

Transient Memory in Gene Regulation

Calin Guet², Thomas A. Henzinger², Claudia Iglér², Tatjana Petrov^{1(✉)},
and Ali Sezgin³

¹ Department of Computer and Information Sciences/Centre for the Advanced Study
of Collective Behaviour, University of Konstanz, 78464 Konstanz, Germany
tatjana.petrov@gmail.com

² IST Austria, Am Campus 1, 34000 Klosterneuburg, Austria

³ Aselsan, Ankara, Turkey

Abstract. The expression of a gene is characterised by its transcription factors and the function processing them. If the transcription factors are not affected by gene products, the regulating function is often represented as a combinational logic circuit, where the outputs (product) are determined by current input values (transcription factors) only, and are hence independent on their relative arrival times. However, the simultaneous arrival of transcription factors (TFs) in genetic circuits is a strong assumption, given that the processes of transcription and translation of a gene into a protein introduce intrinsic time delays and that there is no global synchronisation among the arrival times of different molecular species at molecular targets.

In this paper, we construct an experimentally implementable genetic circuit with two inputs and a single output, such that, in presence of small delays in input arrival, the circuit exhibits qualitatively distinct observable phenotypes. In particular, these phenotypes are long lived transients: they all converge to a single value, but so slowly, that they seem stable for an extended time period, longer than typical experiment duration. We used rule-based language to prototype our circuit, and we implemented a search for finding the parameter combinations raising the phenotypes of interest.

The behaviour of our prototype circuit has wide implications. First, it suggests that GRNs can exploit event timing to create phenotypes. Second, it opens the possibility that GRNs are using event timing to react to stimuli and memorise events, without explicit feedback in regulation. From the modelling perspective, our prototype circuit demonstrates the critical importance of analysing the transient dynamics at the promoter binding sites of the DNA, before applying rapid equilibrium assumptions.

Tatjana Petrov's research was supported by SNSF Advanced Postdoc. Mobility Fellowship grant number *P300P2_161067*, the Ministry of Science, Research and the Arts of the state of Baden-Württemberg, and the DFG Centre of Excellence 2117 'Centre for the Advanced Study of Collective Behaviour' (ID: 422037984). Claudia Iglér is the recipient of a DOC Fellowship of the Austrian Academy of Sciences. Thomas A. Henzinger's research was supported in part by the Austrian Science Fund (FWF) under grant Z211-N23 (Wittgenstein Award).

1 Introduction

The fundamental conceptual breakthroughs related to how a gene is turned on and off, have inspired a large body of theoretical and experimental work on gene regulation, including the explanation of stochastic switching between lysis and lysogeny of phage [25], all the way to more complex logic gate formalisms that attempt to abstract more complex biological behaviour. Synthetic biology enthusiasts often use analogies with how electronic circuits are manipulated by computers [13, 24], and have demonstrated success in engineering simple genetic circuits that are encoded in DNA and perform their function *in vivo*. However, such digital (in the sense that the expression states are encoded through Boolean values) and combinational design (in the sense that the output is a pure function of present input only, different to the *sequential* design) quickly becomes infeasible in experiment, because the cellular environment is resource-limited and highly crosstalk-prone. The effective engineering of biological systems needs to take into account the intrinsic properties of the biological medium, so as not to fight against the principles of tinkering that characterise biology [16], but rather to make use of them. Significant conceptual challenges remain related to the still unsatisfactory quantitative but also qualitative understanding of the underlying processes [20]. Understanding time-dependent phenomena is fundamental in this complex picture of the cell that unravels itself at the molecular scale, especially since cells do not have computer-like clocking mechanisms, beyond circadian and cell cycle ones. A major question emerges as to what are the macroscopic effects of small delays in the arrival times of different molecules at molecular targets.

Gene expression in a single cell is modelled by a stochastic process which captures the stochastic switching among possible configurations at the DNA (the architectural configuration of which is often termed *promoter logic*, e.g. shown in Fig. 4), and their effect on the copy number of other species involved in regulation, such as mRNA, proteins and transcription factors (TFs). The switching mechanism depends on the binding affinities of the TFs and RNA polymerase to their respective binding sites as well as the concentrations of those proteins in the cell. Stochastic dynamics of such gene regulatory process typically has a single equilibrium, as a consequence of reversibility of reactions occurring at the DNA binding sites. Sometimes, the transient regime of the distribution among DNA configurations is rapid and robust to possible delays in arrival of TFs. In such cases, it is satisfactory to use the statistical thermodynamics model, which has shown unquestionable success (e.g. [1, 37]). It estimates the probability of being in any of the possible DNA binding configurations from their relative binding energies (Boltzmann weights) and the protein concentrations, both of which can often be experimentally accessed. While this model takes into account the stochasticity inherent to the DNA binding configurations (unlike the also widely used deterministic limit [18]), it neglects the transient probabilities in

the promoter logic before the equilibrium is reached. The question arises: In which ways does the transient regime at the DNA (promoter) affect the shape and duration of observable protein dynamics? Can it happen that the observable transients move towards the unique equilibrium so slowly, that they are mistaken for steady state dynamics?

In this paper, as a proof-of-concept, we construct a prototype genetic circuit based on two different transcription factors that regulate the same gene, without feedback. Our circuit demonstrates that, for gene regulation, qualitatively distinct transients may take extraordinarily long times to disappear, and the observable phenotype in the transient can be highly sensitive to the order of arrival of TFs in the system. In particular, the transient phenotype may appear to be stable even though it is not, creating an effect of *long lived transients*[7]. Our prototype circuit is realistic, experimentally implementable in the sense that the mechanism can be implemented by the current technology and kinetic rate values are in realistic ranges. The behaviour of this circuit suggests that the genetic circuit can memorise the order of arrival of TFs, although there is no explicit feedback at the gene regulatory level.

2 Preliminaries and Background

A gene is expressed at a basal rate, whenever the RNA polymerase (RNAP) is bound to its promoter region at the DNA. *Activators* are transcription factors (TFs) that bind to specific locations on the DNA, or to other TFs, and enhance the expression of gene g by promoting the binding of RNAP. *Repressors* reduce the expression of gene g , by directly blocking the binding of RNAP, or indirectly, by inhibiting the activators, or promoting direct repressors. The mechanism of how and at which rates the molecular species are interacting is transparently written in a list of reactions. Reactions are equipped with the stochastic semantics which is valid under mild assumptions [12].

Definition 1. A reaction system; is a pair (S, R) , such that $S = \{S_1, \dots, S_n\}$ is a finite set of species, and $R = \{r_1, \dots, r_r\}$ is a finite set of reactions. The state of a system can be represented as a multi-set of species, denoted by $\mathbf{x} = (x_1, \dots, x_n) \in \mathbf{N}^n$. Each reaction is a triple $r_j \equiv (\mathbf{a}_j, \boldsymbol{\nu}_j, c_j) \in \mathbf{N}^n \times \mathbf{N}^n \times \mathbb{R}_{\geq 0}$, written down in the following form:

$$a_{1j}S_1, \dots, a_{nj}S_n \xrightarrow{c_j} a'_{1j}S_1, \dots, a'_{nj}S_n, \text{ such that } \forall i. a'_{ij} = a_{ij} + \nu_{ij}.$$

The vectors \mathbf{a}_j and \mathbf{a}'_j are often called respectively the *consumption* and *production* vectors due to j th reaction, and c_j is the respective *kinetic rate*. If the j th reaction occurs, after being in state \mathbf{x} , the next state will be $\mathbf{x}' = \mathbf{x} + \boldsymbol{\nu}_j$. This will be possible only if $x_i \geq a_{ij}$ for $i = 1, \dots, n$.


Stochastic Semantics. The species multiplicities follow a continuous-time Markov chain (CTMC) $\{X(t)\}_{t \geq 0}$, defined over the state space $S = \{\mathbf{x} \mid \mathbf{x}$

is reachable from \mathbf{x}_0 by a finite sequence of reactions from $\{r_1, \dots, r_r\}$. In other words, the probability of moving to the state $\mathbf{x} + \boldsymbol{\nu}_j$ from \mathbf{x} after time Δ is

$$P(X(t + \Delta) = \mathbf{x} + \boldsymbol{\nu}_j \mid X(t) = \mathbf{x}) = \lambda_j(\mathbf{x})\Delta + o(\Delta),$$


with λ_j the propensity of j th reaction, assumed to follow the principle of mass-action: $\lambda_j(\mathbf{x}) = c_j \prod_{i=1}^n \binom{x_i}{a_{ij}}$. The binomial coefficient $\binom{x_i}{a_{ij}}$ reflects the probability of choosing a_{ij} molecules of species S_i out of x_i available ones.

Example 1 (basal gene expression). Basal gene expression with RNAP binding can be modelled with four reactions, where the first reversible reaction models binding between the promoter site at the DNA and the polymerase, and the second two reactions model the protein production and degradation, respectively:



The state space of the underlying CTMC $S \cong \{0, 1\} \times \{0, 1, 2, \dots\}$, such that $s_{(1,x)} \in S$ denotes an active configuration (where the RNAP is bound to the DNA) with $x \in \mathbf{N}$ protein copy number, as depicted in Fig. 2.

Example 2 (adding repression). Repressor blocking the polymerase binding can be modelled by adding a reaction



In this case, there are three possible promoter configurations, that is, $S \cong \{\text{DNA}, \text{DNA.RNAP}, \text{DNA.}R\} \times \{0, 1, 2, \dots\}$, where $D_0 = \{\text{DNA}, \text{DNA.}R\}$ are inactive promoter states.

Computing the Transient. Using the vector notation $\mathbf{X}(t) \in \mathbf{N}^n$ for the marginal of process $\{X(t)\}_{t \geq 0}$ at time t , we can compute this transient distribution by integrating the *chemical master equation* (CME). Denoting by $p^{(t)}(\mathbf{x}) = P(\mathbf{X}(t) = \mathbf{x})$, the CME for state $\mathbf{x} \in \mathbf{N}^n$ reads

$$\frac{d}{dt}p^{(t)}(\mathbf{x}) = \sum_{j=1, \mathbf{x}-\boldsymbol{\nu}_j \in S}^r \lambda_j(\mathbf{x} - \boldsymbol{\nu}_j)p^{(t)}(\mathbf{x} - \boldsymbol{\nu}_j) - \sum_{j=1}^r \lambda_j(\mathbf{x})p^{(t)}(\mathbf{x}). \quad (1)$$


The solution may be obtained by solving the system of differential equations, but, due to its high (possibly infinite) dimensionality, it is often statistically estimated by simulating the traces of $\{X_t\}$, known as the stochastic simulation algorithm (SSA) in chemical literature [12]. As the statistical estimation often remains computationally expensive for desired accuracy, for the case when the deterministic model is unsatisfactory due to the low multiplicities of many molecular species [19], different further approximation methods have been proposed, major challenge to which remains the quantification of approximation accuracy (see [36] and references therein for a thorough review on the subject).

