorithms is further constrained by the cellular hardware: transcription factors and gene readout complexes have to arrive at the activation sites and assemble—transcription and translation take time.

Exploring different functions has identified many molecular trade-offs. Inspired by early work of Berg and Purcell [10], previous work has asked about the constraints that precision of gene expression imposes on molecular designs [10,22–39]. Others [17,40] have considered the limits that biochemical noise imposes on the readout of positional gradients, leading to the existence of optimal gradient length scales that maximize readout precision. Bialek and coworkers [18,41] proposed the principle of information maximization between the input (the maternally laid Bcd gradient) and the output (the gap genes) in order to explain the precise, experimentally measured expression profile of Hunchback [24] and then other gap genes [41]. Others [42,43] further explored the idea suggested by Gregor *et al.* [24] of spatial averaging of the protein concentrations as a mechanism of noise reduction. As an alternative to fixed-time decision making, Wald's sequential probability ratio test (SPRT) [44] applied to biochemical networks was used to show that receptors [45] and gene

*These authors contributed equally to this work.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

regulatory networks [21] can reach large precision in a shorter time than allowed by Berg-Purcell [10,24,30] or maximum likelihood estimators [34–36].

These examples show us that we can imagine different functional pressures acting on promoters. Different encoding of the signal can lead to different optimal designs, even if we impose the same objective: a fast and precise readout. Here we explore the question whether different encoding of the signal results in different optimal promoter architectures, assuming the signal is decoded in an optimal way. What constraints does limited readout time impose on molecular implementations of transcription? Specifically, focusing on the simplest gene regulatory networks that can function out of equilibrium, we ask how much faster decisions can be made if the whole series of binding and unbinding events is exploited rather than the accumulated product of the transcription process and how that distinction as well as physical constraints affect the architecture of the optimal network.

In the same spirit as the Berg and Purcell scheme [10], and the line of work descending from it [22–39], we adopt the point of view which separates finding the physical limits under which biology operates from identifying the mechanisms it may use to reach these limits. Practically, to obtain performance bounds that no physical scheme can surpass, we study an idealized observer that has access to the full history transcription, with no assumption about whether such an observer is actually realized biologically. Whether regulatory networks decode using the full activity, or an incomplete, distorted, cumulative, or otherwise processed signal, remains an important question. If the process is encoded transcriptionally, then the first stage of this information processing pipeline is the encoding of concentrations in mRNA transcription activity, which is the object of our study.

We start by unifying previous works that interpret the readout of a biochemical network as an on-the-fly decision [21,22,45] and those that view it as a fixed-time estimate of concentration [10,24,30,34–36] under the common framework of Fisher information (first two sections of the results). We further highlight the differences in decoding with accumulated and full-trace evidence. We then study optimal architecture for a single ligand and a single binding site and generalize to multiple binding sites and to the additional presence of spurious ligands. We end with a discussion of the energetic costs associated with optimal networks.

II. RESULTS

A. Decision time in a gene regulation model

We consider a simple nonequilibrium model of gene regulation, where the promoter may be active (on), meaning that it gets bound by RNA-polymerase to make transcripts, or not active (off). It may also be bound or not by a transcription factor [Fig. 1(b)]. This results in a four-state system summarized in Fig. 1(a). The binding rates $k_{b|off}c$, $k_{b|on}c$ of the transcription factor (TF) onto the DNA scale linearly with the TF concentration c . They also depend on the activity state on/off, as do the unbinding rates $k_{u|off}$, $k_{u|on}$. Likewise, the activation $k_{on|b}$, $k_{on|u}$ and deactivation $k_{off|b}$, $k_{off|u}$ rates depend on the binding state. Later we shall consider generalizations

of this model to multiple binding sites for the transcription factor, as well as the possibility of nonspecific binding. Unless otherwise specified, we assume that the system has reached a steady state.

We ask how long an ideal observer of the sequence of activity would need to distinguish between two similar concentrations c and $c + \delta c$ of the transcription factor regulating the gene [21,45]. While there is no guarantee that cells can implement decisions as well as an ideal observer, this puts a hard physical bound on the precision of any process downstream of gene expression [Fig. 1(e)]. Formally, the ideal observer has access to a temporal trace of transcriptional activity [gray lines in Fig. 1(c)], which can be encoded by a sequence of times spent in the active and inactive states up to time t , $O(t) = (T_1^-, T_1^+, T_2^-, \dots)$. For each concentration c , one can define the log-likelihood of these observations given c :

$$\mathcal{L}(c; t) = \ln P(O(t)|c). \quad (1)$$

The SPRT, originally proposed by Wald [44], consists of tracking the log-likelihood ratio between the two competing hypotheses (c and $c + \delta c$):

$$\mathcal{R}(t) = \ln \frac{P(O(t)|c)}{P(O(t)|c + \delta c)} = \mathcal{L}(c; t) - \mathcal{L}(c + \delta c; t). \quad (2)$$

When the true concentration is c , $\mathcal{R}(t)$ follows a stochastic process that is biased towards positive values, reflecting the accumulation of evidence from the observations [Fig. 1(d)]. A decision is made when evidence has reached some threshold of certainty, $|\mathcal{R}(t)| > K \equiv \ln((1 - \varepsilon)/\varepsilon)$, where ε is the error level, which we fix in all numerical experiments at $\int_{|x|>1} dx/\sqrt{2\pi}e^{-x^2/2} = 0.32$, the probability of making an estimate with an error greater than one standard deviation, which is well approximated by a Gaussian distribution in the hard decision regime [21,22]. The observer determines that the concentration is c if $\mathcal{R}(t) > K$ and $c + \delta c$ if $\mathcal{R}(t) < -K$. The time of this decision T_{dec} is itself a random variable that depends on the stochastic trajectory of $\mathcal{R}(t)$.

The smaller δc the harder the two hypotheses are to distinguish and the longer it takes for that evidence to accumulate. In that limit, and under reasonable assumptions that the active and inactive interval durations are not correlated over long times, the law of large numbers allows us to approximate the evolution of the log-likelihood ratio as a random walk [45]:

$$\frac{d\mathcal{R}}{dt} = V + \sqrt{2V}\eta, \quad (3)$$

where $\langle \eta(t)\eta(t') \rangle \sim \delta(t - t')$. This equation may be integrated explicitly, so that $\mathcal{R}(t)$ is normally distributed with mean Vt and variance $2Vt$. Note that the diffusion term is equal to the drift in (3), which is made necessary by the condition $\langle e^{-\mathcal{R}} \rangle = \sum_O P(O|c)P(O|c + \delta c)/P(O|c) = 1$ [21].

Assuming that c is the true concentration, V is as follows:

$$V = \lim_{t \rightarrow \infty} \frac{\langle \mathcal{R}(t) \rangle_c}{t} = \lim_{t \rightarrow \infty} \frac{D_{\text{KL}}[P(O(t)|c) || P(O(t)|c + \delta c)]}{t}, \quad (4)$$

where $\langle \cdot \rangle_c$ denotes an average over $P(O|c)$ and D_{KL} is the Kullback-Leibler divergence. In the limit where decisions are

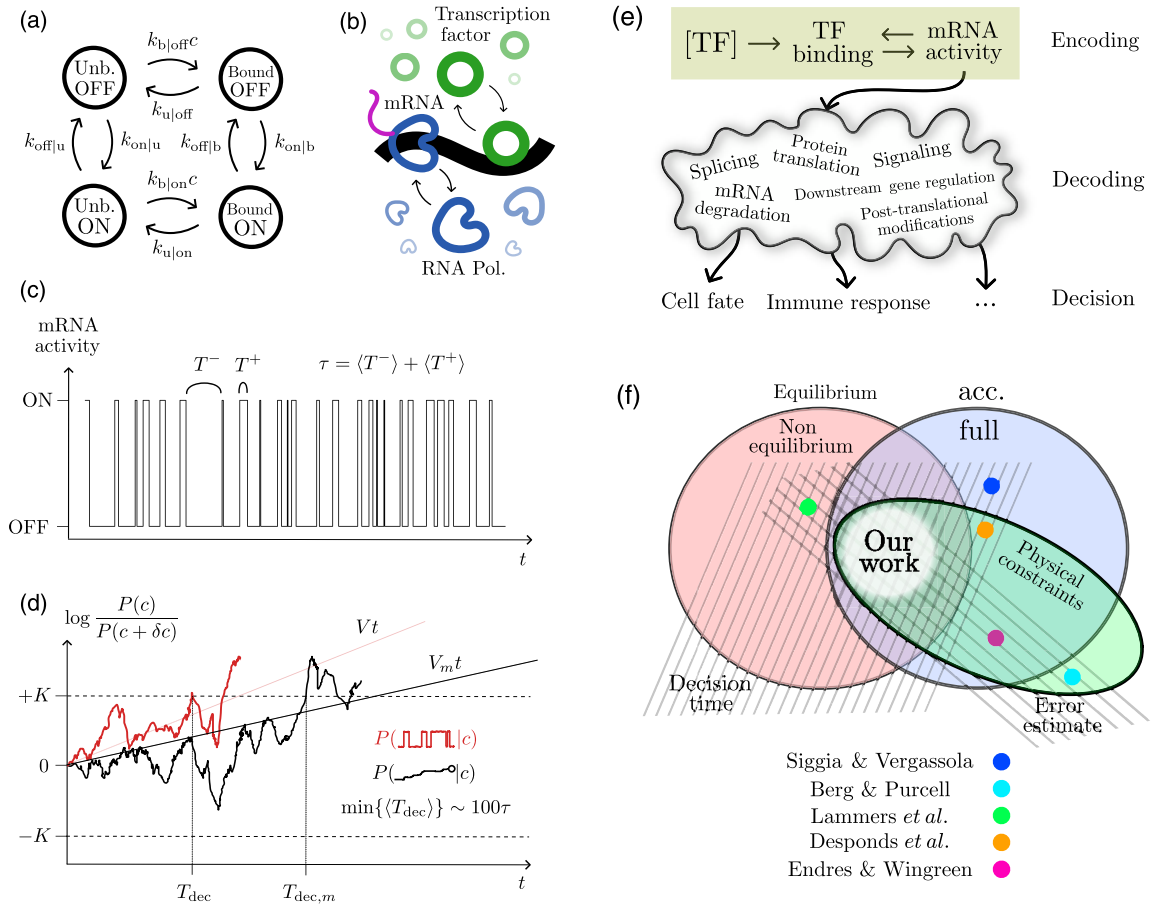


FIG. 1. (a) A model of a receptor to which a ligand can bind and activate a biochemical response (ON-OFF). Binding activity is coupled to the response through eight transition rates. (b) A promoter with one binding site regulating a gene, an example of a system described by (a). (c) Trace of transcription activity over time. (d) Log likelihood ratio for a given activity trace for two readout mechanisms: full activity (red trace) and cumulated activity decoding (black trace), the stochastic SPRT decision times T and T_m , and the SPRT drift terms Vt and $V_m t$ associated with each scheme. Minimum decision times for the circuit in (a) are on the order of 100 activity cycles. (e) A schematic view of the information flow from a transcription factor concentration all the way to a biochemical decision. Our work’s scope is at the level of mRNA transcription activity, focusing on the best possible decoding from that stage. (f) A visual literature review highlighting key connections to previous work. Sensing of concentrations has been considered in equilibrium and nonequilibrium setting, assuming full transcriptional traces (full) and accumulated mRNA concentrations (acc.) as readout. Error estimates and decision times have been calculated in the presence and absence of physical constraints. Here we combine all of these elements and show connections between previous formulations.

hard, $\delta c \ll c$, we may expand this expression at small δc :

$$V \approx \frac{1}{2} \left(\frac{\delta c}{c} \right)^2 \dot{\mathcal{F}}(c), \tag{5}$$

where

$$\dot{\mathcal{F}}(c) = - \lim_{t \rightarrow \infty} \frac{1}{t} \left\langle \frac{\partial^2 \mathcal{L}(c)}{(\partial \ln c)^2} \right\rangle, \tag{6}$$

is the Fisher information rate for the log-concentration. Even for a simple model such as in Fig. 1, that rate does not have a general analytical expression. We developed a numerical procedure based on a combination of Monte Carlo sampling and dynamic programming to compute $\dot{\mathcal{F}}(c)$ efficiently for any Markov system with active states (see Sec. IV A).

The decision time is given by the mean first-passage time for this biased random walk in the log-likelihood space with

decision boundaries set by K . For $\delta c \ll c$, it is [21,45]:

$$\langle T_{\text{dec}} \rangle = \frac{K}{V} \tanh \left(\frac{K}{2} \right). \tag{7}$$

Large deterministic bias V reduces the decision time [Fig. 1(d)]. This result is general and does not depend on the studied problem. In the limit of low errors $\varepsilon \rightarrow 0$, $K \rightarrow \infty$, and $\langle T \rangle_{\text{dec}} \sim K/V$.

