



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Physica A 332 (2004) 585–616

PHYSICA A

www.elsevier.com/locate/physa

Bridging genetic networks and queueing theory

Arnon Arazi^a, Eshel Ben-Jacob^{b,*}, Uri Yechiali^a

^a*Department of Statistics & Operations Research, School of Mathematical Sciences, Raymond & Beverly Sackler Faculty of Exact Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel*

^b*School of Physics and Astronomy, Raymond & Beverly Sackler Faculty of Exact Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel*

Abstract

One of the main challenges facing biology today is the understanding of the joint action of genes, proteins and RNA molecules, interwoven in intricate interdependencies commonly known as genetic networks. To this end, several mathematical approaches have been introduced to date. In addition to developing the analytical tools required for this task anew, one can utilize knowledge found in existing disciplines, specializing in the representation and analysis of systems featuring similar aspects. We suggest queueing theory as a possible source of such knowledge. This discipline, which focuses on the study of workloads forming in a variety of scenarios, offers an assortment of tools allowing for the derivation of the statistical properties of the inspected systems. We argue that a proper adaptation of modeling techniques and analytical methods used in queueing theory can contribute to the study of genetic regulatory networks. This is demonstrated by presenting a queueing-inspired model of a genetic network of arbitrary size and structure, for which the probability distribution function is derived. This model is further applied to the description of the *lac* operon regulation mechanism. In addition, we discuss the possible benefits stemming from queueing theory from the interdisciplinary dialogue with molecular biology—in particular, the incorporation of various dynamical behaviours into queueing networks. © 2003 Elsevier B.V. All rights reserved.

PACS: 87.22.B; 02.50.H; 87.10

Keywords: Genetic networks; Queueing theory; G-Networks; Product form; Periodic behaviour

1. Introduction

With the completion of several genome sequencing projects, concerning various organisms (see for examples Refs. [1–4]), increasing attention is now turned to the

* Corresponding author. Tel.: +972-3-6407845; fax: +972-3-642-5787.
E-mail address: eshel@tamar.tau.ac.il (E. Ben-Jacob).

understanding of the interactions between genes, proteins and RNA molecules [5]. In particular, the regulation of gene expression is studied extensively. Alongside the more orthodox biological methods, mathematical tools are applied for the gaining of both qualitative and quantitative insights into regulatory systems (see Refs. [6,7] for surveys).

Although such efforts have now seen their fourth decade, no single modeling framework seems to capture all of the observed aspects of genetic regulatory networks [8]. The use of differential equations for their description (see Ref. [9] for a pioneering work), the most common approach employed so far, has been criticized for assuming that the modeled quantities change continuously, whereas in reality the inspected processes usually involve small numbers of molecules. Boolean networks [10], another widespread formalism, were argued against for viewing genes as binary elements, being either “off” or “on”; while this can be considered as an acceptable approximation for some genes, it has been shown that there are occasions where intermediate levels of expression should be regarded explicitly. Both approaches have been criticized for treating genetic networks as deterministic mechanisms, ignoring the stochastic noise intrinsic to them.

During the last decade, the use of *discrete* and *probabilistic* models for the description of genetic networks has become increasingly popular. As the analytical investigation of such schemes is quite complicated, most researchers employ simulations in their study (see, for example, Ref. [11]). Even so, some analytical results have been obtained (e.g. [12–15]).

In this paper we suggest the use of queueing theory for the study of genetic networks. Queueing theory is a branch of Operations Research, dedicated to the study of systems of queues, or workloads, forming in various scenarios; usually, these are assumed to incorporate stochastic phenomena. A well-established field, it has been successfully utilized for the modeling and analysis of several real-world problems.

Generalizing the standard concepts of queueing theory, a queue can be considered as the accumulation of customers, due to a series of random events, whose occurrence is subject to the set of rules defining the system. The customers can represent practically anything—people standing on a line in front of a clerk, airplanes waiting to land, computer jobs waiting to receive CPU time, etc. We suggest here to use this general notion for the representation of *molecules* of some sort. In this case, the queue length is simply the number of such molecules currently present in the system. Alternatively, one can view the queue length as something more abstract, such as the expression level of some gene; a customer then represents a “notch” in the scale of gene activity.