2.1 Transients in Gene Expression Without Feedback

We will further focus on regulation of single gene without feedback. This allows a circuit-view, where activators and repressors are inputs, and the average transient protein expression is the output. Since there is a single DNA molecule per cell, each state counts one copy of the current DNA configuration, and zero copies of all other DNA binding configurations. Hence, the expression state for a single gene of interest consists of two layers: the proteins that we see, and the regulatory configuration of the DNA (for example, two activators and polymerase are bound) (Fig. 3). Such two-layer model allows us to study the transient of coupled promoter state and protein count. In order to focus our analysis on the effect of input timing perturbations, yet to keep our model simple, we chose not to involve further mechanistic details, such as the steps involving mRNA.

The following observation states that in general, when there is no feedback, computing the output does not require integrating the Master equation for the entire CTMC, but only for a CTMC controlling the switching among the DNA configurations (depicted left in Fig. 3).

Lemma 1. Let $\{X(t)\}_{t \geq 0}$ be the CTMC for a model of single gene regulation without feedback, over the state space $S = S_0 \uplus S_1 = (D_0 \uplus D_1) \times \{0, 1, \dots\}$, where $D_0 = \{D_{01}, D_{02}, \dots\}$ are inactive DNA configurations, and $D_1 = \{D_{11}, D_{12}, \dots\}$ are active DNA configurations (RNAP bound). Let the reaction system; (S, R) be such that all reactions are of one of the following types (for some $i \geq 0$ and $j \geq 0$):



Then, the average amount of protein in a population follows the differential equation

$$\frac{d}{dt} \langle x_p(t) \rangle = P(\mathbf{X}(t)|_{\mathbb{D}} \in D_1) \alpha - \beta \langle x_p(t) \rangle, \quad (2)$$

where $\langle x_p(t) \rangle$ denotes the average amount of the protein molecules at time t , and process $\mathbf{X}(t)|_{\mathbb{D}}$ is the projection of process $\mathbf{X}(t)$ to states at the promoter, that is $\mathbf{X}(t)|_{\mathbb{D}} = d$ if and only if $(\mathbf{X}(t) \in \cup_{i \geq 0} (d, i))$. In other words, $P(\mathbf{X}(t)|_{\mathbb{D}} \in D_1)$ denotes the marginal probability that the promoter is in active state (bound RNAP) at time t .

The proof is discussed in Appendix, Sect. 1.C.

Corollary 1. Let $\pi_1 = \lim_{t \rightarrow \infty} P(\mathbf{X}(t)|_{\mathbb{D}} \in D_1)$ denote the probability of active promoter at stationarity. Then, whenever the initial probability equals that of the stationary, i.e. $P(\mathbf{X}(0)|_{\mathbb{D}} \in D_1) = \pi_1$, the average protein dynamics follows the differential equation

$$\frac{d}{dt}\langle x_p(t) \rangle = \pi_1 \alpha - \beta \langle x_p(t) \rangle. \quad (3)$$

When DNA is modelled with one binding site, the promoter can be in only two states, and the analytic solution to Eq. (2) is tractable. In general, as activators and repressors bind to different regions (operator sites) of the same DNA molecule, the respective number of regulatory configurations at the promoter grows combinatorially with the number of operator sites. For instance, one hypothesised mechanism in λ -phage, containing only three left and three right operators, leads to 1200 different DNA configurations [35]¹! In such cases, the simplification based on the argument of fast equilibrium is often employed, meaning that the transient protein dynamics is computed according to Eq. (3), thus neglecting the transient changes in probability distribution among the DNA regulatory configurations.

Fast equilibrium assumption is a prerequisite to applying a widely popular *statistical thermodynamics* model [38]. Assuming that the DNA regulatory configurations mix rapidly, this model allows to experimentally estimate the free energies of each promoter configuration, and then, subsequently, to derive the equilibrium constants² for each of the reactions [28,38]. As the absolute and precise values of kinetic rates are rarely available in practice, this method is powerful, because it allows to predict the dynamics of a genetic circuit from a scarcely available experimental data. However, the statistical thermodynamics model is applicable only when the assumption of rapid equilibrium at the promoter is valid.

In the following, we showcase a simple, experimentally realisable genetic circuit which demonstrates an interesting situation where the long transient at the promoter creates phenotypes that are qualitatively distinct from the phenotypes created when the promoter configurations start at the equilibrium. In particular, these phenotypes are long lived transients: they all converge to a single value, but so slowly, that they seem stable for an extended time period, longer than typical experiment duration.

3 Problem Statement

We focus on a single gene regulation without feedback, where activators and repressors are inputs, and the average protein expression is the output. Assuming that a fixed amount of activators and repressors are added to the system with a possible time lag, our reference scenarios are (Fig. 4 in Appendix):

- $\mathcal{X}_{A||R}$ in which activator and repressor are introduced together,
- $\mathcal{X}_{A \rightarrow R}(\Delta)$ in which the activator is introduced Δ time units before the repressor, and
- $\mathcal{X}_{R \rightarrow A}(\Delta)$ in which the repressor is introduced Δ time units before the activator.

¹ Models used in this paper will count 23 and 6 distinct DNA binding configurations.

² The ratio between the binding and unbinding rate.

Our goal is to construct a genetic circuit with the following requirements: (i) it is realistic, that is - experimentally implementable in the sense that the mechanism can be implemented by the current technology and kinetic rate values are in realistic ranges, (ii) the scenarios provide striking differences in the shape and duration of transient protein output. To quantify the latter, we introduce two quantitative measures:

- *amplitude*, the maximum distance in phenotype of the scenario with delay from the scenario without delay, that is

$$\alpha_{|s} := \max_{t \geq t_0} |\langle x_p \rangle(t|s) - \langle x_p \rangle(t|\mathcal{X}_{A||R})|,$$

where s refers to the scenario in question ($\mathcal{X}_{A \rightarrow R}(\Delta)$ or $\mathcal{X}_{R \rightarrow A}(\Delta)$) and $x_p(t|s)$ denotes the average protein number in a population at time t in scenario s , and

- *half-life*, the time the system takes from the moment of reaching the amplitude, to the moment when the distance from the phenotype without delay disappears, that is

$$t_{1/2|s} := \arg \min_{t \geq t_{\alpha|s}} \{t \mid |\langle x_p \rangle(t|s) - x_p(t|\mathcal{X}_{A||R})| < \frac{1}{2}\alpha_{|s}\},$$

where $t_{\alpha|s}$ denotes the moment when the amplitude is reached in scenario s .

In summary, the *amplitude* reflects how observable is the sensitivity to the delay among inputs, and the second measure, *half-life*, reflects how slow is the convergence to the real equilibrium after the amplitude has been observed. Long lived transients are characterised by a large amplitude relative to the basal expression and a half-life exceeding several cell division cycles.

4 Searching for Long Lived Transients

We develop and analyse models for two promoter architectures (drafted in Fig. 5):

- **Model without looping.** A basic mechanism for activation and repression is assumed: repressor R competes with RNAP, and the activator A recruits the polymerase RNAP and binds independently of the repressor and the polymerase (configurations shown in Fig. 6).
- **Model with looping.** In the model with looping (Fig. 5, right), two activators and two repressors can bind the DNA; Binding of the second activator (resp. repressor) promotes looping of the DNA in the active (resp. repressed) state, thereby excluding binding of the other TF. This small mechanistic change leads to the blow-up of the state space of the CTMC to 23 states as a composition of two sub-models (Fig. 7).

Biological Context. The CTMC for the DNA switching of our prototype circuit with looping is inspired by the very well-characterized regulatory mechanisms of the *lac* operon and of the bacteriophage lambda genetic switch [17, 30].

Mathematical Context. In Markov chains, the time to be ϵ -close to equilibrium, *mixing time*, varies depending on the initial distribution, the chain connectivity, and the rate parameters. In particular, long mixing times are prominent for chains with a large spectral gap of the underlying generator matrix, and can be guaranteed for chains with large connectivity diameter, suggesting that more states and sparse connectivities generally can prolong the mixing time [21]. Still, tightly estimating bounds on the mixing time for a given chain is an open problem, beyond the scope of this manuscript. Intuitively, DNA looping architecture is a good candidate for creating large mixing times, because the looped states are quick to reach when only activators or only repressors are present, but, once entered, they are then hard to exit ('dynamically trapped states').

4.1 Model Implementation

Implementation. The models are written and analysed within the rule-based modelling framework Kappa which allows us to represent the mechanistic model concisely and to run an efficient stochastic simulation algorithm [3, 11]. Source-code of the rule-based models is given in Appendix 1.B. Parameter exploration and additional output analysis were performed with Python.

Simulation. We simulated multiple samples of the stochastic model, and we statistically estimated the first two moments of protein expression. In the model with looping, we used 1000 individual cells for a time of 36000 s = 10 h, that equals around 20 average cell doubling times, where inputs are added from time point $t_0 = 5400 = 1.5$ h (see Table 2 for all simulation parameters).

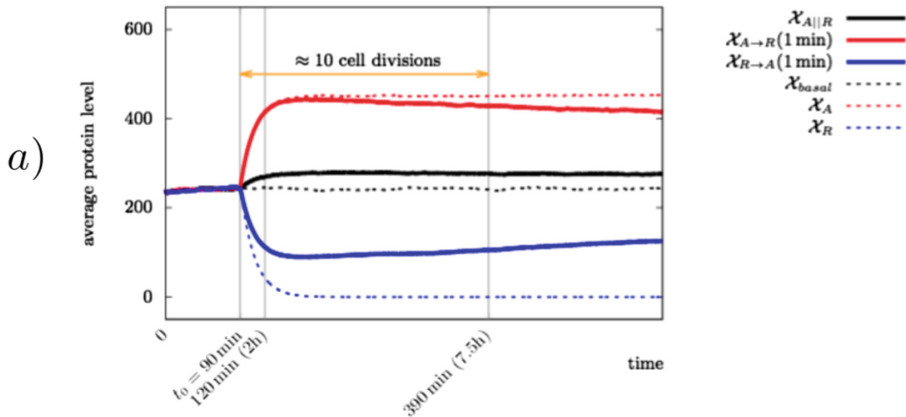
Kinetic Rates. All model parameters are in realistic ranges taken from the literature, given in Table 1 and further explained in the Appendix. The mechanism for the activator is inspired by the λ -phage. The mechanism for repressor is inspired by the *lac* operon. Further values that were tested to show the generality of our approach came from other well-characterized TFs such as CRP. The chosen parameter values were found in the literature, both for the scenario without looping [6, 15] and for the scenario with looping [32–34, 39, 40].

Parameter Search. We implemented a grid search of the viable parameter space (for different levels of eleven kinetic rate parameters, and the amounts of activator, repressor and RNAP), where we compute the average protein expression, amplitude and half-life for a subset of all parameter combinations. In our implementation, the user specifies a range for each parameter, and the models are executed, figures drawn for each possible parameter combination.