Estimates of the decision time (7) assume full knowledge of the activity trace. Alternatively, one may assume that observations are incomplete or reduced. For instance, Lammers *et al.* [22] studied SPRT where observations were reduced to the cumulated time of active transcription, $T_{\text{on}}(t) = \sum_i T_i^+$. Doing so leads to an expression for the reduced Fisher information rate $\dot{\mathcal{F}}_m(c)$, obtained using the same Eq. (6), but using $\mathcal{L}_m(c; t) = \ln P(T_{\text{on}}(t)|c)$ instead of $\mathcal{L} = \ln P(O(t)|c)$. While it only provides a lower bound to the full information rate, Lammers *et al.* [22] showed that it can be computed

analytically as:

$$\dot{\mathcal{F}}_m(c) = \frac{s^2}{\dot{v}} \leq \dot{\mathcal{F}}(c), \quad (8)$$

where

$$s = c \frac{\partial p}{\partial c} \quad (9)$$

is the logarithmic sensitivity of the activation curve $p(c) \equiv P(\text{on}|c) = \lim_{t \rightarrow \infty} T_{\text{on}}(t)/t$ and

$$\dot{v} = \lim_{t \rightarrow \infty} \frac{1}{t} \text{Var}(T_{\text{on}}(t)) \quad (10)$$

is the rate of increase of the variance of the cumulated activity time. Both s and \dot{v} can be computed from the transition rates of the model, using asymptotic formulas for Markov processes [46].

We will compare this partial readout mechanism with the one based on the full likelihood, $P(O|c)$. We will show not only that the latter extracts more information about concentration per unit time, leading to faster decisions under SPRT, but also that minimizing decision time under one decoding scheme or the other yields different optimal gene activation circuits. This highlights how conclusions drawn from information-optimization arguments may depend intricately on the nature of the available information, the details of the functions being optimized, and under which constraints. We discuss these aspects in the next sections.

B. Link to the accuracy of concentration sensing

In their pioneering paper, Berg and Purcell calculated the accuracy with which a receptor can sense the concentration c of a freely diffusing ligand within a fixed time T [10]. Here we show how to link these accuracy bounds to the SPRT framework for decision times.

Within our framework, the question of accuracy may be addressed by computing the error made when estimating c from the on-off observations $O(t)$ using maximum likelihood, $\hat{c}(t) = \arg \max_c P(O(t)|c)$. At long times, the scaling of the error $\delta\hat{c} = \hat{c}(t) - c$ is given by the Cramér-Rao bound, which states that it is governed by the Fisher information rate:

$$\frac{\langle \delta\hat{c}^2 \rangle}{c^2} \approx \frac{1}{T \dot{\mathcal{F}}(c)}. \quad (11)$$

As a result, both the SPRT decision time and the accuracy of concentration estimate at fixed time are governed, in the limit of hard decision and long times, by the same quantity: the Fisher information rate $\dot{\mathcal{F}}(c)$. In other words, making an accurate estimate of concentration on the one hand and making a fast binary decision between two competing hypotheses on concentration on the other, while they are different tasks (making an estimate at fixed time versus making a decision on the fly at variable time) best carried out by different decoding schemes, correspond to the same encoding objective. Having made this link, the rest of the paper adopts the Fisher rate as the quantity of interest, regardless of whether the readout is used to make a decision between two concentrations or to sense a concentration with high precision. In the following sections we shall focus on optimizing that quantity for various choices and constraints on the gene regulation model.

Classical results about readout accuracy [10,35] may be recovered using (11). We can solve analytically a particular case of our simple gene regulatory model, in which the gene is active only when bound and inactive when unbound. That case is achieved by considering the limit $k_{\text{on}|b}, k_{\text{off}|u} \rightarrow \infty$, and $k_{u|\text{off}=0}, k_{b|\text{on}} = 0$. The system then makes transitions from an inactive, unbound to an active, bound state, and back, with effective rates $k_b \equiv k_{b|\text{off}}$ and $k_u \equiv k_{u|\text{on}}$. This model of a two-state system has been widely used, beyond the context of gene regulation, to model the binding of ligands by receptors. The Fisher information rate may be computed analytically [35]:

$$\dot{\mathcal{F}}(c) = (1 - p(c))k_b c, \quad (12)$$

where $p(c) = (1 + k_u/k_b c)^{-1}$ is the probability of being active. Note that this rate is exactly the average number of binding events per unit time, with the $1 - p$ term ensuring that the binding site must be unoccupied for a binding event to occur. Further assuming that the binding rate is diffusion limited, with the binding site modeled as a circle of radius a , and the transcription factor having diffusivity D , we have $k_b = 4Da$ and

$$\frac{\langle \delta\hat{c}^2 \rangle}{c^2} \approx \frac{1}{4Dac(1 - p)t}, \quad (13)$$

which is the result of Ref. [35].

The reduced information rate, where only the total bound time is used, may also be computed exactly using (8), yielding $\dot{\mathcal{F}}_m(c) = (1 - p)k_b c/2$ [22], which is exactly half of the full rate. This leads to an increased error in the estimate of the concentration at fixed time, $\hat{c}_m(t) = \arg \max_c \mathcal{L}_m(c, t)$:

$$\frac{\langle \delta\hat{c}_m^2 \rangle}{c^2} \approx \frac{1}{2Dac(1 - p)t}, \quad (14)$$

which is exactly the result obtained by Berg and Purcell for a single receptor [10].

In summary, \mathcal{F} and \mathcal{F}_m correspond to two estimators previously discussed in the literature: maximum likelihood and mean occupancy. The estimator $\hat{c}_m(t)$ is in fact obtained by inverting the relation $T_{\text{on}}(t)/t = p(c)$, which is the decoding scheme originally studied by Berg and Purcell. The twofold difference in accuracy between using \hat{c} versus \hat{c}_m , which was first noted in Ref. [35], is here reframed in terms of using the likelihood of the full activity trace versus the average activity. Physically, the difference stems from the fact that the information rate $\dot{\mathcal{F}}$ only relies on unbinding times to discern concentrations, while the reduced information rate $\dot{\mathcal{F}}_m$ is subject to noise in both the bound and unbound time periods, the former containing no information about concentration in the two-state system. This doubles the sources of uncertainty, as discussed in detail in Refs. [26,35]. The relation of our results to previous literature is summarized in Fig. 1(f).

Since in this section we have shown that the following definitions are equivalent, in the rest of the paper we will use the terms ‘‘optimal’’ or ‘‘highest precision’’ to mean (i) the minimal error at a given mean decision time, (ii) the shortest mean decision time for given error level, and (iii) the maximum Fisher information rate.

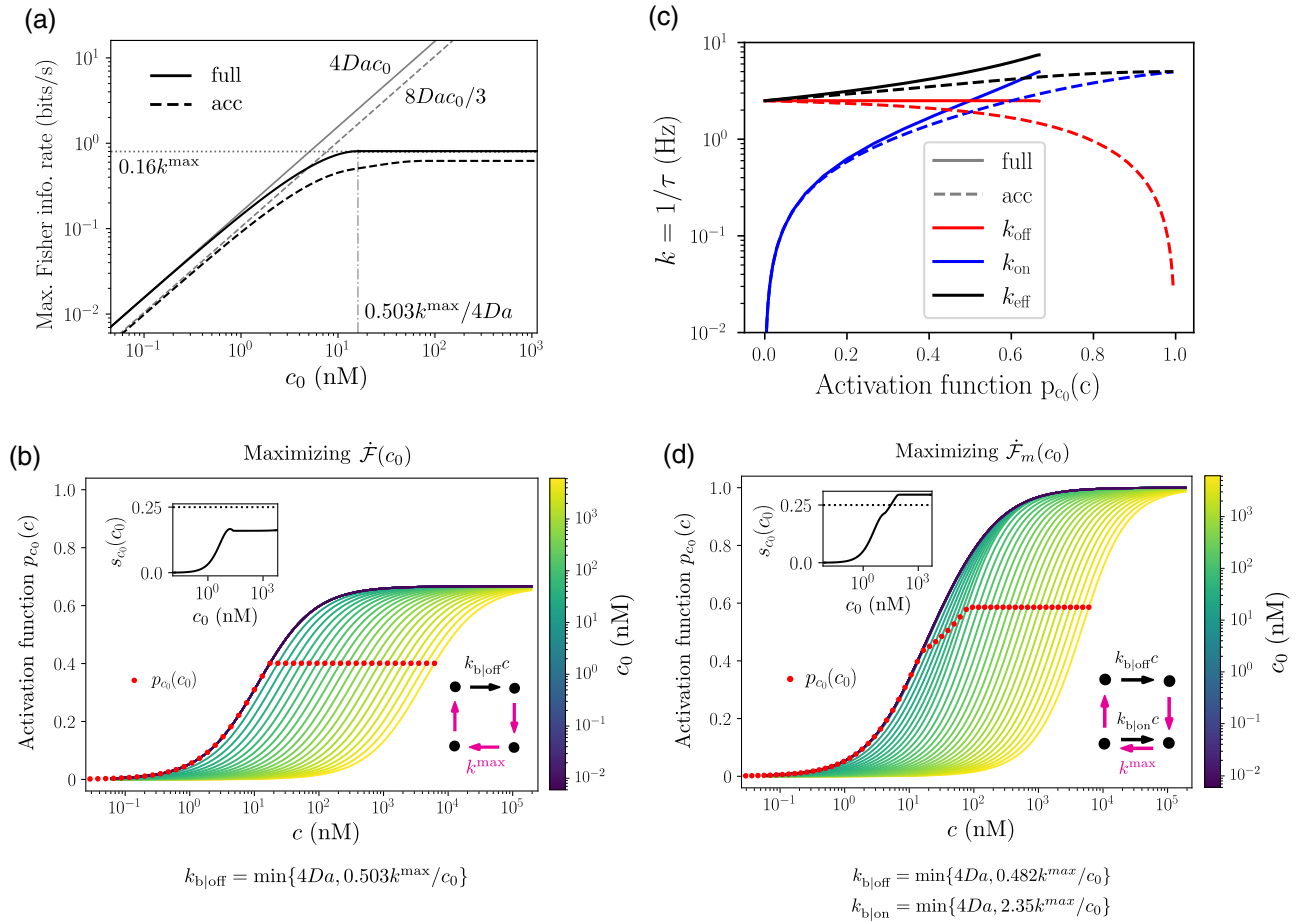


FIG. 2. Minimizing decision time under diffusion, activation and unbinding time limits. (a) Optimal Fisher information rates at different concentrations. D is the diffusion constant of the ligand in square micrometers per second and a is the length of the receptor in micrometers, giving $4Da$ units of $M^{-1}s^{-1}$. k^{max} is the maximum activation rate per second. Throughout the text, labels “full” and “acc” refer, respectively, to the full trace and accumulated mRNA decoding scheme. (b) Activation functions $p_{c_0}(c)$ of a set of circuits maximizing $\dot{\mathcal{F}}(c_0)$ at different concentrations c_0 . Inset: Sharpness of the activation function $s_{c_0}(c_0) = c_0 \partial p_{c_0}(c) / \partial c|_{c_0}$ at the true concentration. Dotted line at 0.25 indicates a Hill coefficient of 1. The decrease compared to the inset of (d) shows that $\dot{\mathcal{F}}$ -optimal circuits do not optimize for sharpness at high concentration. The architecture of optimal circuits is shown on the lower-right corner in (b) and (d). States and arrows follow the layout of Fig. 1(a). Arrows are colored black if their rate equals $k_{b|off}$ and purple if it equals k^{max} . Arrows not shown saturate k_{min} . (c) Effective rates of on ($k_{on} = 1/\tau_{on}$, blue) and off ($k_{off} = 1/\tau_{off}$, red) activity cycles as a function of mean activity $p_{c_0}(c)$ of a set of circuits optimized at $c_0 = 3.36$ nM, evaluated at different values of c . Black lines indicate the correlation time between bursts, $k_{eff} = 1/\tau_{on} + 1/\tau_{off}$. (d) Equivalent figure for $\dot{\mathcal{F}}_m$ -maximizing circuits.

C. Different optimal encodings at high and low concentration

We now set out to find the circuits that best discriminate between transcription factor concentrations around $c = c_0$. This translates into finding the set of transition rates that maximize $\dot{\mathcal{F}}(c_0)$ or, alternatively, $\dot{\mathcal{F}}_m(c_0)$ if one assumes the decoder can only use the cumulated activity.

Decisions can be sped up by increasing indefinitely all the rates to very high values. However, physics imposes constraints on the rates. Binding events are often limited by the diffusion time of the ligand. Diffusion rates in turn are limited by the cellular environment and have accurately measured bounds, e.g., in the context of early fly development [38,47–49]. We estimate upper bounds for unbinding rates from the minimum bound time required to trigger a downstream mechanism, such as activating a gene (which is an assumption made in Ref. [21]). As a result, in a realistic system, all four states have nonzero occupancy probability.

In addition to making the models more physically realistic, these constraints play a central role in determining the optimal encoding resulting from each of the two readout mechanisms we consider.

Throughout the text, unless otherwise noted, optimizations are done under the constraints imposed by ligand diffusion and the minimum time for activation and unbinding involved in transcription. Diffusion sets the upper bound on binding factors: $k_{b|on,off} \leq 4Da$. All other concentration-independent rates are upper-bounded by k^{max} . We also impose a general lower limit k_{min} on all rates for numerical stability.

The maximization of the information rate under these constraints and at a given c_0 reveals two distinct behaviours in the low and high concentration regimes, as illustrated in Fig. 2(a) for both $\dot{\mathcal{F}}$ and $\dot{\mathcal{F}}_m$: first a linear increase of information with c_0 and then a plateau at high c_0 .

The circuits maximizing $\dot{\mathcal{F}}(c_0)$ all have the same cyclic structure [Fig. 2(b), lower inset]. These structures are strongly out of equilibrium, as we will discuss further in Sec. II F. Once bound, the system traverses the (bound, on) state and (unbound, on) states to go back to the starting point (unbound, off) as quickly as possible, maximizing all the corresponding rates at their upper bound, while all reverse rates saturate the lower bound. Within the constraints on rates, this circuit approaches the most closely the two-state system achieving the bound of Eq. (13). The sharp transition observed in Fig. 2(a) is driven by a sharp transition in the optimal binding rate $k_{b|\text{off}}$ as a function of c_0 . When ligands are relatively scarce (small c_0), the waiting time for binding an inactive gene is the rate limiting step, and the binding rate saturates the diffusion-limited bound, $k_{b|\text{off}} = 4Da$. When ligands are abundant (large c_0), TF are no longer constrained by the diffusion limit. However, the optimal binding rate $k_{b|\text{off}}c_0$ saturates to about half of the maximum value of the nonbinding rates, $k_{b|\text{off}}c_0 \approx 0.503k_{\text{max}}$ (see Appendix A for a derivation of the prefactor, which is the solution of a transcendental equation). The Fisher information rate is then $\dot{\mathcal{F}} \approx 0.16k_{\text{max}}$; that limit is plotted as a horizontal line in Fig. 2(a). The reason for this saturation is that, while the optimal circuit seeks to complete binding and activation cycles as quickly as possible, it also tries to remain as sensitive as possible to small changes in concentration and limits its binding rate to avoid being always active. In sum, we have $k_{b|\text{off}}(c_0) = \min(4Da, 0.503k_{\text{max}}/c_0)$, with the transition between the two regimes happening when $c_0 \approx 0.503k_{\text{max}}/4Da$. To summarize, the optimal binding rate is strictly determined by the physical constraints on diffusion and activation time. The latter dominates at high concentration, while the binding rate saturates the former at low concentration.

Across both regimes, the gene activation curve of the network optimized for discriminability around c_0 [Fig. 2(b)] is as follows:

$$p_{c_0}(c) = \frac{2}{3 + k_{\text{max}}/(k_{b|\text{off}}(c_0)c)}. \quad (15)$$

The curves all saturate at $\sim 2/3$, as the network spends a third of its time in the inactive, bound state as $c \rightarrow \infty$. Since binding only occurs in the inactive state, only the mean inactive time, $\tau_{\text{off}} = 1/4Da + 1/k_{\text{max}}$, depends on concentration, while the mean active time, $\tau_{\text{on}} = 2/k_{\text{max}}$, is a constant [Fig. 2(c)]. In the low concentration limit ($4Dac_0 \ll k_{\text{max}}$), the optimal circuit, operating at the concentration c_0 for which it was optimized, spends most of its time in the unbound, off state, $p_{c_0}(c_0) \approx 8Dac_0/k_{\text{max}}$. Since ligands are limited, the system tries to process each binding event as quickly as possible, to make itself available to the next one. In effect, this strategy emulates the two-state system considered previously [Eq. (12)], with effective off rate $k_{\text{max}}/2$. The optimized information rate yields the same value as in that system, $\dot{\mathcal{F}}(c_0) \approx 4aDc_0$, plotted as a gray solid line in Fig. 2(a). When $Dac_0 > 0.503k_{\text{max}}$, the optimal expression sets on $p_{c_0}(c_0) \approx 0.401$ [red dots in Fig. 2(b)]. Note that optimal networks in general do not try to maximize the sensitivity $s = dp_{c_0}/dc(c_0)$, which remains small even at large c_0 [inset of Fig. 2(b)].

The optimal circuits that maximize the reduced Fisher information rate $\dot{\mathcal{F}}_m(c_0)$ are markedly different from those that maximize the full rate $\dot{\mathcal{F}}(c_0)$ [Fig. 2(d), lower inset]. The

key difference is that, over a wide range of concentrations, the binding rate while active maximizes the diffusion bound, $k_{b|\text{on}} = 4Da$. Intuitively, since the reduced observable $T_{\text{on}}(t)$ is “forced” to read off the duration of active periods whether they are informative about c or not and is thus subject to their stochasticity, optimal circuits seek to make the best of it by packing information into those periods. Having a nonvanishing back rate $k_{b|\text{on}}$ introduces a dependence on concentration of the duration of active periods, which circuits can exploit. As a result, the durations of active and inactive periods both depend on c [Fig. 2(c), blue and red dashed lines] but with opposite trends, so that the effective autocorrelation inverse time, $k_{\text{eff}} = 1/\tau_{\text{on}} + 1/\tau_{\text{off}}$, depends much more weakly on activation [Fig. 2(c), dashed black line], in contrast with $\dot{\mathcal{F}}$ -optimal networks (solid black line).

In the low concentration limit, however, $4Dac_0 \ll k_{\text{max}}$, that back rate is negligible and the circuit behaves effectively like an irreversible cycle, just like the optimal circuit for $\dot{\mathcal{F}}(c_0)$. However, its reduced information rate is limited by the stochasticity of active periods, in a similar way to the two-state system discussed earlier [Eq. (14)]. As a result, $\dot{\mathcal{F}}_m(c_0) \approx (2/3)\dot{\mathcal{F}}(c_0)$. Recall that, in the two-state system, the loss factor relative to the full trace $\dot{\mathcal{F}}$ was $1/2$. Here it is only $2/3$ because the active periods are less stochastic than in the two-state system, as they now involve two steps to deactivation: first unbinding, and then deactivation itself, each with rate k_{max} . This implies $\langle T^+ \rangle = 2/k_{\text{max}}$ and $\langle (\delta T^+)^2 \rangle = 2/(k_{\text{max}})^2$, i.e., coefficient of variation $\langle (\delta T^+)^2 \rangle / \langle T^+ \rangle^2 = 1/2$, instead of 1 in the two-state system. As a result, an observer of the cumulated activity is subject to the sum of fluctuations of the binding time, with coefficient of variation of 1, which is also shared by the observer of the full trace, and of the active time, with coefficient of variation of $1/2$, which the observer of the full trace ignores since it contains no information on c . As a result, the total stochasticity in \dot{v} , defined by Eq. (10), is 50% larger than for an observer of the full activity trace; hence the $2/3$ factor in $\dot{\mathcal{F}}$. A more formal derivation is provided in Appendix A.

For higher concentrations, the system goes through two transitions as a function of c_0 . First, for similar reasons as for $\dot{\mathcal{F}}$, the binding rate while inactive $k_{b|\text{off}}$ saturates a bound of $0.482k_{\text{max}}/c_0$. Second, $k_{b|\text{on}}$ knows the same fate but at a higher value, $2.35k_{\text{max}}/c_0$. Those saturations are necessary to keep the activity at c_0 in the range of reasonable sensitivity.

In contrast to the optimal circuits for $\dot{\mathcal{F}}$, the response curves of the $\dot{\mathcal{F}}_m$ -optimal circuits [Fig. 2(d)] have the property $p_{c_0}(c) \rightarrow 1$ as $c \rightarrow \infty$: Since both binding rates are nonvanishing, the system spends almost all its time in the (bound,on) state as $c \rightarrow \infty$. Mean gene activity $p_{c_0}(c_0)$ (red dots) and sensitivity $s(c_0)$ (inset) are both small at low concentrations c_0 and increase for larger ones. In the large concentration regime, gene activity reaches a maximum $p_{c_0}(c_0) \approx 0.586$ and so does the sensitivity $s = 0.295$. That sensitivity is larger than for the $\dot{\mathcal{F}}$ -optimal circuits, and also than the $1/4$ of a classic activation curve obtained in a two-state system (Hill with cooperativity of 1) meaning that these circuits exploit the four-state architecture to implement an effective cooperativity. However, it is still smaller than the maximum sensitivity of $1/2$ (Hill with cooperativity 2) reported in Ref. [22].

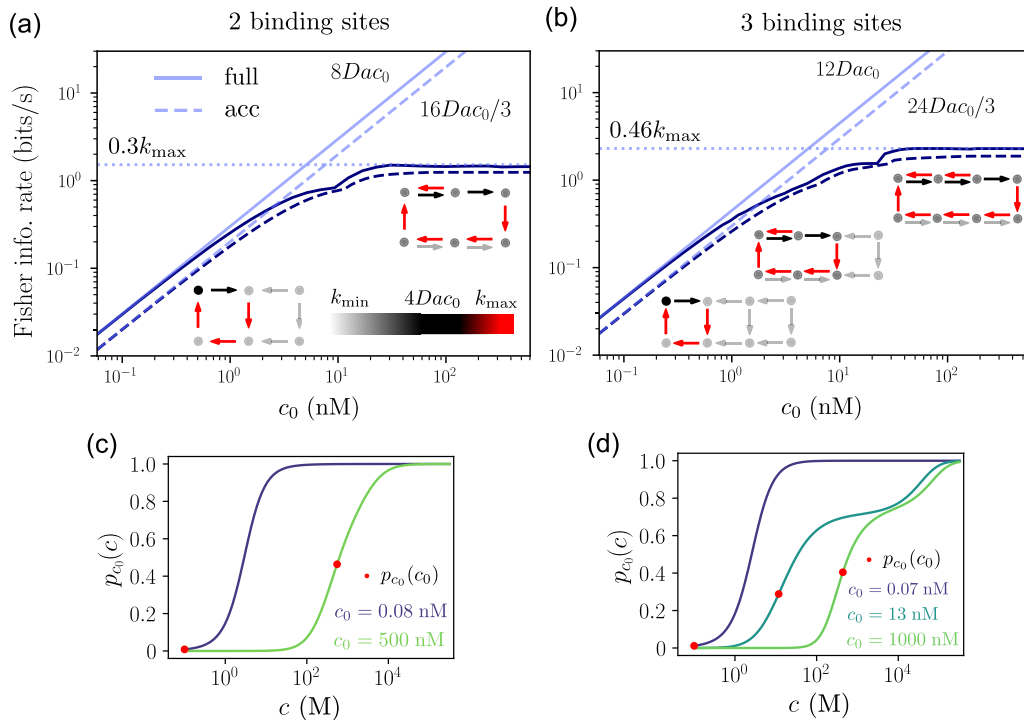


FIG. 3. Onset of cooperativity for optimal receptors with multiple binding sites. (a) Optimal Fisher information rates at concentrations ranging from 0.06 to 600 nM. Full trace decoding indicated by solid lines and accumulated mRNA decoding by dashed lines. D is the diffusion constant of the ligand in square micrometers per second, a is the length of the receptor in micrometers, giving $4Da$ units of $M^{-1}s^{-1}$. k_{max} is the maximum activation rate per second. Optimal promoter architectures for the full trace decoding scheme are shown as insets. Darker shading of states indicates higher occupation probability. On the left-hand side, optimal circuits take on a diffusion-limited, noncooperative configuration where the two sites bind ligands independently, and receptors remain mostly unoccupied. On the right, the Fisher information rate is limited by activation time $1/k_{max}$ and the two sites bind cooperatively. (b) Corresponding plot for three binding sites, showing two transitions to cooperativity. [(c) and (d)] Activation functions of optimal circuits for two and three binding site receptors, respectively, for circuits optimized at two concentrations c_0 corresponding to the regimes of cooperativity outlined in (a) and (b).

To summarize, when ligands are limited, optimal circuits want to stay as inactive as possible to capture the maximum number of binding events. That result holds for both the full observations (\mathcal{F}) or for the cumulated activity (\mathcal{F}_m). In that regime, the only constraint driving the behavior is the diffusion limit, all other constraints being essentially nonlimiting. When ligands are no longer limited, physical constraints on the rates of gene operation take over. The strategy for the full observation is markedly different than for the cumulated activity, favoring less-active circuits with a more shallow response curve, and in which the duration of activity periods does not depend on concentration. By contrast, circuits optimized for cumulated activity exploit the duration of active periods to enhance sharpness and utilize the full range of possible activity levels.

D. Transition to cooperativity with multiple binding sites

The regulatory architecture coupling TF binding and gene activity can be generalized to multiple identical binding sites, by defining binding states corresponding to the number of bound TF, and corresponding binding rates: $k_{1b|off}$ for the first binding event, $k_{2b|off}$ for the second, and so on, and likewise for the binding rates while active, as well as unbinding rates [Fig. 3(a)]. Binding rates are now limited by $4Da$ times the

number of free binding sites and similarly unbinding rates by k_{max} times the number of occupied binding sites.

We optimized \mathcal{F} and \mathcal{F}_m numerically with two binding sites, varying the reference concentration c_0 [Fig. 3(a)]. The optimization revealed three regimes, corresponding to effectively two simple network architectures (insets).