5 Results

In further text, by *phenotype*, we mean the average protein expression in a population of 1000 cells. All three scenarios $\mathcal{X}_{A||R}$, $\mathcal{X}_{A \rightarrow R}(\Delta)$ and $\mathcal{X}_{R \rightarrow A}(\Delta)$ have the

Long lived transients (reference parameter set)



50 single cell traces and the respective average

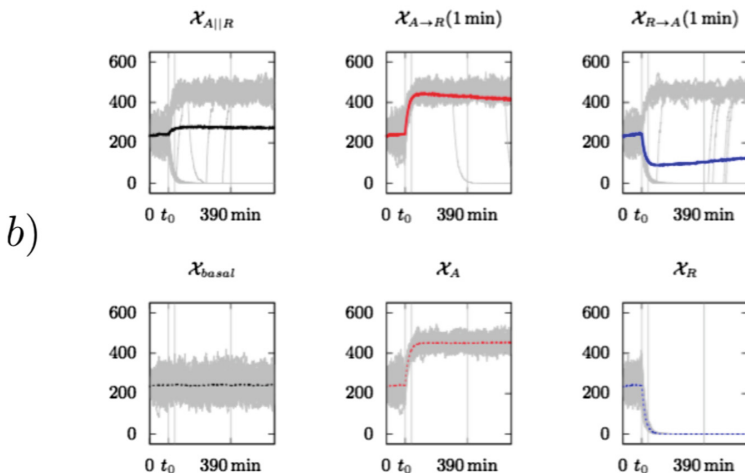


Fig. 1. A small delay in arrival times of TFs can give rise to qualitatively opposite, stable transient phenotypes for a long period of time (a) Average protein level for a population of 1000 cells, in three input scenarios (full lines) and three reference scenarios (dotted lines). (b) 50 single cell traces (grey lines) and the respective average, for each of the six modes.

same phenotypes eventually. As a reference, we also analyse the scenarios where no TFs are input (\mathcal{X}_{basal}), only activators (\mathcal{X}_A) or only repressors are input (\mathcal{X}_R). We investigated the phenotype in the three scenarios for a large range of parameter combinations. (there are $2^{11} > 2000$ combinations when only two values for each parameter are set). We choose one parameter set as the *reference parameter set* (shown in Fig. 1a), where the phenotypes are symmetric with

respect to the $\mathcal{X}_{A||R}$ scenarios in the sense that the protein expression deviates in the same amount from the phenotype of $\mathcal{X}_{A||R}$ and the rate of reaching the phenotype of $\mathcal{X}_{R||A}$ is of the same scale.

A Small Delay in Arrival Times of TFs Can Give Rise to Qualitatively Opposite, Stable Transient Phenotypes for a Long Period of Time. In Fig. 1a, we plot the observable phenotype – the mean of protein expression for a given population of cells – in the three regimes of interest (full lines) and the three reference regimes (dotted lines). Three distinct transient phenotypes are observed:

- for $\mathcal{X}_{A||R}$, the expression is close to the level of basal expression,
- the transient regime for the input $\mathcal{X}_{A\rightarrow R}(1\text{ min})$ shows high expression for multiple average cell doubling times, while
- the transient regime for the input $\mathcal{X}_{R\rightarrow A}(1\text{ min})$ shows low expression for multiple average cell doubling times.

The transient for the input $\mathcal{X}_{A||R}$ lasts roughly for one average cell doubling time (30 min), while both phenotypes for $\mathcal{X}_{A\rightarrow R}(1\text{ min})$ and $\mathcal{X}_{R\rightarrow A}(1\text{ min})$ last well over 10 average doubling times. Therefore, the delay in arrival times of TFs can result in long lived transient regimes with qualitatively opposite phenotypes (both differing significantly from the equilibrium phenotype), depending on which TF arrives first. Moreover, each of the phenotypes seems stable at the time-scale of multiple cell lifetimes. In other words, the small delays, hence two different histories of input, produce substantially different routes to the equilibrium, and the routes are so slow that they appear as steady state behavior at the timescale of most experiments.

In Fig. 1b, we see that individual cells exhibit ‘all-or-none’ behaviour: an individual cell either has high or low expression and the phenotype depends on whether the cell entered the active looped state or the repressed looped state. The expected time that a cell spends in one looped state is long. The protein expression for 50 randomly chosen single cells is displayed in Fig. 1b for each of the three regimes. In regime $\mathcal{X}_{A||R}$, an individual cell either has high expression at around 400 proteins or low expression, being fully inhibited. The noise around the low expression value is not observable in the plot, because the low expression is fully inhibited most of the time. If the DNA unloops and subsequently loops towards a different regulatory state, eg from looped repressed to looped active state (or vice-versa), the protein expression will change from low to high expression (or vice-versa). In the taken time window (10 h), three (out of 50) displayed traces switching from the high to low expression level and one trace switching from low to high expression level. As expected, the average expression in a given population (thick line in respective color) follows a continuous line; It is saturated at around 270 protein molecules. In regime $\mathcal{X}_{A\rightarrow R}$, all of the displayed 50 cells enter the active looped state before the repressors are input, but, due to the slow unlooping, the high expression profile is long-preserved, resulting in slow switching towards the low-expression state, and hence long transient time towards the average expression. In regime $\mathcal{X}_{R\rightarrow A}$, even though repressors

are input first, some cells are activated, but most of the cells are repressed. Similarly as in the profile $\mathcal{X}_{A \rightarrow R}$, since DNA unlooping from the repressed state is slow, the transient of the average protein expression is also slow. The reference scenarios - \mathcal{X}_{basal} , \mathcal{X}_A and \mathcal{X}_R show the expected behaviour.

Long Lived Transients are Robust to Changes in Kinetic Parameters.

Are long lived transients a consequence of system regulatory architecture or a careful tuning of kinetic parameters? To tackle this question, we chose six different parameter combinations, listed in Table 1, and we reproduce the plot shown in Fig. 1a for each of the parameter combinations, each for three time delays - 1 min, 5 min and 15 min. \mathbf{p}_1 is the reference parameter set (the plot shown in Fig. 1). Results, shown in Fig. 8, confirm that the long lived transients are preserved with the chosen parameter changes. Higher unlooping rate for either activators or repressors results in shortening the transient and moving the average expression level to lower and higher value respectively (\mathbf{p}_2 and \mathbf{p}_3). While decreasing the number of activators does not change the phenotype much (\mathbf{p}_4), decreasing the number of repressors results in complete dominance of activation effect when both TFs are input simultaneously (\mathbf{p}_5). Still, the delay of activator input shows full repression profile for a long period of time. When RNAP rates are scaled so that the binding and unbinding rates are both ten times slower, the duration of transients shortens and the three input regimes show the same output after much shorter time (≈ 10 h, \mathbf{p}_6).

To quantify the effect of long-lived transients, in Fig. 12 (up, model with looping), we see that the reference parameter set (\mathbf{p}_1) has a halflife longer than 20 hours no matter if the delay occurs in favour of the activator or repressor. The halflife decreases significantly in cases when the unlooping rate is decreased (one at a time - \mathbf{p}_2 and \mathbf{p}_3), or when RNAP binding and unbinding rate is scaled down (\mathbf{p}_6), while the change in the number of activators/repressors reflect more on the amplitude than on the halflife (\mathbf{p}_4 , \mathbf{p}_5).

Long Lived Transients Are Not Observed in the Model Without Looping. We next inquire how changes in regulatory architecture affect the behaviour, i.e. is DNA looping essential for observing the long lived transients? We repeated the experiments on a model without looping. Phenotypes for six parameter combinations, listed in Table 1³, are each plotted for three time delays (Fig. 10). For all parameter combinations, the amplitude and duration of transient regimes is clearly correlated with the duration of delay - the longer delays induce longer transient regimes. The transient phase is significantly shorter than in the model with looping (notice the different time-scale than in Figs. 1 and 8), but they still can last for several cell doubling times (for delays of 15 min up to 2.5 h or 5 average doubling times). However, they are not long lived transients, as the shape of transients clearly reveals that the steady-state regime is going to be reached later on, that is, the transients in this model would not be easily confused with the steady state. The observations above are indicating that looping is essential

³ Notice that these six parameter combinations are different than those used for the model with looping.

for creating the effect of long lived transients. \mathbf{p}_1 is the reference parameter set, which we choose so that the level of expression when both activators and repressors are input is close to basal (the TFs neutralise each-other’s effect overall). As expected, decreasing the recruitment by the activator results in lower stationary expression (\mathbf{p}_2), increasing the number of repressors results in stronger repression (\mathbf{p}_3), weakening the repressor binding results in higher expression (\mathbf{p}_4), weak binding of repressor in combination is not affected by decreasing the recruitment by activator simultaneously (\mathbf{p}_5) and weak binding of repressor in combination with more repressor molecules results in low expression (\mathbf{p}_6).

In the summary of characteristics of long lived transients for the model without looping (shown in Fig. 12 down), we see that, for a delay of 5 min, all parameter combinations achieve the amplitude at comparable scale as that in the case of model with looping. However, the maximal half-life in all tested parameter points is 15 min, a 100-fold difference with respect to the half-life of long lived transients in the model with looping, confirming that adding the looped configurations was essential for the effect of long-lived transients.

Phenotypes in Long Lived Transients can be Modulated by the Delay Between Inputs. We now comment on the dependency on the delay. In Fig. 11, the phenotypes in scenario $\mathcal{X}_{A \rightarrow R}(\Delta)$ are observably equivalent for all chosen values of delay. In particular, they transiently reach the same protein expression value as the scenario \mathcal{X}_A where only activator is present. Therefore, this scenario seems to be independent of delay timing between TFs as long as the delay occurs in favour of the activator.

On the other hand, the difference between phenotypes in scenario $\mathcal{X}_{R \rightarrow A}(\Delta)$ is different for delay $\Delta = 1$ min than for delays $\Delta \in \{5 \text{ min}, 15 \text{ min}\}$. While for all three delays, the effect of long lived transients can be observed (the slope of approaching the limit value is small), the phenotype (protein expression around which the transients seem to stabilise) is different. It appears that, unlike delays longer than 5 min, the delay of 1 min is not long enough for the population to repress protein expression to a value as low as in the scenario \mathcal{X}_R (where only repressor is present). In other words, the lowest gene expression value for delay of 1 min is never as small as in the scenario \mathcal{X}_R . To investigate the dependency of the transient phenotype on the delay, we simulated the scenario for several delay values between 1 sec and 5 min, namely $\Delta \in \{1\text{s}, 20\text{s}, 40\text{s}, 60\text{s}, 120\text{s}, 180\text{s}\}$ and we computed the amplitude for the scenario $\mathcal{X}_{R \rightarrow A}$. The plot in Fig. 11 demonstrates that the amplitude approaches the value of \mathcal{X}_R scenario exponentially fast with increasing delay time.