For $c_0 < c_0^* \approx 0.5k_{max}/8Da$, the optimal network allows for only one binding site to be occupied. Binding of either binding site happens with diffusion limited rate $2 \times 4Da$, at which point the gene activated is with the maximal rate k_{max} and then unbound and deactivated as fast as possible. In effect, this solution is equivalent to the single binding site solution in the diffusion limited regime but with a doubled effective binding rate $8Da$ thanks to the existence of two binding sites. The optimal network for cumulated activity time behaves similarly in that limit but with the same $2/3$ factor relative to the full trace, due to the added uncertainty of the duration of the active period.

For $c > c_0^*$, the optimal solution changes drastically to a network in which activity is triggered only once both binding sites have been bound. This design, in which both binding rates while inactive saturate their diffusion-limited bound, corresponds to a maximally cooperative response. A second transition occurs as c_0 is increased and the optimal binding rates cease to be diffusion limited but instead

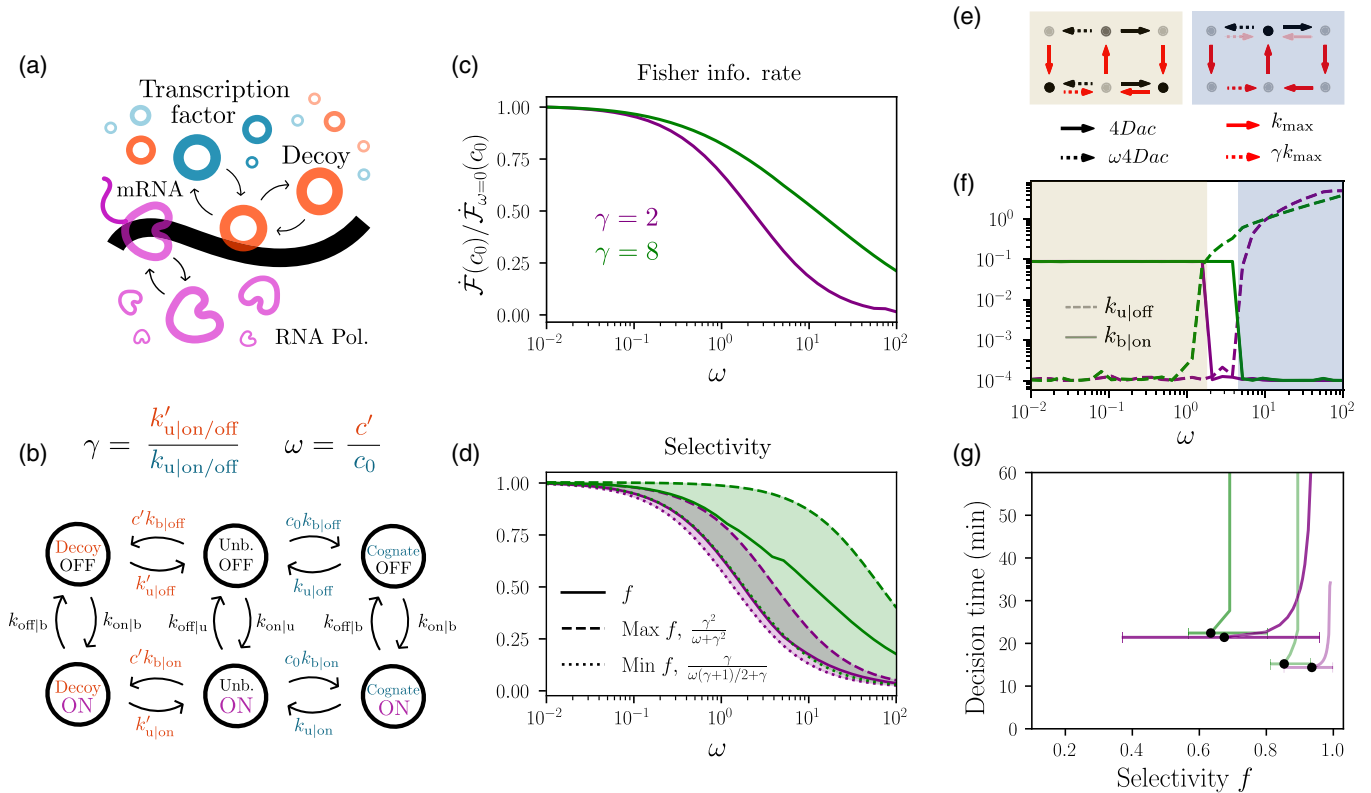


FIG. 4. Maximizing the Fisher information rate in the presence of decoy ligands. (a) A promoter with a single binding site. (b) A general model for a receptor to which both cognate and decoy ligands can bind and activate a downstream response (ON-OFF). The factor γ quantifies how much faster decoys unbind relative to cognate ligands and ω is the ratio of their concentrations, respectively c' and c_0 . (c) Fisher information rate normalized by its value in a solution with no decoys, i.e., $\omega = 0$, for specificities $\gamma \in \{2, 8\}$. (d) Selectivity of the circuit with increasing ω . Circuits that make the fastest decisions in the presence of spurious ligands are neither minimally selective nor maximally selective against the decoys. (e) The two different architectures of optimal circuits in the presence of decoys. Arrows under each circuit indicate transition rates: Diffusion-limited in black and limited by activation-deactivation speed in red. Light red indicates that rates are near the limit. For states, darker shading indicates higher occupation probability. (f) Two key transition rates $k_{u|off}$ and $k_{b|on}$ characterize the circuits shown in (e). (g) Minimum decision time $\langle T_{dec} \rangle$ [Eq. (7)] with a constraint on selectivity, solving $\min_{k_{min} \leq k \leq k_{max}} \{ \langle T_{dec} \rangle + \beta \ln f \}$. Horizontal lines indicate the intervals between $\gamma/(\gamma + \omega(1 + \gamma)/2)$ and $\gamma^2/(\gamma^2 + \omega)$, respectively the minimum and maximum possible selectivity for the model shown in (b), for $\gamma = 2$ and 8 (purple and green) and $\omega = 0.3$ and 3 for $\gamma = 8$ and $\omega = 0.3$ and 1 for $\gamma = 2$ (increasing with color intensity). Black dots indicate the global minima for $\beta = 0$. All simulations are done with $c_0 = 0.6$ nM and $k_{max} = 5$ s⁻¹.

scale with k_{max} . The maximum Fisher information scales with k_{max} .

The result generalizes beyond two binding sites. With three binding sites, the system undergoes two transitions towards cooperativity from a just one binding site being sufficient to activate transcription to two and then three as the reference concentration c_0 is increased [Fig. 3(b)]. We expect similar behaviours for any $k \geq 2$ binding sites: When diffusion limited, the optimal network does not implement cooperativity but rather activates whenever any of the binding site is occupied with rate $k \times 4Da$. When the binding rate becomes comparable to the other rates, the system requires that several binding sites be occupied in order to start transcription, thereby increasing its cooperativity [Figs. 3(c) and 3(d)].

E. Circuits with fastest decision time do not maximize selectivity against decoys

The cellular nucleus typically contains a mixture of transcription factors and other molecules which are not specific

to the promoter's binding sites but may still bind and initiate transcription [Fig. 4(a)]. Assuming a promoter should be adapted to optimally encode information about its cognate ligand into its transcription activity, one may expect spurious binding to hurt that objective. This imposes constraints on how selective the transcription factor binding site may be and how transcription factor concentrations may be distinguished in the face of additional noise coming from the binding of decoy binders. Our next result shows that this is not always the case; under certain conditions promoters may make faster decisions while allowing for spurious binding, even at high concentrations of noncognate ligands.

To model the presence of decoys, we go back to the single binding site and add binding states to the Markov system corresponding to situations where a decoy is bound to the binding site instead of the cognate transcription factor [Fig. 4(b)]. The decoys bind and activate the gene with the same rates $k_{b|on/off}$ as the cognate binder, but they unbind faster, with rates $k'_{u,on/off}$. Decoys are present in concentration c' , which may

be larger than the reference concentration of the transcription factor, c_0 .

The relative abundance of the decoy is set by the ratio $\omega = c'/c_0$. The specificity of the binding site for its cognate transcription factor relative to the decoy is given by the ratio of unbinding rates $\gamma = k'_{u,\text{on/off}}/k_{u,\text{on/off}}$. We numerically searched for networks with maximal \mathcal{F} as a function of these two main parameters, ω and γ . Figure 4(c) shows the information rate \mathcal{F} relative to the case of no decoy ($\omega = 0$) at low c_0 . Increasing the number of decoys (large ω) hurts the precision of concentration estimation; however, increasing specificity γ mitigates this loss, because activation events triggered by decoy binding are then much shorter than by cognate binding, and an ideal observer of the transcriptional activity can use information from the duration of activation events to sort out cognate activation events from decoy ones.

In all the optimal networks that we found numerically, the rate of activation in absence of binding ($k_{\text{on}|u}$) was always best set to its minimal value k_{min} , as these events carry no information about the ligands. Consequently, activation always occurs after binding of either the cognate or decoy ligand. We define the selectivity f as the fraction of time the system is active following a cognate binding event relative to the total active time. As expected, this selectivity decreases with the amount of decoys but increases with specificity, γ [Fig. 4(d)].

In a general equilibrium system, selectivity cannot be higher than $k_{\text{on}c}/(k_{\text{on}c} + k'_{\text{on}c'}) = \gamma/(\omega + \gamma)$. Ninio and Hopfield [50,51] pointed out that selectivity could be improved in nonequilibrium circuits through kinetic proofreading. The idea is to force the system to clear multiple checkpoints, each of which gives an opportunity to test specificity (with efficiency γ), before reaching a productive state. In our circuits, the system allows for two such checkpoints: binding and activation, each of which may be terminated specifically by unbinding, with rates $k_{u,\text{off}}$ and $k_{u,\text{on}}$. An architecture that optimizes these checkpoints can reach a maximal selectivity $f = \gamma^2/(\omega + \gamma^2)$ (see Appendix B). As a point of comparison, it is also useful to consider the simplest nonkinetic proofreading scheme in which all back rates are minimal. Note that while this network is highly nonequilibrium, it corresponds to the extension to decoys of the simplest irreversible loop of Fig. 2(b). Its selectivity, $f = \gamma/(\omega(1 + \gamma)/2 + \gamma)$, is worse than the equilibrium one because of the minimal amount of time $\sim k_{\text{max}}^{-1}$ spent unbound but still active, even when activation was triggered by decoy binding.

We find that in general networks that are optimal for precision do not always make the most of the kinetic proofreading mechanism: Plotting the realized selectivity f in the optimal networks in Fig. 4(d), we find that it is far away from the maximally achievable one (dashed lines), especially if ligands are hard to discriminate ($\gamma = 2$). However, optimal networks are still more selective than the simplest, nonproofreading network (dotted lines).

Different equilibria along the sensitivity-selectivity trade-off are reflected in the architectures of optimal networks, which undergo two transitions with increasing decoy concentration [Figs. 4(e) and 4(f)]. These transitions are governed by two key rates: For large-enough ω , the back rate $k_{u|\text{off}}$ activates the proofreading mechanism [blue shade in Fig. 4(f)], which makes bursts caused by cognate and noncognate binding more

distinguishable. This comes at the expense of longer waiting times between activation bursts and hence lesser throughput. Note that in this regime, unlike in Fig. 2, the Fisher rate does rely on ON times, because burst durations are informative about the identity of the ligand (cognate vs noncognate). However, when $\omega \lesssim \gamma$, selectivity is high enough for the circuit to exploit an additional c_0 dependence of ON durations for concentration sensing by switching on $k_{b|\text{on}}$ [beige shade in Fig. 4(f); see Appendix B]. Although we find this rate saturates its diffusion limit ($4Da$) in optimal circuits and visibly affects f , its contribution to the Fisher rate is negligible when that rate is small compared to the deactivation rate k_{max} .

Conversely, networks with the largest selectivity are far from optimal in terms of precision, as measured by \mathcal{F} . This is illustrated in Fig. 4(g), which shows the best mean decision time, obtained from \mathcal{F} through Eqs. (6) and (7) as a function of a fixed value of the selectivity f (obtained by constrained optimization with Lagrange multipliers, see Appendix C). The fastest decisions are obtained for relatively poor selectivity, relative to what is achievable through kinetic proofreading. The reason why optimal networks do not care much about selectivity is that decoy activations can be statistically distinguished from cognate ones, mitigating the harm of their interference. But kinetic proofreading does come at the cost of reducing the number of fruitful binding events leading to expression. Faced with this trade-off between purity of ligand and throughput, optimal circuits favor throughput.

F. Quantifying irreversibility of optimal networks

All the optimal networks we have discussed so far are strongly out of equilibrium, as they typically follow irreversible paths across states, breaking microscopic time symmetry. How far a circuit is from equilibrium may be quantified rigorously through the entropy production rate:

$$\text{EPR} = \sum_{i,j} \pi_i q_{ij} \ln \frac{\pi_i q_{ij}}{\pi_j q_{ji}}, \quad (16)$$

where q_{ij} denote the transition rate from state i to j and π_i denotes the steady-state distribution of that system, which satisfies $\sum_j (q_{ji}\pi_j - q_{ij}\pi_i) = 0$. The entropy production rate (EPR) multiplied by RT corresponds the minimal amount of energy per mol and per unit time (or power) that must be dissipated to maintain the observed irreversibility.