Plotting the phenotypes for scenario $\mathcal{X}_{A \rightarrow R}(\Delta)$ for delays between 1 s and 1 min shows that the same activated gene expression levels are observed even for delays as small as 1 s (plots not shown). The explanation for different sensitivity of transient phenotypes to the delay in scenarios $\mathcal{X}_{A \rightarrow R}(\Delta)$ and $\mathcal{X}_{R \rightarrow A}(\Delta)$ are the different mechanisms implementing the activation and repression. When activator is input first, it quickly binds both operator sites and the probability of being in the looped active state almost instantaneously increases to the maximum value (as fast as within 1 s), and then starts decreasing only very

slowly towards the equilibrium as soon as the repressor is present as well. On the other hand, when the repressor is input first, it does not bind both operator sites as quickly, because it is competing with the abundant RNAP, even while the activator is not in the system. Only if there is enough time for the repressor to reach the looped repressed state with a probability nearly as high as in case of repressors only, the maximally repressed expression level will be observable. Otherwise, as soon as the activator is in the system, the probability of being in the looped repressed state starts shifting slowly towards the equilibrium point, and, consequently, the protein expression in the population starts increasing.

Long Lived Transients in Protein Expression Follow the Long Lived Transients (Mixing Times) in Promoter Activity. In Fig. 9, we plot the probability of the active regulatory configuration of the promoter for six different parameter combinations (listed in Table 1). Plots show the expected agreement with those in Fig. 8.

6 Discussion

Given that the processes of transcription and translation of a gene into a protein introduce intrinsic time delays and that there is no global synchronization among the arrival times of different molecular species at molecular targets, the simultaneous arrival of TFs in genetic circuits is a strong assumption. We subjected this assumption to a perturbation analysis, where the perturbed parameters are the relative arrival times of the TFs (different to the usual choices of perturbation parameters being the kinetic rates). We simulated a simple and realistic genetic circuit with two inputs and we showed that, in presence of small perturbations in the arrival of inputs (shorter than 1 min), the circuit can exhibit three qualitatively distinct phenotypes which are stable for as long as any typical experiment would last (longer than 20 cell doubling-times). This has wide implications.

First, while our showcase example was constructed with the goal of demonstrating that long lived transients can appear in gene regulation, there are reasons to believe that many other gene regulatory schemes also exhibit long lived transients and implement multiple phenotypes by modulating the timing of inputs. To see this, consider that the number of potential phenotypes grows factorially with the number of inputs per gene as it is determined by the number of possible input orderings, meaning that, for instance, only 5 inputs would require us to analyse $5! = 120$ different input scenarios. Moreover, our analysis indicates that long lived transients are possible in promoters with many configurations and certain states that are easy to reach but hard to exit. For instance, genomic regulation of the development of sea urchin embryo shows potential for long lived transients. The relevance of transient TF production has already been determined in this system [2,44]: multiple TFs regulate a single gene which in turn has multiple targets, and there is clear differentiation between upstream and downstream components in the network. Therefore, considering long lived transients might clear up some puzzling observations like the discrepancy of TF

interactions between endogenous promoters and minimal promoters controlled by three Endo16 regulatory modules [43].

Secondly, our proof-of-concept case study suggests that any modelling approach which assumes perfectly synchronous arrival of TFs or assumes rapid equilibrium at the promoter, may fail to explain a variety of phenotypes and raise false conclusions. To illustrate this point, think of an experimentalist who observes the system which seems equilibrated, but is a long lived transient (e.g. in our case study, a delay in favour of the activator occurs). Assuming that what she sees is an equilibrium, following the approach of statistical thermodynamics, she would proceed by estimating the free energies of binding configurations, but these estimates would be wrong, as the real equilibrium is much further away. Moreover, the obtained model would explain a single phenotype, and not the variety of quasi-stable phenotypes such as the ones we see in our showcase example. In summary, one cannot ignore the order of stimulating a cell, even when the GRN under consideration is assumed to be feedback-free. Similarly, one cannot assume what is observed towards the end of the life cycle is close to equilibrium even when the system seems relatively stable, e.g. growth at steady state in bulk, even when the stimulation was completed very early on. This opens further important questions such as how can an experimentalist who observes a stable phenotype for the chosen experiment duration, distinguish between a long lived transient and a real equilibrium? One immediate insight is the critical importance of experimentally measuring the kinetic rates as accurately as possible, and taking the timing of inputs into account.

Finally, our case study opens the possibility that GRNs are exploiting event timing to perform desired behaviours - it suggests that the cell does not compute with equilibrium dynamics - as is widely assumed in the field (with the exception of ‘well behaved’ limit cycle behaviours or pulsatile behaviour [22]), but uses the transients to react to stimuli and to memorise events. The DNA may be encoding more behaviours and thus phenotypes than an understanding based on the conventional input to output mapping suggests. In particular, as our analysis of delay timing between TFs shows (Fig. 11), a whole range of different stable gene expression levels can be encoded in the event timing of inputs. More broadly, this aspect may provide an explanation to why an organism can display so many more phenotypes, though the number of genes is limited, as the complexity of the organism increases, e.g. number of genes in bacteria and human vary by a factor of only 4!

Our primary goal was to show that a simple gene regulation without feedback, with realistic parameters, can exhibit long lived transients. We hypothesised that the promoter architecture with looping will have the desired feature, and, in order to find the feature, we performed a search over the 11-dimensional parameter space, which allowed us to display and discuss a range of parameterisations showing interesting behaviour. One of the compelling questions for future work is formalisation and computation of robustness of a given promoter architecture wrt. property of long-lived transients, as well as its sensitivity to a specific (group of) parameters. To this end, we believe that the ideas of parameter synthesis for

stochastic chemical reaction networks, extensively studied in [4, 9, 10] (where the properties of the CTMC assigned to general biochemical reaction networks are expressed in continuous signalling logic (CSL)), would be a useful starting point.

Related Work. The consequences of combined effects of time delay and intrinsic noise on gene regulation has been studied in [45]. In more recent works [23], the authors elucidate the importance of relative timing of TF activation in combinatorial gene regulation with pulsatile signals. Like Lin et al. [23], our work shows that relative timing between TFs may be used by the cell to implement responses to different environments and therefore has to be taken into consideration for modelling gene expression patterns. However, while the authors in [23] suggest that the phenotypes differ in pulsatile regulation patterns, our study reveals the existence of long lived transients. From a dynamical system point of view, the effect of long lived transients that we present here can be seen through the prism of general theoretical frameworks such as proposed in [31, 42], where the authors discuss how to detect and automatically compute the meta-stable states from only the topology and timescales of the network; It would be interesting to see how precisely these methods could be used to detect the long lived transients we showcase in this paper.

Of relevance for synthetic biology, our construction based on looping suggests a way to implement memory units, though they may be leaky, in the sense that the signal is slowly being lost. In a broader context, cellular memory refers to systems whose present phenotype is dependent on the history of input stimuli and therefore the trajectory by which it has been reached [8]. The molecular mechanisms associated with such memory effects are usually based on feedback loops (e.g. the *E. coli* lac operon), DNA methylation patterns (e.g. temperate phage, pilus synthesis, cell differentiation) or inversions catalysed by site-specific recombinases (e.g. the *Salmonella* Hin system or the *E. coli* Fim system) [8, 29]. The long lived transient behaviour observed in our simulations differs from the mentioned memory mechanisms as it is purely relying on dynamical trapping of the transcriptional state. Different to the usual references to cellular memory, the long lived transients require no stabilisation of the phenotype through strong (covalent) modification of the DNA or any kind of feedback of the output on the promoter state (which is generally considered necessary for cellular memory).

The nature of the observed long lived transient states confer an epigenetic nature to these states. Methylation of histones is widely used in eukaryotic gene regulation as a modulator of gene activity that confers memory and stability to gene expression states. However, unlike methylation that requires a sleuth of specialised proteins that expend energy in order to form covalent bonds of methyl groups to histones, the long lived transients arise simply as a dynamical property of the system.

Acknowledgements. We are very grateful to Moritz Lang, Tiago Paixao and Jakob Ruess, for their feedback during the manuscript preparation.

Appendix 1.A Parameter Values

Table 1 lists the parameter ranges used for our case study example. We next explain the choice of each of the parameters with respect to their biological context.

1.A.1 Stochastic Scaling Constant

The stochastic scaling of rates and concentrations is done with a standard scaling rate for *E. coli* cell $N = 10^9$ [26].

1.A.2 Protein Production and Degradation

The protein production is taken 0.5 molec.s^{-1} ([41], caption of Fig. 2) and the degradation rate is taken 0.001 s^{-1} (corresponding to the half-life of 12 min, consistent with [26]).

1.A.3 RNAP Rates

On rate, off rate and number of RNAP molecules are consistent with the orders of values reported in [5, 14, 35].

1.A.4 Activator

The activation mechanism is inspired by the activation of the PRM promoter in the lysogenic state by protein CI in the regulation of λ -phage: CI competes with Cro to bind to the promoter sites, and, when bound, it recruits RNAP (increases PRM activity). The mechanism with looping, explained at mechanistic detail level in [35], contains three left and three right operators, leading to 1200 different DNA binding states. We model a mechanism with two states for the activator without looping ('bound' or 'not bound') and with four binding states for the

Table 1. Parameter combinations tested in the model with looping.

Parameter set ID ^a	RNAP			A binding		R binding		unloop A	unloop R	# A	# R
	on ^b	off	ifA	on	off	on	off	u_A	u_R	x_A	x_R
Reference set (P ₁)	$10^4 N^{-1}$	$1.6 \cdot 0.01$	9×	$8.8 \cdot 10^7 N^{-1}$	0.0264	$8.8 \cdot 10^7 N^{-1}$	0.016	1000	1000	275	350
increase unloop A (P ₂)	–	–	–	–	–	–	–	100	–	–	–
increase unloop R (P ₃)	–	–	–	–	–	–	–	–	100	–	–
decrease # A (P ₄)	–	–	–	–	–	–	–	–	–	10	–
decrease # R (P ₅)	–	–	–	–	–	–	–	–	–	–	10
downscale RNAP rates (P ₆)	10×	10×	–	–	–	–	–	–	–	–	–

^aThe stochastic scaling of rates and concentrations is done with $N = 10^9$. The choice of this and other parameters is detailed in the main text of the appendix.

^bAll on-rates are given in units $\text{molec.}^{-1}\text{s}^{-1}$, off-rates in units s^{-1} . The unlooping rate is specified relative to the unbinding of the respective transcription factor - eg. it means that the unlooping rate is 1000 times smaller than the unbinding of the TF A. x_A , x_R are given in molecule numbers.

activator with looping (see Fig. 5). The on-rate, off-rate well as the number of activators is taken from [35] (page 82) When activator is bound, the recruitment of RNAP is increased by factor 10 or 50 ([34] and [27] respectively) (Table 3).