To assess the importance of irreversibility for precision, in Figs. 5(a) and 5(b) we optimized simple networks for a single binding site with no decoys [Fig. 1(a)] for \mathcal{F} and \mathcal{F}_m , but under a fixed EPR budget, using Lagrange multipliers, for a low [Fig. 5(a)] and a high [Fig. 5(b)] value of the reference concentration c_0 . Note the thermodynamic uncertainty relations [52], which put bounds on the fluctuations of currents as a function of the entropy production rate, do not apply here since \mathcal{F} is not a current.

We see that irreversibility contributes substantially to the performance, as equilibrium networks (EPR = 0, circles) perform 30% to 50% less well than the nonequilibrium ones. To evaluate the order of magnitude of the energy that is required to achieve near-maximum performance, we evaluate the EPR at which the Fisher information rate reaches 90% of its maximal value. This EPR, which is represented as function of c_0 in

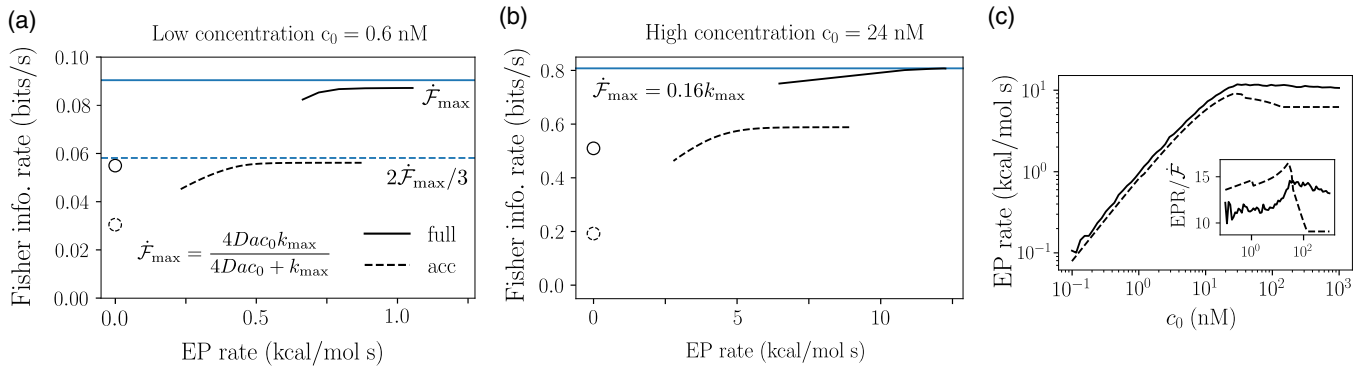


FIG. 5. Entropy production and circuit performance across decoding schemes. (a) Trade-off between EPR and Fisher information rate for the full decoding scheme (solid) and accumulated activity decoding (dashed) at low concentration. Curves delimit Pareto fronts which we compute using the technique of Lagrange multipliers. Circles indicate optimal equilibrium solutions. Increasing the allowable entropy production produces a discontinuous jump from equilibrium to nonequilibrium with no optima in between. Horizontal lines indicate upper limits of each decoding scheme. (b) Analogous plot at high concentration. (c) Entropy production rate at 90% of the maximum over a range of concentrations. Since the absolute maximum EPR depends closely on the lower bounds on rates, which are chosen for numerical convenience, we plot the value at the onset of the plateau shown in (a) and (b). Inset: Entropy production rate per decision versus concentration.

Fig. 5(c), grows with concentration. It goes from vanishingly small consumptions at low c_0 , to ~ 10 kcal/mol/s at large c_0 . The architecture of the optimal nonequilibrium circuit is given in Fig. 2(b), while the optimal equilibrium circuit has nonvanishing rates for all reactions. Intuitively, optimal circuit architectures for decision allow the system to return to the OFF-unbound state as quickly as possible, while passing through the ON states, in which mRNA is produced. Going through this cycle produces entropy. Satisfying the equilibrium requirement means the return to the OFF-unbound state takes longer, allowing for fewer independent measurements of concentration and hence lower precision.

At low concentrations, the gene gets activated rarely, with diffusion-limited rate $4Dac_0$, each time consuming the energy of one activation cycle, which explains the scaling of EPR with c_0 . To obtain a normalized measure of dissipation that accounts for that, we plotted the ratio EPR/\dot{F} , which is proportional to the amount of dissipation per decision [Fig. 5(c), inset]. That normalized dissipation is in fact larger at small concentrations than at high ones, but it peaks at the transition concentration c_0^* between the two regimes.

III. DISCUSSION

Biological systems must often make decisions between discrete fates. An example from early fly development is the establishment of gap gene expression with single-cell resolution, in which small differences in TF concentration result in distinct readouts and gene expression programs [24,38]. Given an unlimited amount of time, it is theoretically possible to distinguish any pair of concentrations as they will result in slightly different mean expression levels. However, development happens on relatively fast timescales, over which observations of gene expression activity are noisy because of the stochastic nature of transcription. The SPRT approach developed in this paper had previously been applied to show that decisions about hunchback expression in the early development of the fly embryo can be made in under 3 minutes [21]. Those calculations assumed perfect coupling

between binding and activation, ignoring the recruitment of the RNA polymerase. Modeling gene activation explicitly and separately from TF binding leads to a technical complication in the computation of the likelihood, due to the fact the durations of active and inactive periods are correlated through the TF binding state, which is hidden from the observer of gene activity. In this paper we showed how to compute the likelihoods necessary for implementing SPRT using dynamic programming. This in turn allowed us to predict the architecture of optimal networks for the task of distinguishing two nearby concentrations of the transcription factor.

We have shown that the information encoded in the full trace of activation and inactivation times is larger than the information encoded in the cumulated time of activation. This result formalizes and generalizes the previous observation that fundamental physical bounds on concentration sensing, first studied by Berg and Purcell [10], could be improved by a factor 2 by using maximum likelihood estimators [34]. To show this, we established a mathematical equivalence between physical bounds on concentration sensing and decision times under the SPRT. We showed that both are governed, in the limit of small concentration differences or equivalently large observation times, by the Fisher information rate \dot{F} . The difference between the two schemes is whether the observer has access to the whole history of gene activation or just the cumulated activity.

This opens the question of how can cells access that historical information and decode the signal. Accumulating mRNA to approximate the total time of activation is one possible decoding mechanism [22]. However, regardless of optimality arguments, we know that cells do not accumulate mRNA: mRNA has a short finite lifetimes, shorter than the division times of cells, even in prokaryotes. Decoding may be done at the next level—protein production, which may itself be subject to active post-transcriptional and post-translational regulations. When these proteins are themselves transcription factors, as it is the case for *hunchback* in fly development, decoding could happen further downstream through the regulation of multiple downstream genes. In sum, many

mechanisms are available to process the raw information contained in the full temporal trace of gene activity, beyond the mere cumulated activation time. Some of these schemes were explored in Desponds *et al.* [21] for simple binding models. In fact, an ideal decoder that computes the log-likelihood ratio $\mathcal{R}(t)$ in the SPRT framework would only need to store as many variables as there are hidden nodes in the forward algorithm (see the Methods section), which in our model is just the number of binding sites + 1. This is true even for full-trajectory decoding; there is never any need to store historical data about mRNA activity. It would be interesting to explore how these mechanisms could be used to make efficient decisions downstream of the optimal networks that we predict.

Our work shares concepts and tools with that of Lamers *et al.* [22], who optimized networks for distinguishing nearby concentrations using SPRT. However, there are three key differences with our approach, which affects the predicted form of the optimal networks. First, we optimized for the SPRT over the full trace, which contains more information than the cumulated activity, as already discussed, but also changes the form of the predicted optimal networks and corresponding responses. Second, in Ref. [22] information rates are normalized by the cycle time, i.e., the typical time it takes to complete an activation event. By contrast, we optimized raw rates expressed in bits per second. Then, without further constraints, the optimal network is always to send all rates to arbitrarily large values to increase information rate. Since rates are finite and often diffusion limited, these schemes lead to unphysical solutions. This brings us to the third difference: We imposed upper bounds on rates of binding, unbinding, activation, and inactivation. Physically, binding rates are limited by the availability of transcription factors and by the speed with which they find their binding site on the DNA through diffusion. Unbinding rates cannot be arbitrarily large if the TF is to bind specifically to DNA. Bounds on (in)activation rates correspond to the minimum bound time required to trigger downstream mechanisms, such as activating a gene, recruiting RNAP, or shutting it off [21]. As it turns out, these constraints are key to shaping the response properties of optimal networks. For instance, we saw that, at low concentrations, optimal networks aim to complete transcriptional bursts as quickly as possible within those constraints, making them inactive most of the time, in contrast to the results of Ref. [22], in which it is optimal to be at half of the maximum activity. The other effect of constraints on rates other than binding is that they limit the ability of nonequilibrium four-state circuits to approach equivalent equilibrium two-state circuits, for which precision bounds were previously derived [10,35]. Another difference is that, unlike Ref. [22], we find that optimal networks do not have sharp response functions.

In the presence of decoys, we find that optimal circuits generically use kinetic-proofreading mechanisms to discriminate cognate from noncognate ligands, consistent with previous theoretical analyses [53]. However, discriminability is only one of the contributor of precision, and optimal circuits balance it against throughput. It may be beneficial to initiate transcription even upon binding the wrong ligands, if it allows for processing more binding events. This is an example of the classical speed-accuracy trade-off, where in our case the role of accuracy is played by specificity. A recent

study, in the very different context of DNA replication [54], argued that this trade-off may be inverted in the high error regime, thanks to the stalling effect: Kinetic proofreading could then evolve from selecting on replication speed only. In the regime of very low mutation rates, however, the usual accuracy-speed trade-off is still predicted to apply. Since there is no stalling in our setting, we find no equivalent effect. Note that, unlike for DNA replication where speed is defined irrespectively of accuracy, in our context speed implies information rate. It is intrinsically linked to the presence of the correct ligand, meaning that the analogy with DNA replication is very limited.

We find that optimal networks are generically out of equilibrium, and that imposing constraints on dissipation (quantified by entropy production) degrades precision. Microscopic analyses of the transcription dynamics at a single gene copy of yeast have shown that the regulatory process for transcription is cyclic and irreversible [55]. Functioning out of equilibrium was suggested as a mechanism for reconciling fast regulatory kinetics and high regulatory specificity. Past work has tried to address how limiting entropy production constrains optimal circuits [21,22,56,57]. However, it worth recalling that the energetic costs incurred by that dissipation are dwarfed by those of mRNA transcription, which involves two phosphate bond hydrolyses per added nucleotide, i.e., of the order of thousands of kilocalories per mole per transcribed mRNA, not counting the costs of synthesizing nucleotides. By contrast, the dissipation caused by irreversible effects in our optimal networks is of the order of 10 kcal/mol at best. This suggests that while being out of equilibrium helps, the amount of energy required to realize that improvement is negligible. This discrepancy between the information-theoretic costs of computation and actual metabolic costs also highlights the fact that the amount of dissipation depends on the level of coarse graining [57–60]. Our simple models do not account for the full complexity of the processes of gene transcription and may thus underestimate the true dissipation. Conversely, it is worth noting that, when coarse-grained, our models may yield very different amounts of energy dissipation. For instance, the irreversible loop of binding-activation-unbinding-deactivation that we find in the simplest setup may be reduced, when (in)activation is fast, to a simpler system with two states—bound-active, unbound-inactive—which has no dissipation at all.

Our optimization for precision makes several predictions that can be compared to data. We find that, when the binding rate of TF to the DNA is limiting, $4Dac_0 \ll k_{\max}$, optimal networks should remain minimally active by completing each activation cycle as soon as possible in order to be available for new informative binding events. This prediction holds even when considering an observer that only has access to the cumulated activity time. However, when kinetic constraints on the rates of other processes are large, so that diffusion is no longer the main limit, the model predicts that genes should be relatively well expressed at the decision point. In the early stages of fly development, genes are substantially active at positions along the anterior-posterior axis where cell fate is decided [61,62]. Our theory would then suggest that the TF binding rate is not the only limiting factor but that the time necessary to complete each activation cycle, $\sim k_{\max}^{-1}$, also limits

precision. It also predicts that, in the presence of multiple binding sites, it is only beneficial to use cooperativity when the binding rates are not limiting ($4Dac_0 \sim k_{\max}$). The fact that the response to maternal gradients is often cooperative [24,38,63] further indicates that TF binding rates may not be the main limiting factor to precision but rather that they are balanced by other rates.

Other biochemical constraints may add to the effect of limits on biochemical rates. A limitation of our approach, also shared by Ref. [22], is the assumption that mRNA production tracks gene activity deterministically. In reality, mRNA transcription is discrete, so that the number of mRNA products is a linear but noisy function of the activity time. While this approximation is acceptable in the regime of reasonably high gene transcription activity, it breaks down when activation periods are short, which is often the case in our optimal networks. With transcription rates of 15–30 mRNAs per minute [62], an optimally fast activation cycle of duration $\sim 2k_{\max}^{-1} < 0.5$ s would produce at most one mRNA transcript per binding event. Transcription would then be very stochastic, which would limit the benefit of short activation cycles predicted by our theory. A possible way to account for this effect would be to constrain the rate of deactivation to very low values, to ensure that each activation event leads to at least a few transcripts, so that $k_{\max} \sim 0.25\text{--}1$ s⁻¹ (instead of 4 s⁻¹ as in the figures). An interesting future direction would be to explore the impact of transcription stochasticity on precision and on the corresponding optimal networks.

Another prediction of our theory is the form of the regulation function and more precisely how on and off times depend on gene activity [Fig. 2(c)]. Experimental data on the gap genes in the fly [62,64] show both on and off times depend on the gene activity but in a way that their harmonic mean is constant. This is consistent with our predictions, provided that we consider the cumulated activity, rather than the full trace, as the output used by the decoder. We take the inverse of ON and OFF times as a proxy for the ON and OFF rates reported experimentally in Refs. [62,64].

More broadly, the question of how cells respond to environmental signals and exploit information can be broken down into two parts: how they encode and store information about the input and how they decode it. Decoding matters, yet the way information is stored already can limit the ability of cells to exploit information and respond. Additionally, different types of cellular hardware may be better suited for specific encoding schemes, as they evolve together. Our results show that physical and structural constraints play an important role in possible encoding schemes in signaling circuits.

IV. METHODS

A. Likelihood of expression traces

To compute decision time, we need to estimate numerically the Fisher information rate defined in (6). In this section we give an outline of how to perform that computation.

The system explores a set of states $\sigma \in \mathcal{A} \cup \mathcal{I}$, where $\mathcal{A} = \{A_i, i = 1, \dots, n\}$ are the active states (transcription is on) and $\mathcal{I} = \{I_i, i = 1, \dots, n\}$ the corresponding set of inactive states (transcription is off). The transition rates between state σ and

σ' are encoded in the matrix $Q(c) = (q_{\sigma',\sigma}(c))$. The likelihood for a particular trajectory of states, where the system spends t_1 in state σ_1 , t_2 in state σ_2 , and so on, reads:

$$\begin{aligned} P((t_1, \sigma_1), (t_2, \sigma_2), \dots | c) \\ = e^{-t_1 \sum_{\sigma} q_{\sigma, \sigma_1}(c)} q_{\sigma_2, \sigma_1}(c) \\ \times e^{-t_2 \sum_{\sigma} q_{\sigma, \sigma_2}(c)} q_{\sigma_3, \sigma_2}(c) \dots \end{aligned} \quad (17)$$

However, this sequence of visited states is not accessible to the transcription state of the system. Instead, the output of the system is the temporal trace of transcriptional activity [Fig. 1(c)], which can be encoded by a sequence of times spent in the active and inactive states, $T_1^-, T_1^+, T_2^-, \dots$, where $T_1^- = \sum_{i \in s_1^-} t_i$, $T_1^+ = \sum_{i \in s_1^+} t_i$, and so on, where s_1^- is the sequence of inactive states first visited before visiting any active state, s_1^+ the following sequence of active states visited before being inactivated again, s_2^- the sequence of inactive states visited after that, and so on.

We need to compute the likelihood of this sequence of residence times for a given concentration c :

$$P(O|c) = P(T_1^-, T_1^+, T_2^-, \dots | c). \quad (18)$$

The likelihood is a sum of the probabilities (17) of all possible scenarios $\{t_i, \sigma_i\}$ that are consistent with the output $(T_1^-, T_1^+, T_2^-, \dots)$. The difficulty is that these durations are correlated with each other, through the particular realizations of the states through which the transitions between active and inactive states happen. Therefore, the likelihood in (30) cannot be written as a product over those transitions.

We can make progress by keeping some, if not all, information about the visited hidden states, namely only the states to which the system transitions when it switches between active and inactive states. Call σ_i^- the first visited state of s_i^- , σ_i^+ the first state of s_i^+ , and so on. Then the likelihood of the trace augmented by that information does factorize:

$$\begin{aligned} P(\{T_i^-, T_i^+, \sigma_i^-, \sigma_i^+\}) \\ = P(\sigma_1^-) \prod_{i \geq 1} P(T_i^-, \sigma_i^+ | \sigma_i^-) P(T_i^+, \sigma_{i+1}^- | \sigma_i^+), \end{aligned} \quad (19)$$

where $P(T_i^-, \sigma_i^+ | \sigma_i^-)$ is the probability that, starting at inactive state σ_i^- , it takes time T_i^- for the system to reach its first active state and that that active state is σ_i^+ ; and similarly for $P(T_i^+, \sigma_{i+1}^- | \sigma_i^+)$. Note that, with that additional conditioning on the entry state, waiting times are now uncorrelated, so that probabilities may be factorized. Then the likelihood (30) may be rewritten as:

$$P(O|c) = \sum_{\sigma_1^-, \sigma_1^+, \sigma_2^-, \dots} P(\{T_i^-, T_i^+, \sigma_i^-, \sigma_i^+\}). \quad (20)$$

This sum involves 2^{2n} terms, where $2n$ is the number of transitions between active and inactive states; however, it can be computed efficiently using dynamic programming approaches. To do this, we define intermediate, conditional sums up to the k th activation cycle:

$$Z_{k, \sigma_k^+}^+(c) = P(T_1^-, T_1^+, \dots, T_k^-, \sigma_k^+ | c), \quad (21)$$

$$Z_{k, \sigma_k^-}^-(c) = P(T_1^-, T_1^+, \dots, T_{k-1}^+, \sigma_k^- | c). \quad (22)$$

Then we have the recursion:

$$Z_{k,\sigma_k^+}^+(c) = \sum_{\sigma_k^-} P(T_k^-, \sigma_k^+ | \sigma_k^-) Z_{k,\sigma_k^-}^-(c) \quad (23)$$

$$Z_{k,\sigma_k^-}^-(c) = \sum_{\sigma_{k-1}^+} P(T_{k-1}^+, \sigma_k^- | \sigma_{k-1}^+) Z_{k-1,\sigma_{k-1}^+}^+(c), \quad (24)$$

and the final likelihood is obtained as the sum over the final (n th) state:

$$P(O|c) = \sum_{\sigma_n^+} Z_{n,\sigma_n^+}^+(c). \quad (25)$$

The conditional probabilities for the on and off durations in Eqs. (23) and (24) may be computed by first-passage time distribution techniques [65]. To recall, $P(T_k^-, \sigma_k^+ | \sigma_k^-)$ is the probability density of the k th period of inactivity T_k^- , the time taken for the state to reach an active state σ_k^+ for the first time, starting from inactive state σ_k^- . It can be obtained by building a matrix $Q^-(c) = (q_{\sigma',\sigma}^-(c))$ in which all rates in the matrix $Q(c)$ that transition out of the active states (A_i) are set to zero, making \mathcal{A} an absorbing set,

$$q_{\sigma',\sigma}^-(c) = \begin{cases} 0 & \text{if } \sigma' \in \mathcal{A} \text{ and } \sigma \in \mathcal{I}, \\ q_{\sigma',\sigma}^-(c) & \text{otherwise.} \end{cases} \quad (26)$$

The transition probability matrix of the system thus obtained reads

$$P^-(t, c) = \exp[tQ^-(c)]. \quad (27)$$

The density of the first-passage time T_k^- can then be read off from entries of the time derivative of (27),

$$P(T_k^-, \sigma_k^+ | \sigma_k^-) = (Q^-(c) \exp[tQ^-(c)])_{\sigma_k^-, \sigma_k^+}. \quad (28)$$

The same procedure can be used, setting the inactive states as absorbing, to obtain $P(T_{k-1}^+, \sigma_k^- | \sigma_{k-1}^+)$.

B. Numerical optimization

Although computing the likelihood of a single trace can be done efficiently as described above, our study aims to optimize the SPRT drift, Eq. (4). This is a trajectory average over observable traces or, since it is a self-averaging quantity, an estimate with a single observation trace of long-enough duration. Obtaining small estimate variances for the purposes of optimization remains computationally expensive for generic circuits but necessary, since one cannot rule out that correlations between bursts might contain information about concentration.

We carried out the optimizations of Figs. 2–5 in two stages. First, the full trajectory average in Eq. (4) was estimated by Monte Carlo sampling with a Julia script, using a single simulated trace $\tilde{O}(t)$ for a large t ,

$$\dot{\mathcal{F}} \simeq \frac{1}{t} \ln \frac{P(\tilde{O}(t)|c)}{P(\tilde{O}(t)|c + \delta c)}. \quad (29)$$

This was in turn maximized numerically using the differential evolution algorithm from Julia's optimization package.

We then observed that $\dot{\mathcal{F}}$ -maximizing circuits always verify $k_{\text{on}lu} = k_{\text{min}} \simeq 0$, which aligns with the intuitive fact that a transcript initiated without TF binding does not carry any information about its concentration. In a second iteration,

we repeated all optimizations imposing this equality, whose consequence is that an ON-OFF burst always starts in the (unbound, OFF) state. In that case, the likelihood may be factorized as a product over each activation cycle:

$$P(O|c) = \prod_i P(T_i^-, T_i^+ | c). \quad (30)$$

Each term of this product may be computed exactly by using (28) with $\sigma_i^- = \text{unbound}$ and summing over $\sigma_i^+ = \text{bound}$ cognate and bound noncognate. This procedure does not require any recursive matrix product or Monte Carlo and is differentiable.

We ran a second round of optimizations with the Fisher information rate instead computed as an exact double integral over ON and OFF durations, yielding a computational speedup of over two orders of magnitude:

$$\dot{\mathcal{F}} = \frac{c_0^2}{\tau} \int_0^\infty \int_0^\infty \frac{(\partial_{c_0} P(T^-, T^+ | c_0))^2}{P(T^-, T^+ | c_0)} dT^- dT^+. \quad (31)$$

This speedup was key in enabling us to explore the effects of varying ω and γ in addition to concentration.

C. Joint Fisher information across cognate and noncognate concentrations

The decoy-prone circuit architectures of Fig. 4 encode information about both c_0 and c' , the concentration of decoy ligands, in the mRNA transcript. In these circuits, measuring the sensitivity to c_0 using the univariate formula for Fisher information (6) amounts to assuming that c' is known perfectly to a downstream observer. Exploiting this knowledge could lead to a better estimation of c_0 , which is unrealistic in real receptors. To rule out such solutions, we modeled the joint uncertainty in both parameters using the inverse Fisher information matrix,

$$\mathbf{C} = \begin{bmatrix} \dot{\mathcal{F}}_{c_0 c_0} & \dot{\mathcal{F}}_{c_0 c'} \\ \dot{\mathcal{F}}_{c_0 c'} & \dot{\mathcal{F}}_{c' c'} \end{bmatrix}^{-1}, \quad (32)$$

with the cross terms given by

$$\dot{\mathcal{F}}_{c_0 c'} = \frac{c_0 c'}{\tau} \int_0^\infty \int_0^\infty \frac{\partial_{c_0} P(c_0, c') \partial_{c'} P(c_0, c')}{P(c_0, c')} dT^+ dT^-. \quad (33)$$

The results of Fig. 4 are then obtained by maximizing

$$\dot{\mathcal{F}}_{c_0} = \frac{1}{[\mathbf{C}]_{c_0 c_0}}, \quad (34)$$

which corresponds to the Fisher information rate for c_0 after accounting for the uncertainty on c' .

ACKNOWLEDGMENTS

The study was supported by Agence Nationale de la Recherche Grant No. ANR-22-CE95-0005-01 ‘‘DISTANT’’ and by the CZI Theory Initiative.

DATA AVAILABILITY

The code and numerical data used to produce the figures are available [66].

APPENDIX A: BOUNDS ON FISHER INFORMATION RATE FOR OPTIMAL FOUR-STATE CIRCUITS (WITHOUT DECOYS)

Our numerical results establish that in the absence of decoys, $\dot{\mathcal{F}}$ -maximizing circuits come in one of two architectures, depending on the decoding scheme. They are shown in the insets of Figs. 2(b) and 2(d). Here we examine the Fisher rate of these architectures in detail, shedding light on the limits of $\dot{\mathcal{F}}$ from Fig. 2(b) at low and high concentration, for accumulated and full activity decoding.

At a given Bicoid concentration c_0 , the parameters of the circuit are the binding rates, $k_{b|\text{off}}c_0$ and $k_{b|\text{on}}c_0$; the rate of unbinding during active transcription $k_{u|\text{on}}$; the rate of activating once Bicoid is bound, $k_{\text{on}|b}$; and, finally, the rate of deactivating once Bicoid unbinds, $k_{\text{off}|u}$. The constraint imposed by diffusion sets $k_{b|\{\text{on},\text{off}\}} \leq 4aD$, with a the linear size of the receptor and D the diffusion constant of Bicoid. All other rates are assumed to take values smaller than k_{max} and we impose a lower limit of 10^{-5} on all rates for numerical stability in simulations.