1.A.5 Repressor

The repression mechanism is inspired by the well-studied transcriptional regulation, there is a word missing after transcription of the *lac* operon, the repressor *Lacl*. We take the binding and unbinding rates for the repressor from ([39], Fig. 4).

1.A.6 Looping Rates

The stability of the looped state is incorporated in the model by scaling down the unlooping rate. We choose the scaling factors of 100 and 1000 based on the computation of the ratio of dissociation rates for the models with and without looping ([40], Table 1; parameter a in [39]). The mechanism proposed in, eg. [39] suggests that the looping increases the binding rate (due to increased local concentration of TFs), while leaving the unbinding rate unchanged. As the scaled on-rates may exceed theoretical limit for diffusion-limited reactions, in our model, we incorporate the same effect by leaving the binding rate identical, and scaling down the unlooping rate.

Table 2. Parameter combinations tested for the model without looping.

Parameter set ID	RNAP			A binding		R binding		# A	# R
	on	off	ifA	on	off	on	off	x_A	x_R
Reference set (\mathbf{p}_1)	$10^4 V^{-1}$	$1.6 \cdot 0.01$	$49 \times$	$8.8 \cdot 10^7 V^{-1}$	0.0264	$8.8 \cdot 10^7 V^{-1}$	0.016	275	10
low recruitment by A (\mathbf{p}_2)	—	—	$9 \times$	—	—	—	—	—	—
increase # R (\mathbf{p}_3)	—	—	—	—	—	—	—	—	350
weak R binding (\mathbf{p}_4)	—	—	—	—	—	—	0.19	—	—
low recruitment by A, weak R binding (\mathbf{p}_5)	—	—	$9 \times$	—	—	—	0.19	—	—
weak R binding, increase # R (\mathbf{p}_6)	—	—	—	—	—	—	0.19	—	350

Table 3. Simulation parameters: all models were run for three different delays and in six different regimes.

Input time	Total time	Time delays	Input schemes	Simulation points	# samples
5400 s	36000 s	60 s, 300 s, 900 s	$\{\mathcal{X}_{basal}, \mathcal{X}_{both}, \mathcal{X}_A, \mathcal{X}_R, \mathcal{X}_{AR}, \mathcal{X}_{RA}\}$	2000	1000

Appendix 1.B Kappa Models

```
##### MODEL 1 (no looping) #####
```

```
#### Signatures
```

```
%agent: D(a,d) # Declaration of agent representing DNA with two binding
sites: 'a' for
binding the activator and 'd' for binding the Polymerase or the repressor
%agent: RNAP(d) # Declaration of Polymerase with binding site named 'd'
%agent: A(a) # Declaration of activator A with binding site named 'a'
%agent: R(d) # Declaration of repressor R with binding site named 'd'
%agent: P() # Declaration of protein P
```

```
#### Rules
```

```
# numbers after the '!' sign denote bond identifiers
# for bimolecular reactions, the rate is scaled with the average number
of molecules in
the cell 'N' in order to convert from units 'per Mol per sec' to 'per
molecule per sec')
```

```
# POLYMERASE
```

```
RNAP(d), D(d) -> RNAP(d!1), D(d!1) @ 'on_rnap' # RNAP binds, bimolecular
reaction
RNAP(d!1), D(d!1) -> RNAP(d), D(d) @ 'off_rnap' # RNAP unbinds
```

```
# PROTEIN
```

```
RNAP(d!1), D(d!1) -> RNAP(d!1), D(d!1), P() @ 'p_on' # P is expressed
when RNAP is bound
P() -> @ 'p_off'
```

```
# ACTIVATION
```

```
A(a), D(a) <-> A(a!1), D(a!1) @ 'on_a', 'off_a' # A binds to D
A(a!1), D(a!1,d), RNAP(d) -> A(a!1), D(a!1,d!2), RNAP(d!2) @
'on_rnap_if_a' # A recruits
RNAP, that is, if A is bound, RNAP binds with larger affinity
```

```
# INHIBITION
```

```
R(d), D(d) <-> R(d!1), D(d!1) @ 'on_b', 'off_b' # repressor binds the
'd' site of D;
Since 'd' is also the site for binding RNAP, when the repressor binds to
site 'd', it
prevents the RNAP from binding and hence inhibits the protein expression.
```

```

#### Variables
## rates

%var: 'N' 10^9 # the average number of molecules in the cell
%var: 'on_rnap' 10^4/'N' # division because the reaction is bimolecular
%var: 'off_rnap' 1.6*0.01
%var: 'on_rnap_if_a' 49*'on_rnap'
%var: 'on_b' 8.8*10^7/'N'
%var: 'off_b' 0.016
%var: 'on_a' 8.8*10^7/'N'
%var: 'off_a' 0.0264
%var: 'p_on' 0.5
%var: 'p_off' 0.001
%var: 'rnap0' 1500
%var: 'a_add' 275
%var: 'b_add' 10

%var: 'p0' 240 # initial number of proteins
%var: 'b0' 0 # initial number of B molecules
%var: 'a0' 0

```

Observables

```

%obs: 'protein' P()
%obs: 'd_active' D(d!1),RNAP(d!1)

```

Perturbation

```

%mod: [T]= 5400 do $ADD 'a_add' A(a)
%mod: [T]= 5400 do $ADD 'b_add' R(d)

```

Initial conditions

```

%init: 1 D(d,a)
%init: 'rnap0' RNAP(d)
%init: 'b0' R(d)
%init: 'a0' A(a)
%init: 'p0' P()

```

MODEL 2 (with looping)

Signatures

```

%agent: D(a1,a2,d,b2,loop~0~1) # Declaration of agent representing DNA
with four binding
sites: 'a1' and 'a2' for binding the activators, 'd' and 'b2' for
binding the repressor
(both) or Polymerase (site 'd'), and site 'loop' which indicates whether
the DNA is looped

```

```

or not.
%agent: RNAP(d) # Declaration of Polymerase with binding site named 'd'
%agent: A(a) # Declaration of actovatrpr A with binding site named 'a'
%agent: R(d) # Declaration of repressor B with binding site named 'd'
%agent: P() # Declaration of protein P

#### Rules
# numbers after the '!' sign denote bond identifiers
# for bimolecular reactions, the rate is scaled with the average number
  of molecules in
the cell 'N' in order to convert from units 'per Mol per sec' to 'per
  molecule per sec')

# POLYMERASE
RNAP(d), D(d) -> RNAP(d!1), D(d!1) @ 'on_rnap' # RNAP binds, bimolecular
  reaction
RNAP(d!1), D(d!1) -> RNAP(d), D(d) @ 'off_rnap' # RNAP unbinds

# PROTEIN
RNAP(d!1), D(d!1) -> RNAP(d!1), D(d!1), P() @ 'p_on' # P is expressed
  when RNAP is bound
P() -> @ 'p_off'

# ACITIVATION

#A binds to the site 'a1' or site 'a2' of DNA whenever it is not looped
A(a), D(a1,a2,loop~0) <-> A(a!1), D(a1!1, a2, loop~0) @ 'on_a', 'off_a'
A(a), D(a1,a2,loop~0) <-> A(a!1), D(a1, a2!1, loop~0) @ 'on_a', 'off_a'

# lopping is immediate when the second activator binds
A(a!1), D(a1!1,a2,loop~0), A(a) <-> A(a!1), D(a1!1,a2!2,loop~1), A(a!2)
  @ 'loop_a', 'unloop_a'
A(a!1), D(a1,a2!1,loop~0), A(a) <-> A(a!1), D(a1!2,a2!1,loop~1), A(a!2)
  @ 'loop_a', 'unloop_a'

# if A is bound to site 'a1', it recruits RNAP
A(a!1), D(a1!1,d), RNAP(d) -> A(a!1), D(a1!1,d!2), RNAP(d!2) @
  'on_rnap_if_a'

# INHIBITION
# R binds to site 'd' or site 'b2' of DNA whenever it is not looped
# By binding to site 'd', repressor inhibits the binding of RNAP to the
  same site, and
hence inhibits the expression of the protein indirectly
R(d), D(d,b2,loop~0) <-> R(d!1), D(d!1,b2,loop~0) @ 'on_b', 'off_b'

```

```

R(d), D(d,b2,loop~0) <-> R(d!1), D(d,b2!1,loop~0) @ 'on_b', 'off_b'

# looping is immediate when the second repressor binds
R(d!1), D(d!1,b2,loop~0), R(d) <-> R(d!1), D(d!1,b2!2,loop~1), R(d!2) @
    'loop_b', 'unloop_b'
R(d!1), D(d,b2!1,loop~0), R(d) <-> R(d!1), D(d!2,b2!1,loop~1), R(d!2) @
    'loop_b', 'unloop_b'

#### Variables
## rates
%var: 'N' 10^9 # the average number of molecules in the cell
%var: 'on_rnap' 10^4/'N'
%var: 'off_rnap' 1.6*0.01
%var: 'on_rnap_if_a' 9*'on_rnap'
%var: 'on_b' 8.8*10^7/'N'
%var: 'off_b' 0.19
%var: 'on_a' 8.8*10^7/'N'
%var: 'off_a' 0.0264
%var: 'p_on' 0.5
%var: 'p_off' 0.001
%var: 'rnap0' 1500 # initial number of RNAP molecules
%var: 'a_add' 275
%var: 'b_add' 350
%var: 'unloop_a' 'off_a'/1000
%var: 'unloop_b' 'off_b'/1000

%var: 'p0' 240 # initial number of proteins
%var: 'b0' 0 # initial number of B molecules
%var: 'a0' 0

%var: 'loop_a' 'on_a'
%var: 'loop_b' 'on_b'

#### Observables
%obs: 'protein' P()
%obs: 'd_active' D(d!1),RNAP(d!1)
##### Perturbation

%mod: [T]= 5400 do $ADD 'a_add' A(a)
%mod: [T]= 5400 do $ADD 'b_add' R(d)

#### Initial conditions
%init: 1 D(d,a1,a2,b2,loop~0)
%init: 'rnap0' RNAP(d)
%init: 'b0' R(d)
%init: 'a0' A(a)
%init: 'p0' P()

```

Appendix 1.C Supplementary Theory and Proofs

1.C.1 Deterministic Limit

In the continuous, deterministic model of a chemical reaction network, the state $\mathbf{z}(t) = (z_1, \dots, z_n)(t) \in \mathbb{R}^n$ is represented by listing the concentrations of each species. The dynamics is given by a set of differential equations in form

$$\frac{d}{dt} z_i = \nu_{ij} \sum_{j=1}^r k_j \prod_{i=1}^n z_i(t)^{a_{ij}}, \quad (4)$$

where k_j is a deterministic rate constant, computed from the stochastic one and the volume N according to $k_j := c_j N^{|\mathbf{a}_j| - 1}$ ($|\mathbf{x}|$ denotes the 1-norm of the vector \mathbf{x}). The deterministic model is a limit of the stochastic model when all species in a reaction network are highly abundant [19]: by scaling the species multiplicities with the volume: $Z_i(t) = X_i(t)/N$, adjusting the propensities accordingly, in the limit of infinite volume $N \rightarrow \infty$, the scaled process $\mathbf{Z}(t)$ follows an ordinary differential Eq. (4).