1. Low concentration regime

The diffusion-limited regime is characterized by few binding events, their rates being small in comparison to the rates of activating and deactivating transcription. This limit simplifies the analysis of multistate circuits, even when computing likelihood ratios of full transcription activity. At low concentration, the Fisher rate of an optimal four-state receptor approaches that of a two-state receptor that activates as soon as a ligand binds and deactivates when it unbinds. In that limit $k_{\text{on}|b} \gg k_{b|\text{off}}c_0$, and OFF times are distributed as $P(T_{\text{off}}) \sim e^{-k_{b|\text{off}}c_0 T_{\text{off}}}$. The distribution of ON times is a convolution of two exponentials, but this complexity is bypassed by noting that ON times are independent of concentration in the optimal architecture and therefore do not contribute to the estimator.

With accumulated decoding, however, a four-state circuit has a 4/3 improvement in sensitivity relative to a two-state architecture and does not reduce to it even at low concentrations. The general expression of the reduced Fisher rate [22] is given by $\dot{\mathcal{F}}_m = s^2/\dot{v}$, where $\dot{v} = \lim_{t \rightarrow \infty} \text{Var}(T_{\text{on}}(t))/t$.

Fluctuations in $T_{\text{on}}(t)$ come from binding and activation events. This was studied in Ref. [36] for a two-state receptor, but their results hold for our setting in the low concentration limit, when times of inactivity are approximately exponentially distributed with rate $k_{b|\text{off}}c_0$ (see Eq. S63 in the Supplemental Material [36]):

$$\dot{v} = [1 + \langle (\delta T^+)^2 \rangle / \langle T^+ \rangle^2] \frac{k_{b|\text{off}}c_0 / \langle T^+ \rangle}{(1/\langle T^+ \rangle + k_{b|\text{off}}c_0)^2}. \quad (\text{A1})$$

The ON binding rate $k_{b|\text{on}}$ is negligible compared to $k_{\text{off}|u}$ at low concentration. The mean time spent ON then reads

$$\langle T^+ \rangle \simeq 1/k_{u|\text{on}} + 1/k_{\text{off}|u}, \quad (\text{A2})$$

and the coefficient of variation of ON time is a sum of the CVs of two exponential waiting times

$$\langle (\delta T^+)^2 \rangle / \langle T^+ \rangle^2 = \frac{(k_{u|\text{on}})^2 + (k_{\text{off}|u})^2}{(k_{u|\text{on}} + k_{\text{off}|u})^2}. \quad (\text{A3})$$

It reaches a minimum of 1/2 when $k_{u|\text{on}} = k_{\text{off}|u} = k_{\text{max}}$, making the leftmost factor on the right-hand side of (A1) equal to 3/2.

Besides we have

$$s^2 = \pi_{\text{on}}^2 (1 - \pi_{\text{on}})^2 = \left[\frac{k_{b|\text{off}}c_0 / \langle T^+ \rangle}{(1/\langle T^+ \rangle + k_{b|\text{off}}c_0)^2} \right]^2, \quad (\text{A4})$$

Putting it all back into $\dot{\mathcal{F}}_m = s^2/\dot{v}$, we get:

$$\dot{\mathcal{F}}_m = \frac{2}{3} \frac{k_{b|\text{off}}c_0 / \langle T^+ \rangle}{1/\langle T^+ \rangle + k_{b|\text{off}}c_0} \simeq 2k_{b|\text{off}}c_0/3 = 2\dot{\mathcal{F}}/3. \quad (\text{A5})$$

2. High concentration regime

Our numerical results indicate that above a precise threshold of concentration c_0 , the optimal OFF binding rate $c_0 k_{b|\text{off}}$ is constant with respect to c_0 for MLE decoding, and both ON and OFF binding rates are constant for BP decoding, as is the Fisher rate. This places the circuit far from any limits that enable two-state simplifications, and so a full analysis of the four-state circuit is required. In this section we perform this analysis for the full trajectory decoding scheme.

We begin by noting that for optimal circuits shown in Fig. 2, the bound inactivation rate $k_{\text{off}|b}$ and the unbound activation rate $k_{\text{on}|u}$ saturate their lower bounds, rendering their effect negligible. This makes activation and inactivation times T_i^+ and T_i^- independent in optimal circuits, and greatly simplifies the likelihood:

$$P(\{T_i^-, T_i^+\}_{i=1}^n | c) = \prod_{i=1}^n P_-(T_i^-) P_+(T_i^+), \quad (\text{A6})$$

$$P_-(t) = \frac{ax}{x-a} (e^{-at} - e^{-xt}), \quad \begin{cases} x = c_0 k_{b|\text{off}} \\ a = k_{\text{on}|b} \end{cases}, \quad (\text{A7})$$

$$P_+(t) = \frac{bc}{s_2 - s_1} (e^{-s_1 t} - e^{-s_2 t}), \quad \begin{cases} b = k_{u|\text{on}} \\ c = k_{\text{off}|u} \end{cases}, \quad (\text{A8})$$

$$s_{1/2} = \frac{1}{2} (-y - b - c \pm \sqrt{\Delta}), \\ \times \begin{cases} \Delta = (y + b + c)^2 - 4bc \\ y = c_0 k_{b|\text{on}} \end{cases}. \quad (\text{A9})$$

As a result of independence, the Fisher information per burst cycle splits into two contributions,

$$\mathcal{F}_{\text{cyc}}(c) = - \left\langle \frac{\partial^2 \ln P_-(t)}{(\partial \ln c_0)^2} \right\rangle_c - \left\langle \frac{\partial^2 \ln P_+(t)}{(\partial \ln c_0)^2} \right\rangle_c \\ = \mathcal{F}_{\text{off}}(c) + \mathcal{F}_{\text{on}}(c). \quad (\text{A10})$$

Moreover, $\mathcal{F}_{\text{on}}(c)$ is negligible in optimal architectures, since ON times are independent of concentration. To obtain the first term, the second derivative of the logarithm of (A7) reads

$$\frac{\partial^2 \ln P_-(t)}{(\partial \ln c_0)^2} = -1 + \left(\frac{x}{x-a} \right)^2 - e^{(x+a)t} \left(\frac{xt}{e^{xt} - e^{at}} \right)^2, \quad (\text{A11})$$

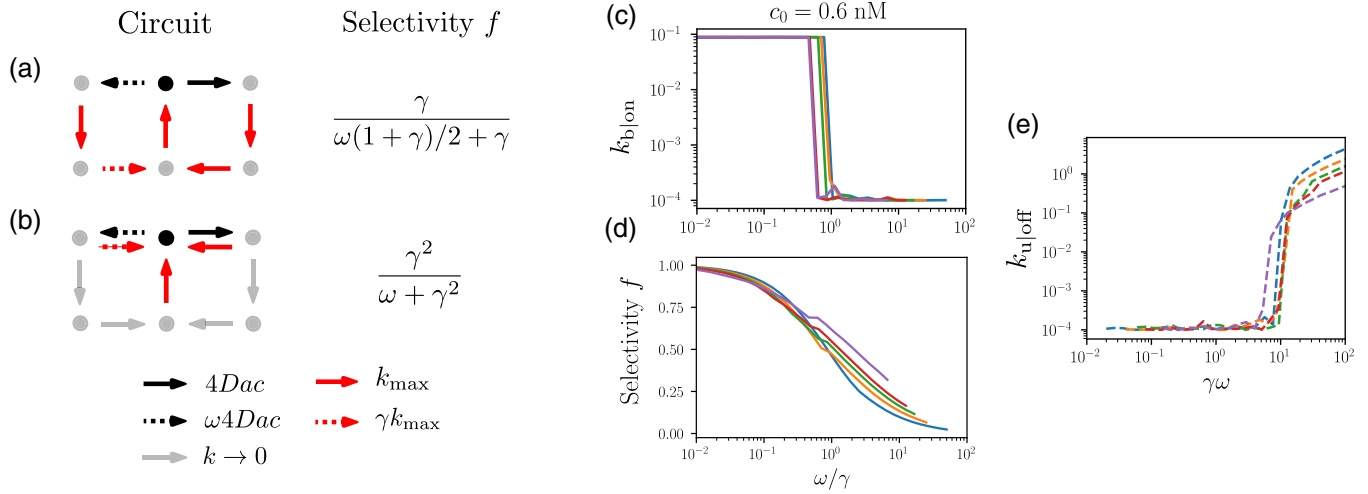


FIG. 6. Minimum and maximum selectivity. [(a) and (b)] Circuit architectures with minimum (a) and maximum selectivity (b). The ON binding rates (c) and selectivity (d) collapse when plotted against ω/γ , hinting that the ON-mediated proofreading mechanism turns off when selectivity is below half, coinciding with $\omega \simeq \gamma$. Colors represent $\gamma \in \{2, 4, 6, 8, 15\}$.

from which the Fisher information of an OFF time is obtained by averaging over $P_-(t)$:

$$\mathcal{F}_{\text{off}}(c_0) = 1 - \left(\frac{x}{x-a}\right)^2 + \frac{ax}{x-a} \int_0^\infty e^{(x+a)t} \times (e^{-at} - e^{-xt}) \left(\frac{xt}{e^{xt} - e^{at}}\right)^2 dt, \quad (\text{A12})$$

$$= 1 - \left(\frac{x}{x-a}\right)^2 + \frac{ax^3}{x-a} \int_0^\infty \frac{t^2}{e^{xt} - e^{at}} dt, \quad (\text{A13})$$

$$= 1 - \left(\frac{x}{x-a}\right)^2 + \frac{ax^3}{x-a} \times \left[\frac{2}{(x-a)^3} \zeta\left(3, \max\left\{\frac{x}{x-a}, \frac{a}{a-x}\right\}\right) \right], \quad (\text{A14})$$

$$= 1 - \alpha^2 + \frac{2a}{x} \alpha^4 \zeta(3, \beta), \quad (\text{A15})$$

where $\alpha = \frac{x}{x-a}$ and $\beta = \alpha$ if $x > a$ and $\beta = 1 - \alpha$ if $x < a$.

We numerically evaluated (A15) in the simulations discussed in the main text. From there, maximizing $\dot{\mathcal{F}}(c_0) = \mathcal{F}_{\text{cyc}}(c_0)/\tau$ over x, y, a, b , and c yields

$$\dot{\mathcal{F}}^*(c_0) = \min\{4aDc_0, \gamma_1 k_{\max}\}, \quad (\text{A16})$$

$$x^* = \min\{4aDc_0, \gamma_2 k_{\max}\}, \quad (\text{A17})$$

$$y^* = k_{\min}, \quad (\text{A18})$$

$$a = b = c = k_{\max}, \quad (\text{A19})$$

$$\gamma_1 \simeq 0.161649, \quad (\text{A20})$$

$$\gamma_2 \simeq 0.503427. \quad (\text{A21})$$

APPENDIX B: PROOFREADING ARCHITECTURES OF THE DECOY-PRONE RECEPTOR

The ability of a receptor to initiate transcription due to the correct ligand (selectivity) in the presence of decoys is limited by its architecture. Two types of proofreading schemes emerge from optimizing the Fisher rate at various decoy concentrations and specificities.

The minimum and maximum selectivity circuits are shown in Figs. 6(a) and 6(b), respectively. We define selectivity as the ratio of time spent transcribing due to a correct binding event. This time includes the residence time in state (on, cognate bound), and the time spent in (on, unbound) weighed by the probability of entering that state due to a cognate binding event.

The upper and lower bounds on f are straightforward to obtain. For Fig. 6(a),

$$f = \frac{2/k_{\max}}{2/k_{\max} + \omega(1/k_{\max} + 1/\gamma k_{\max})} = \frac{\gamma}{\omega(\gamma + 1)/2 + \gamma}, \quad (\text{B1})$$

In Fig. 6(b), the ideal proofreading scheme makes the event of entering a transcribing state due to decoy binding weigh ω/γ relative to cognate binding, instead of ω . This limit of maximum selectivity is only achieved when $k_{\text{on}|b}, k_{\text{u}|on} \rightarrow 0$, where the Fisher information is also zero. In addition, time spent in the (on, decoy bound) state has weight $1/\gamma$ relative to (on, cognate bound). This leads to odds ratio of finding the system in the active state with the cognate ligand rather than the decoy to be γ^2/ω , leading to $f = \gamma^2/(\omega + \gamma^2)$.

When selectivity is high, ON burst times depend strongly on TF concentration; an optimal circuit is therefore incentivized to use them as an information source for sensing c_0 . This is mediated by $k_{b|on}$, which activates this concentration at a transition point around $f \simeq 0.5$, which corresponds to $\omega \simeq \gamma$ [Figs. 6(c) and 6(d)].

Moreover, we find that OFF-mediated proofreading becomes active when $\gamma\omega$ crosses a threshold [Fig. 6(e)]. OFF

times carry more information about concentration when burst time statistics differ sufficiently between cognate and decoy binding events. This is controlled primarily by γ . At high decoy concentrations, optimal circuits also have an incentive to proofread to reduce the total time spent transcribing due to the wrong binding events.