1.C.2 Expected Output in the Transient

The CME implies that the expectation of the marginal distribution of $\{X_t\}$ satisfies the equations

$$\frac{d}{dt} \mathbb{E}(\mathbf{X}_t) = \sum_{j=1}^r \boldsymbol{\nu}_j \mathbb{E}(\lambda_j(\mathbf{X}_t)). \quad (5)$$

To check (5), observe a transition from \mathbf{x} to $\mathbf{x} + \boldsymbol{\nu}_j$. The term $\lambda_j(\mathbf{x})\mathbb{P}(\mathbf{X}_t = \mathbf{x})$ appears exactly once when summing up for the state $\hat{\mathbf{x}} = \mathbf{x}$ as the outflow probability, and exactly once when summing up for the state $\hat{\mathbf{x}} = \mathbf{x} + \boldsymbol{\nu}_j$, as the inflow probability. This gives the term $(\mathbf{x} + \boldsymbol{\nu}_j) - \mathbf{x} = \boldsymbol{\nu}_j \cdot \lambda_j(\mathbf{x})p^{(t)}(\mathbf{x})$. It is worth noting that, upon scaling the rate constants, the equations for $\mathbb{E}(\mathbf{X}_t)$ are equivalent to (4) only if all rate functions are linear, that is, when all reactions are unimolecular.

1.C.3 Proof for Lemma 1

We first notice that the process $\mathbf{X}(t)|_{\mathbb{D}}$ is indeed Markovian, because all states of $\mathbf{X}(t)$ projected to the same state in $\mathbf{X}(t)|_{\mathbb{D}}$ are behaviourally indistinguishable (bisimulation equivalent), due to rates between lumped states not depending on protein count. From (5), it follows that

$$\frac{d}{dt} \langle x_{\mathbb{P}}(t) \rangle = -\beta \langle x_{\mathbb{P}}(t) \rangle + \sum_{j=1}^r 1 \cdot \mathbb{E}(\alpha \cdot x_{\mathbb{D}1_j}(t)) = -\beta \langle x_{\mathbb{P}}(t) \rangle + \alpha \sum_{\text{all } j} \langle x_{\mathbb{D}1_j}(t) \rangle,$$

where $\langle x_{\mathbb{D}1_j}(t) \rangle$ denotes the expected value of being in one of the active promoter configurations. The latter equals (2), since in every reachable state $\mathbf{x} \in (D_0 \uplus D_1) \times \{0, 1, \dots\}$, exactly one DNA configuration takes value 1.




Fig. 2. Transitions of the CTMC underlying basal gene expression. The state space $S \cong \{0, 1\} \times \{0, 1, 2, \dots\}$, such that $s_{1,x}$ denotes an active configuration (where the RNAP is bound to the DNA) and $x \in \mathbf{N}$ protein molecules.




Fig. 3. Each binding configuration of the DNA can be active (green, polymerase bound) or inactive (gray, polymerase not bound). Protein count can increase only when the DNA configuration is active. (Color figure online)

Appendix 1.D Supporting Figures

See Fig. 13.




Fig. 4. Searching for long lived transients in gene regulation without feedback: three modelled scenarios. We demonstrate that small delays Δ can raise qualitatively different phenotypes, which are stable for cell lifetime.




Fig. 5. Two prototype GRNs and their promoter logic: (a) Model without looping: regulatory architecture (promoter logic), (b) Model with looping: regulatory architecture (promoter logic). Mechanistic models are listed in (Appendix 1.B).




Fig. 6. Model without looping: the CTMC regulating six different DNA configurations. Thicker blue line denotes that the recruitment of RNAP is faster when the activator is bound. (Color figure online)




Fig. 7. Model with looping: CTMC regulating the DNA configurations has 23 different states. It is naturally represented as a composition of two sub-models: (left) the switching among configurations with respect to activator binding to its main and auxiliary binding sites (O_{Am} and O_{Ax} respectively, and (right) the switching among configurations with respect to repressor binding to its main and auxiliary binding sites (O_{Am} and O_{Ax} respectively). The unlooping rates (thicker blue lines) are typically much weaker than the TF unbinding. Any combination of the states in the two sub-models can be observed (reachable), except the state where both repressor and activator are looped. (Color figure online)

The effect of kinetic parameters on the shape and duration of transients for a model with looping

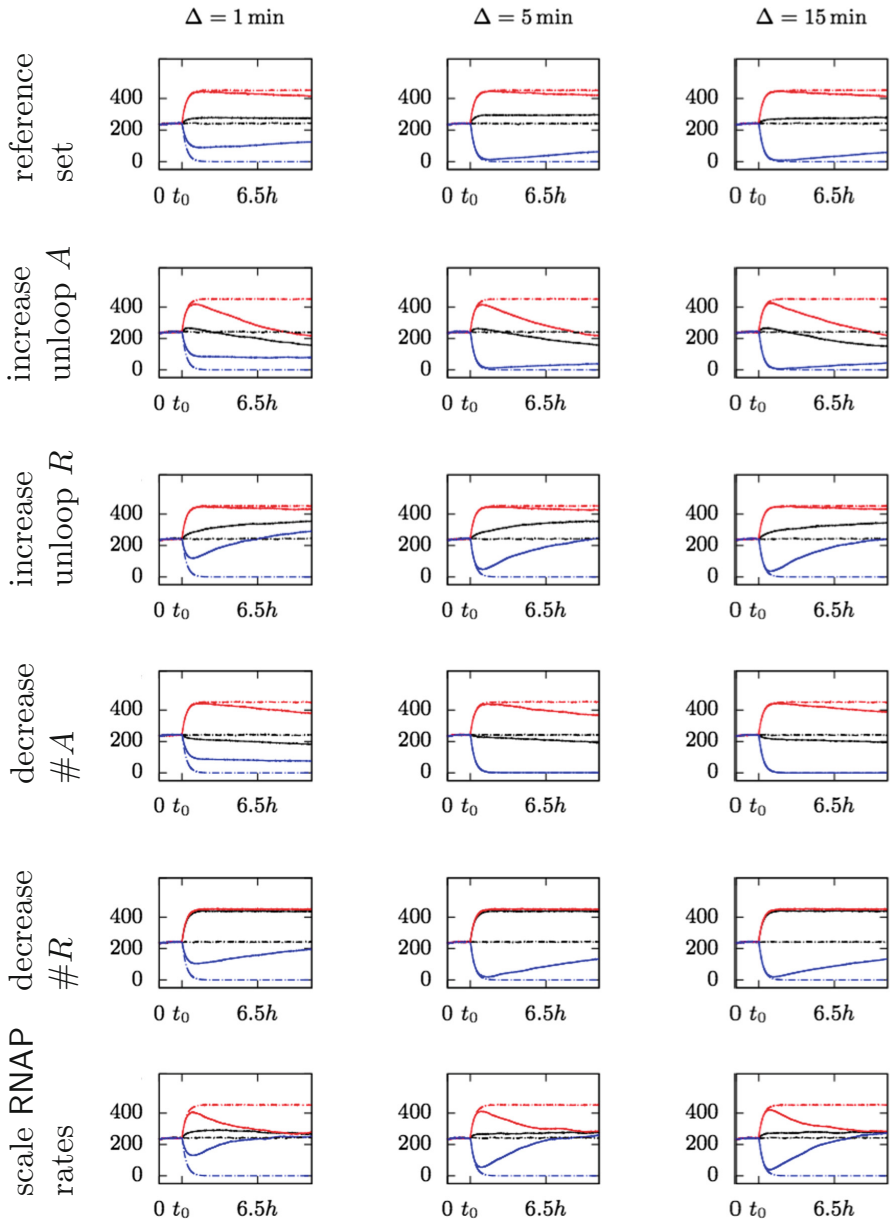


Fig. 8. The effect of kinetic parameters on the shape and duration of transients for a model with looping (for six parameter values listed in Table 1 and time delays of 1 min, 5 min and 15 min respectively).

Long lived transients in protein expression are due to long lived transients in promoter activity

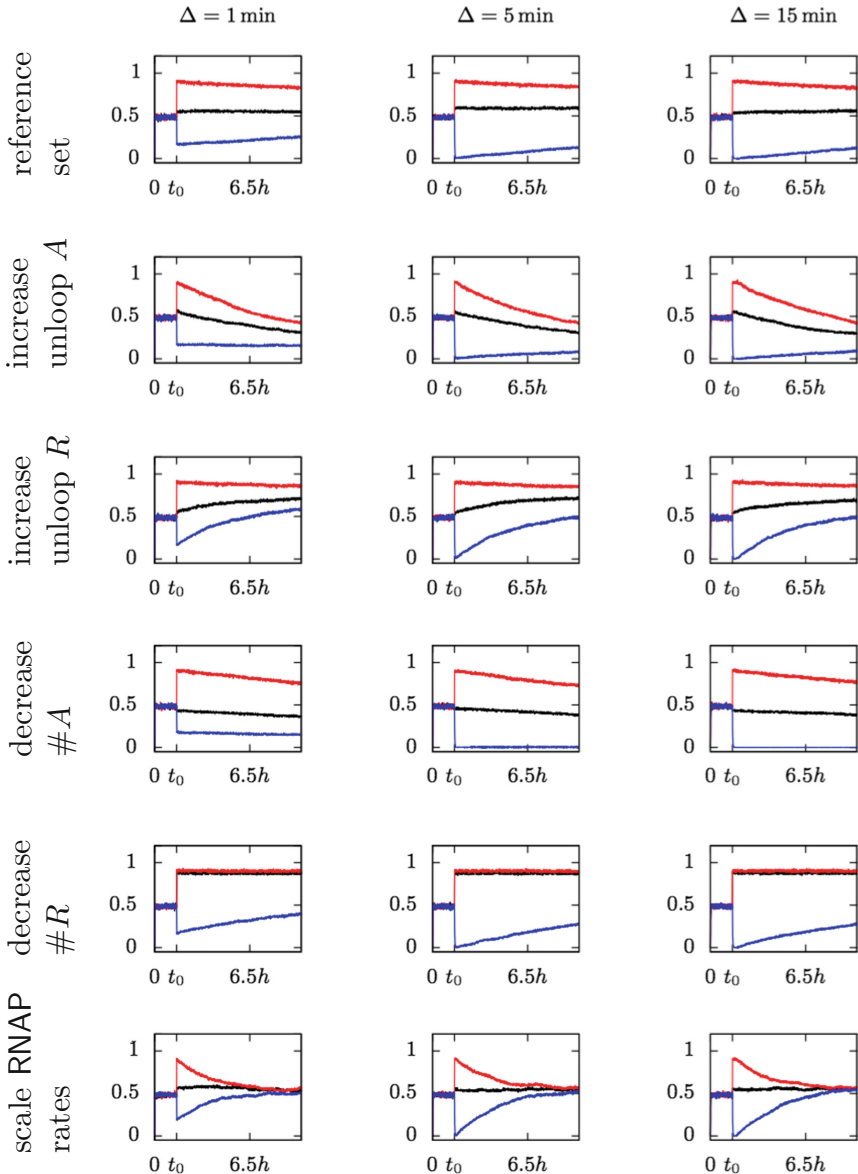


Fig. 9. Long lived transients in protein expression follow the long lived transients (mixing times) in promoter activity. We plot the statistically inferred probability of promoter logic being in the active state for six different parameter combinations, listed in Table 1 and time delays of 1 min, 5 min and 15 min respectively.

Long lived transients in protein expression are not observed in the model without looping

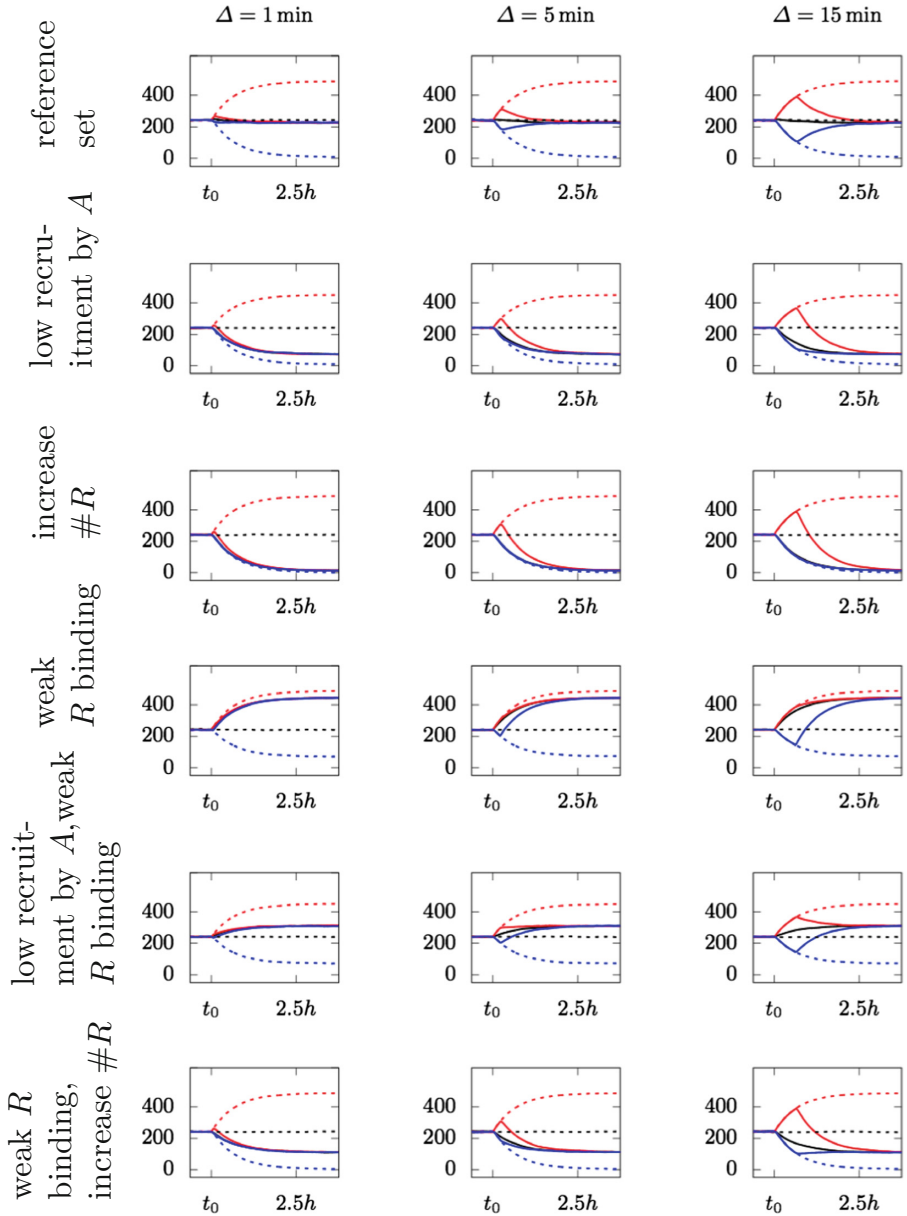


Fig. 10. Long lived transients are not observed in the model without looping (for six parameter values listed in Table 1 and time delays of 1 min, 5 min and 15 min respectively).

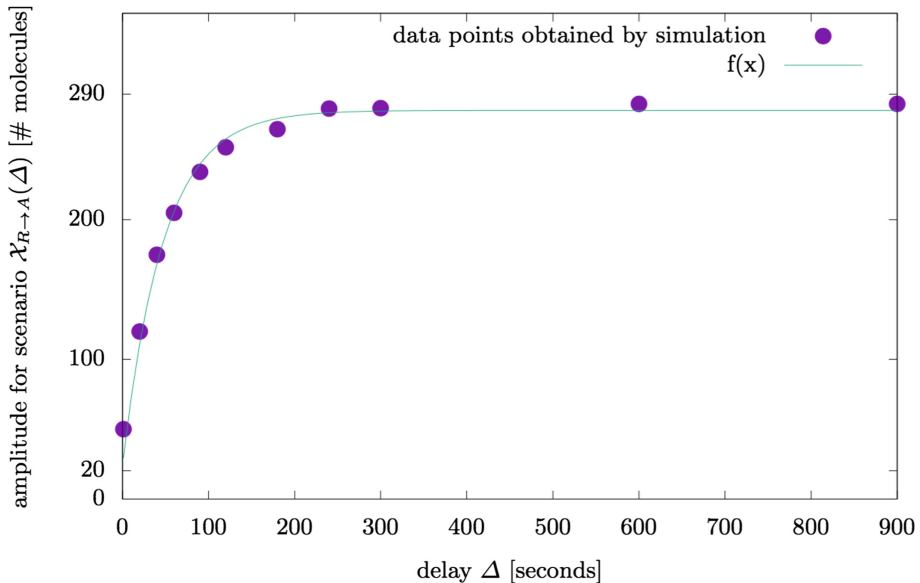


Fig. 11. In the reference parameter set, when there is no delay ($\Delta = 0$), the phenotype in scenario $\mathcal{X}_{R \rightarrow A}(\Delta)$ is equal to the one in scenario $\mathcal{X}_{A||R}$, visibly different than the phenotype \mathcal{X}_R (290 protein molecules). The difference of from the scenario $\mathcal{X}_{A||R}$ (the characteristic we formally termed *amplitude* – see Sect. 5) exponentially grows as the delay increases, that is, it quickly approaches the phenotype of scenario \mathcal{X}_R . The difference of $\mathcal{X}_{R \rightarrow A}(\Delta)$ from \mathcal{X}_R becomes observably negligible already for delays larger than $\Delta = 5 \text{ min} = 300 \text{ s}$ (difference of 10 molecules, 0.035% of the initial difference). We obtained the dependency by fitting the data obtained by simulating the system for $\Delta \in \{1, 20, 40, 60, 120, 180, 240, 300, 600, 900\}$.

Characteristics of long lived transients

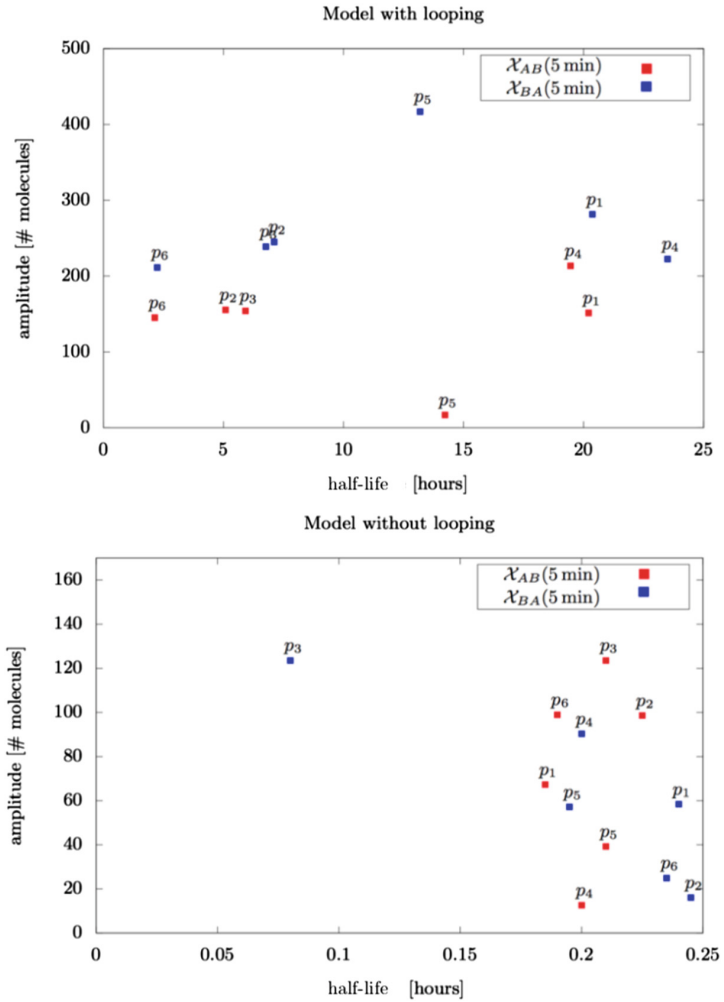


Fig. 12. For chosen parameter sets (Tables 1 and 2) and for a delay $\Delta = 5 \text{ min}$, we plot the amplitude and the half-life (defined in Sect. 2).