APPENDIX C: OPTIMIZATION WITH CONSTRAINTS ON SELECTIVITY AND ENTROPY PRODUCTION RATE

To explore the sensitivity-selectivity trade-off directly we solve the following constrained optimization problem:

$$\max_{k_{\min} \leq k_b \leq 4aD, k_{\min} \leq k_{[u, \text{on}, \text{off}]} \leq k_{\max}} \dot{\mathcal{F}}(c_0) + \beta \ln f. \quad (\text{C1})$$

Solutions are represented in Fig. 4(g), with $\dot{\mathcal{F}}$ substituted by the SPRT decision time by applying Eq. (7), at fixed concentration c_0 and sweeping the trade-off parameter β from 0 to $\tan(\pi/2 - 0.05) \simeq 20$.

The trade-off between sensitivity and entropy production rate in Fig. 5 was obtained with

$$\max_{k_{\min} \leq k_b \leq 4aD, k_{\min} \leq k_{[u, \text{on}, \text{off}]} \leq k_{\max}} \dot{\mathcal{F}}(c_0) - \beta \text{EPR} \quad (\text{C2})$$

for β ranging from 10^{-7} to 0.06 in logarithmic scale.

-
- [1] W. Driever and C. Nüsslein-Volhard, A gradient of bicoid protein in *Drosophila* embryos, *Cell* **54**, 83 (1988).
- [2] L. S. Mosby, A. E. Bowen, and Z. Hadjivasiliou, Morphogens in the evolution of size, shape and patterning, *Development* **151**, dev202412 (2024).
- [3] J. L. Christian, Morphogen gradients in development: From form to function, *WIREs Dev. Biol.* **1**, 3 (2012).
- [4] J. Briscoe and S. Small, Morphogen rules: Design principles of gradient-mediated embryo patterning, *Development* **142**, 3996 (2015).
- [5] O. Feinerman, J. Veiga, J. R. Dorfman, R. N. Germain, and G. Altan-Bonnet, Variability and robustness in T cell activation from regulated heterogeneity in protein levels, *Science* **321**, 1081 (2008).
- [6] G. Altan-Bonnet and R. N. Germain, Modeling T cell antigen discrimination based on feedback control of digital ERK responses, *PLoS Biol.* **3**, e356 (2005).
- [7] P. Francois, G. Voisinne, E. D. Siggia, G. Altan-Bonnet, and M. Vergassola, Phenotypic model for early T-cell activation displaying sensitivity, specificity, and antagonism, *Proc. Natl. Acad. Sci. USA* **110**, E888 (2013).
- [8] S. R. Achar *et al.*, Universal antigen encoding of t cell activation from high-dimensional cytokine dynamics, *Science (NY)* **376**, 880 (2022).
- [9] G. Altan-Bonnet, T. Mora, and A. M. Walczak, Quantitative immunology for physicists, *Phys. Rep.* **849**, 1 (2020).
- [10] H. Berg and E. Purcell, Physics of chemoreception, *Biophys. J.* **20**, 193 (1977).
- [11] R. M. Macnab and D. E. Koshland, The gradient-sensing mechanism in bacterial chemotaxis, *Proc. Natl. Acad. Sci. USA* **69**, 2509 (1972).
- [12] V. Sourjik and N. S. Wingreen, Responding to chemical gradients: Bacterial chemotaxis, *Curr. Opin. Cell Biol.* **24**, 262 (2012).
- [13] C. Janetopoulos and R. A. Firtel, Directional sensing during chemotaxis, *FEBS Lett.* **582**, 2075 (2008).
- [14] Y. Tu, Quantitative modeling of bacterial chemotaxis: Signal amplification and accurate adaptation, *Annu. Rev. Biophys.* **42**, 337 (2013).
- [15] M. B. Elowitz, Stochastic gene expression in a single cell, *Science* **297**, 1183 (2002).
- [16] G. Tkačik and A. M. Walczak, Information transmission in genetic regulatory networks: A review, *J. Phys.: Condens. Matter* **23**, 153102 (2011).
- [17] F. Tostevin, P. R. ten Wolde, and M. Howard, Fundamental limits to position determination by concentration gradients, *PLoS Comput. Biol.* **3**, e78 (2007).
- [18] G. Tkacik, C. G. Callan, and W. Bialek, Information flow and optimization in transcriptional regulation, *Proc. Natl. Acad. Sci. USA* **105**, 12265 (2008).
- [19] G. Lan, S. Schulmeister, V. Sourjik, and Y. Tu, Adapt locally and act globally: Strategy to maintain high chemoreceptor sensitivity in complex environments, *Mol. Syst. Biol.* **7**, 475 (2011).
- [20] G. Lan, P. Sartori, S. Neumann, V. Sourjik, and Y. Tu, The energy–speed–accuracy trade-off in sensory adaptation, *Nat. Phys.* **8**, 422 (2012).
- [21] J. Desponds, M. Vergassola, and A. M. Walczak, A mechanism for *hunchback* promoters to readout morphogenetic positional information in less than a minute, *eLife* **9**, e49758 (2020).
- [22] N. C. Lammers, A. I. Flamholz, and H. G. Garcia, Competing constraints shape the nonequilibrium limits of cellular decision-making, *Proc. Natl. Acad. Sci. USA* **120**, e2211203120 (2023).
- [23] W. Bialek and S. Setayeshgar, Physical limits to biochemical signaling, *Proc. Natl. Acad. Sci. USA* **102**, 10040 (2005).
- [24] T. Gregor, D. W. Tank, E. F. Wieschaus, and W. Bialek, Probing the limits to positional information, *Cell* **130**, 153 (2007).
- [25] T. Mora and I. Nemenman, Physical limit to concentration sensing in a changing environment, *Phys. Rev. Lett.* **123**, 198101 (2019).
- [26] M. Carballo-Pacheco *et al.*, Receptor crosstalk improves concentration sensing of multiple ligands, *Phys. Rev. E* **99**, 022423 (2019).
- [27] V. Singh and I. Nemenman, Universal properties of concentration sensing in large ligand-receptor networks, *Phys. Rev. Lett.* **124**, 028101 (2020).

- [28] V. Singh and I. Nemenman, Simple biochemical networks allow accurate sensing of multiple ligands with a single receptor, *PLoS Comput. Biol.* **13**, e1005490 (2017).
- [29] G. Malaguti and P. R. ten Wolde, Theory for the optimal detection of time-varying signals in cellular sensing systems, [arXiv:1902.09332](https://arxiv.org/abs/1902.09332).
- [30] K. Kaizu *et al.*, The Berg-Purcell limit revisited, *Biophys. J.* **106**, 976 (2014).
- [31] P. Sartori and Y. Tu, Free energy cost of reducing noise while maintaining a high sensitivity, *Phys. Rev. Lett.* **115**, 118102 (2015).
- [32] P. Sartori and Y. Tu, Noise filtering strategies in adaptive biochemical signaling networks, *J. Stat. Phys.* **142**, 1206 (2011).
- [33] A. Eldar, D. Rosin, B. Z. Shilo, and N. Barkai, Self-enhanced ligand degradation underlies robustness of morphogen gradients, *Dev. Cell* **5**, 635 (2003).
- [34] R. G. Endres and N. S. Wingreen, Maximum likelihood and the single receptor, *Phys. Rev. Lett.* **103**, 158101 (2009).
- [35] R. G. Endres and N. S. Wingreen, Accuracy of direct gradient sensing by cell-surface receptors, *Prog. Biophys. Mol. Biol.* **100**, 33 (2009).
- [36] T. Mora and N. S. Wingreen, Limits of sensing temporal concentration changes by single cells, *Phys. Rev. Lett.* **104**, 248101 (2010).
- [37] B. Hu, W. Chen, W. J. Rappel, and H. Levine, Physical limits on cellular sensing of spatial gradients, *Phys. Rev. Lett.* **105**, 048104 (2010).
- [38] H. Tran *et al.*, Precision in a rush: Trade-offs between reproducibility and steepness of the *hunchback* expression pattern, *PLoS Comput. Biol.* **14**, e1006513 (2018).
- [39] J. Estrada, F. Wong, A. DePace, and J. Gunawardena, Information integration and energy expenditure in gene regulation, *Cell* **166**, 234 (2016).
- [40] T. E. Saunders and M. Howard, Morphogen profiles can be optimized to buffer against noise, *Phys. Rev. E* **80**, 041902 (2009).
- [41] M. D. Petkova, G. Tkačik, W. Bialek, E. F. Wieschaus, and T. Gregor, Optimal decoding of information from a genetic network, *Cell* **176**, 844 (2019).
- [42] Y. Okabe-Oho, H. Murakami, S. Oho, and M. Sasai, Stable, precise, and reproducible patterning of bicoid and hunchback molecules in the early *Drosophila* embryo, *PLoS Comput. Biol.* **5**, e1000486 (2009).
- [43] T. Erdmann, M. Howard, and P. R. ten Wolde, Role of spatial averaging in the precision of gene expression patterns, *Phys. Rev. Lett.* **103**, 258101 (2009).
- [44] A. Wald, Sequential tests of statistical hypotheses, *Ann. Math. Stat.* **16**, 117 (1945).
- [45] E. D. Siggia and M. Vergassola, Decisions on the fly in cellular sensory systems, *Proc. Natl. Acad. Sci. USA* **110**, E3704 (2013).
- [46] W. Whitt, Asymptotic formulas for Markov processes with applications to simulation, *Oper. Res.* **40**, 279 (1992).
- [47] T. Gregor, E. F. Wieschaus, A. P. McGregor, W. Bialek, and D. W. Tank, Stability and nuclear dynamics of the bicoid morphogen gradient, *Cell* **130**, 141 (2007).
- [48] A. Abu-Arish, A. Porcher, A. Czerwonka, N. Dostatni, and C. Fradin, High mobility of bicoid captured by fluorescence correlation spectroscopy: Implication for the rapid establishment of its gradient, *Biophys. J.* **99**, L33 (2010).
- [49] A. Porcher and N. Dostatni, The bicoid morphogen system, *Curr. Biol.* **20**, R249 (2010).
- [50] J. J. Hopfield, Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity, *Proc. Natl. Acad. Sci. USA* **71**, 4135 (1974).
- [51] J. Ninio, Kinetic amplification of enzyme discrimination, *Biochimie* **57**, 587 (1975).
- [52] A. Puglisi, A brief introduction to the inference of entropy production through the fluctuations of currents, *Europhys. Lett.* **150**, 67001 (2025).
- [53] S. A. Cepeda-Humerez, G. Rieckh, and G. Tkačik, Stochastic proofreading mechanism alleviates crosstalk in transcriptional regulation, *Phys. Rev. Lett.* **115**, 248101 (2015).
- [54] R. Ravasio *et al.*, A minimal scenario for the origin of nonequilibrium order, [arXiv:2405.10911](https://arxiv.org/abs/2405.10911).
- [55] R. Shelansky *et al.*, Single gene analysis in yeast suggests nonequilibrium regulatory dynamics for transcription, *Nat. Commun.* **15**, 6226 (2024).
- [56] F. Mancini, M. Marsili, and A. M. Walczak, Trade-offs in delayed information transmission in biochemical networks, *J. Stat. Phys.* **162**, 1088 (2016).
- [57] P. Szymańska-Rożek, D. Villamaina, J. Miekisz, and A. M. Walczak, Dissipation in non-steady state regulatory circuits, *Entropy* **21**, 1212 (2019).
- [58] A. Crisanti, A. Puglisi, and D. Villamaina, Nonequilibrium and information: The role of cross correlations, *Phys. Rev. E* **85**, 061127 (2012).
- [59] A. Puglisi, S. Pigolotti, L. Rondoni, and A. Vulpiani, Entropy production and coarse graining in Markov processes, *J. Stat. Mech.* (2010) P05015.
- [60] D. M. Busiello, J. Hidalgo, and A. Maritan, Entropy production for coarse-grained dynamics, *New J. Phys.* **21**, 073004 (2019).
- [61] T. Lucas *et al.*, Live imaging of bicoid-dependent transcription in *Drosophila* embryos, *Curr. Biol.* **23**, 2135 (2013).
- [62] B. Zoller, S. C. Little, and T. Gregor, Diverse spatial expression patterns emerge from unified kinetics of transcriptional bursting, *Cell* **175**, 835 (2018).
- [63] J. O. Dubuis, G. Tkacik, E. F. Wieschaus, T. Gregor, and W. Bialek, Positional information, in bits. *Proc. Natl. Acad. Sci. USA* **110**, 16301 (2013).
- [64] P. T. Chen, M. Levo, B. Zoller, and T. Gregor, Gene activity fully predicts transcriptional bursting dynamics, *Nat. Struct. Mol. Biol.* **32**, 1959 (2025).
- [65] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, University Series in Undergraduate Mathematics, No. 210 (van Nostrand, Princeton, NJ, 1960).
- [66] <https://github.com/statbiophys/ConstrainedPromoters>.