Characteristics of long lived transients

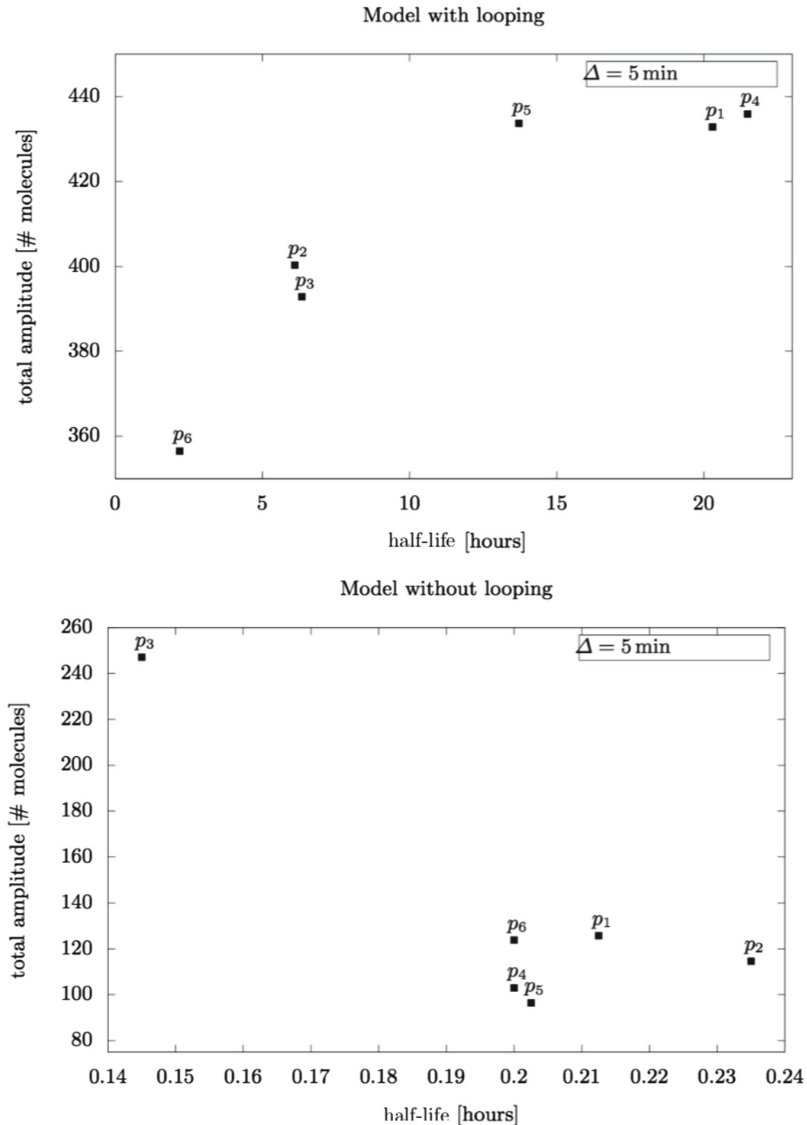


Fig. 13. We define a more global measure of the effect of delay in arrival times of TFs: first, the range of values that can be observed with a delay $\Delta = 5 \text{ min}$ by the measuring the maximum distance between phenotypes the scenarios $\mathcal{X}_{A \rightarrow R}$ and $\mathcal{X}_{R \rightarrow A}$ $\beta := \max_{t \geq t_0} |x_p(t | \mathcal{X}_{A \rightarrow R}) - x_p(t | \mathcal{X}_{R \rightarrow A})|$, and secondly, the halflife of this range $t_{1/2} := \arg \min_{t \geq t_\beta} \{t \mid |x_p(t | \mathcal{X}_{A \rightarrow R}) - x_p(t | \mathcal{X}_{R \rightarrow A})| < \frac{1}{2} \beta\}$.

References

1. Bintu, L., et al.: Transcriptional regulation by the numbers: applications. *Curr. Opin. Genet. Devel.* **15**(2), 125–135 (2005)
2. Bolouri, H., Davidson, E.H.: Modeling transcriptional regulatory networks. *Bioessays: News Rev. Mol. Cell. Devel. Biol.* **24**(12), 1118–1129 (2002)
3. Boutillier, P., et al.: The Kappa platform for rule-based modeling. *Bioinformatics* **34**(13), i583–i592 (2018)
4. Brim, L., Češka, M., Dražan, S., Šafránek, D.: Exploring parameter space of stochastic biochemical systems using quantitative model checking. In: Sharygina, N., Veith, H. (eds.) *CAV 2013. LNCS*, vol. 8044, pp. 107–123. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_7
5. Brunner, M., Bujard, H.: Promoter recognition and promoter strength in the *Escherichia coli* system. *EMBO J.* **6**(10), 3139 (1987)
6. Buchler, N.E., Gerland, U., Hwa, T.: On schemes of combinatorial transcription logic. *Proc. Nat. Acad. Sci.* **100**(9), 5136–5141 (2003)
7. Byers, R., Hansell, R., Madras, N.: Stability-like properties of population models. *Theor. Popul. Biol.* **42**(1), 10–34 (1992)
8. Casadesús, J., D’Ari, R.: Memory in bacteria and phage. *BioEssays* **24**(6), 512–518 (2002)
9. Češka, M., Dannenberg, F., Paoletti, N., Kwiatkowska, M., Brim, L.: Precise parameter synthesis for stochastic biochemical systems. *Acta Informatica* **54**(6), 589–623 (2017)
10. Češka, M., Šafránek, D., Dražan, S., Brim, L.: Robustness analysis of stochastic biochemical systems. *PLoS ONE* **9**(4), e94553 (2014)
11. Danos, V., Feret, J., Fontana, W., Harmer, R., Krivine, J.: Rule-based modelling of cellular signalling. In: Caires, L., Vasconcelos, V.T. (eds.) *CONCUR 2007. LNCS*, vol. 4703, pp. 17–41. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74407-8_3
12. Gillespie, D.: Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**, 2340–2361 (1977)
13. Guet, C.C., Elowitz, M.B., Hsing, W., Leibler, S.: Combinatorial synthesis of genetic networks. *Science* **296**(5572), 1466–1470 (2002)
14. Harada, Y., Funatsu, T., Murakami, K., Nonoyama, Y., Ishihama, A., Yanagida, T.: Single-molecule imaging of RNA polymerase-dna interactions in real time. *Biophys. J.* **76**(2), 709–715 (1999)
15. Hermesen, R., Tans, S., Ten Wolde, P.R.: Transcriptional regulation by competing transcription factor modules. *PLoS Comput. Biol.* **2**(12), e164 (2006)
16. Jacob, F.: Evolution and tinkering. *Science* **196**(4295), 1161–1166 (1977)
17. Jacob, F., Monod, J.: Genetic regulatory mechanisms in the synthesis of proteins. *J. Mol. Biol.* **3**(3), 318–356 (1961)
18. Kurtz, T.G.: Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.* **7**(1), 49–58 (1970)
19. Kurtz, T.G.: Limit theorems for sequences of jump Markov processes approximating ordinary differential processes. *J. Appl. Probab.* **8**(2), 344–356 (1971)
20. Kwok, R.: Five hard truths for synthetic biology. *Nature* **463**(7279), 288–290 (2010)
21. Levin, D.A., Peres, Y., Wilmer, E.L.: *Markov Chains and Mixing Times*. American Mathematical Society (2009)

22. Levine, J.H., Lin, Y., Elowitz, M.B.: Functional roles of pulsing in genetic circuits. *Science* **342**(6163), 1193–1200 (2013)
23. Lin, Y., Sohn, C.H., Dalal, C.K., Cai, L., Elowitz, M.B.: Combinatorial gene regulation by modulation of relative pulse timing. *Nature* **527**(7576), 54–58 (2015)
24. Marchisio, M.A., Stelling, J.: Automatic design of digital synthetic gene circuits. *PLoS Comput. Biol.* **7**(2), e1001083 (2011)
25. McAdams, H.H., Arkin, A.: It’s a noisy business! genetic regulation at the nanomolar scale. *Trends Genet.* **15**(2), 65–69 (1999)
26. Milo, R., Jorgensen, P., Moran, U., Weber, G., Springer, M.: Bionumbers—the database of key numbers in molecular and cell biology. *Nucleic Acids Res.* **38**(suppl 1), D750–D753 (2010)
27. Müller-hill, B.: *Lac Operon*. Wiley Online Library (1996)
28. Myers, C.J.: *Engineering Genetic Circuits*. CRC Press (2009)
29. Nashun, B., Hill, P.W., Hajkova, P.: Reprogramming of cell fate: epigenetic memory and the erasure of memories past. *EMBO J.* **34**(10), 1296–1308 (2015)
30. Ptashne, M.: *A Genetic Switch: Phage Lambda Revisited*, vol. 3. Cold Spring Harbor Laboratory Press Cold Spring Harbor, New York (2004)
31. Radulescu, O., Swarup Samal, S., Naldi, A., Grigoriev, D., Weber, A.: Symbolic dynamics of biochemical pathways as finite states machines. In: Roux, O., Bourdon, J. (eds.) *CMSB 2015. LNCS*, vol. 9308, pp. 104–120. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23401-4_10
32. Saiz, L., Rubi, J.M., Vilar, J.M.G.: Inferring the in vivo looping properties of DNA. *Proc. Nat. Acad. Sci. U.S.A.* **102**(49), 17642–17645 (2005)
33. Saiz, L., Vilar, J.M.: DNA looping: the consequences and its control. *Curr. Opin. Struct. Biol.* **16**(3), 344–350 (2006). *Nucleic acids/Sequences and topology* Anna Marie Pyle and Jonathan Widom/Nick V Grishin and Sarah A Teichmann
34. Saiz, L., Vilar, J.M.: Stochastic dynamics of macromolecular-assembly networks. *Mol. Syst. Biol.* **2**(1) (2006)
35. Santillán, M., Mackey, M.C.: Why the lysogenic state of phage λ is so stable: a mathematical modeling approach. *Biophys. J.* **86**(1), 75–84 (2004)
36. Schnoerr, D., Sanguinetti, G., Grima, R.: Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *J. Phys. A: Math. Theor.* **50**(9), 093001 (2017)
37. Segal, E., Widom, J.: From dna sequence to transcriptional behaviour: a quantitative approach. *Nat. Rev. Genet.* **10**(7), 443–456 (2009)
38. Shea, M.A., Ackers, G.K.: The OR control system of bacteriophage lambda: a physical-chemical model for gene regulation. *J. Mol. Biol.* **181**(2), 211–230 (1985)
39. Vilar, J.M., Leibler, S.: DNA looping and physical constraints on transcription regulation. *J. Mol. Biol.* **331**(5), 981–989 (2003)
40. Vilar, J.M., Saiz, L.: Dna looping in gene regulation: from the assembly of macromolecular complexes to the control of transcriptional noise. *Curr. Opin. Genet. Devel.* **15**(2), 136–144 (2005)
41. Vilar, J.M., Saiz, L.: Suppression and enhancement of transcriptional noise by DNA looping. *Phys. Rev. E* **89**(6), 062703 (2014)
42. Vivek-Ananth, R., Samal, A.: Advances in the integration of transcriptional regulatory information into genome-scale metabolic models. *Biosystems* **147**, 1–10 (2016)

43. Yuh, C.H., Bolouri, H., Davidson, E.H.: Cis-regulatory logic in the *endo16* gene: switching from a specification to a differentiation mode of control. *Devel.* (Cambridge, England) **128**(5), 617–629 (2001)
44. Zeller, R.W., Griffith, J.D., Moore, J.G., Kirchhamer, C.V., Britten, R.J., Davidson, E.H.: A multimerizing transcription factor of sea urchin embryos capable of looping DNA. *Proc. Nat. Acad. Sci.* **92**(7), 2989–2993 (1995)
45. Zhu, R., Salahub, D.: Delay stochastic simulation of single-gene expression reveals a detailed relationship between protein noise and mean abundance. *FEBS Lett.* **582**(19), 2905–2910 (2008)