Integrating multiple, interconnected queues in such a model, leads to a representation of the interactions of several types of molecules and genes. A customer moving from one queue to another due to service completion, can represent the formation of one type of molecule due to some chemical reaction involving the other (e.g., the action of a certain enzyme).

The direct result of applying this modeling scheme is the description of the regulatory network as a discrete, stochastic, asynchronous system. As mentioned above, such a description is well in accordance with current views regarding these networks.

The tools supplied by queueing theory allow for the derivation of the long-term statistical properties of such systems, as well as their dependence on parameter values and on the configuration of queue interactions. Such insights can prove to be extremely useful to the understanding of genetic regulatory systems.

The harnessing of queueing theory to the study of genetic networks can take two forms. One may choose to adopt an existing queueing model, out of the exceedingly wide array of models studied over the years, changing its interpretation to match the biological scene of interest; the immediate benefit from such an approach is the ready validity of existing results found for the original model, in the biological context. Alternatively, an acquaintance with the analytical methods used in queueing theory may be sought; these can be later employed in the study of similar models, tailor-made for the exact biological system at hand.

Some reservations must be noted, however. Most of the results obtained in queueing theory are analytical, backed-up by rigorous mathematical proofs. While this is a major advantage, it also poses restrictions regarding the changes and adjustments that can be made in a given model, without losing the validity of the results. In addition, at least generally speaking, the requirement for mathematical strictness usually accompanying queueing theory, limits the complexity and intricacy of the scenarios that can be handled by its tools.

Finally, it is worth mentioning that some of the analytical work done in the stochastic modeling of genetic networks came close to queueing theory, in that some of the assumptions made (namely those regarding the Markovian character of the processes involved), as well as some of the analytical methods employed in these works, coincide with subsets of queueing theory. In one case (and, to our knowledge, the only case), an explicit analogy was drawn between the model studied and a model taken from queueing theory [13]; doing so, the authors were able to prove that their model converges to a steady state distribution in an exponential rate. However, we are unaware of a more thorough attempt to draw lines of similarity between the two fields.

We also argue that bridging between queueing theory and computational biology can prove to be beneficial for the former as well. Using approximation schemes of stochastic systems, it is possible to derive queueing networks roughly depicting various dynamical behaviours, such as periodical cycles or chaos. While the mathematical knowledge allowing for this possibility has existed for some time, queueing theory literature does not contain examples of such behaviours. To produce these, one can rely on existing dynamical models, and in particular, models describing the dynamics of genetic networks. We suggest that the employment of dynamical systems theory in general, and that of dynamical genetic networks in particular, can significantly enrich the study of queueing systems.

The organization of this paper is as follows. Section 2 provides an overview introducing queueing theory, with examples of some of the themes and models studied, as well as some of the results obtained. Section 3 presents the notion of genetic networks, and reviews the main approaches employed in their modeling, the emphasis being on discrete stochastic models. Section 4 discusses the use of queueing theory for the modeling of genetic networks. An example of a specific model, stemming from such an approach, is given; as a demonstration of the kind of results it produces, this model

is later applied for the description of the *lac* operon regulation mechanism. The final section regards the construction of queueing networks displaying various dynamical behaviours; an example of this is supplied.

2. Queueing theory

2.1. The basics

Queueing theory (see, for example Refs. [16,17]) is one of the subjects explored within the discipline of Operations Research. As the name implies, its main objects of interest are queues, or work loads, forming in front of servers of some sort. It has its roots early in the twentieth century, in the work of the Danish engineer A.K. Erlang, who studied traffic loads occurring in telephone systems. Since then it has developed to answer various real-world challenges, stemming from a variety of areas such as the design and management of industrial production lines, telephony systems, computer networks, motorized vehicles traffic and more, being successful in presenting new solutions. For example, the protocols implemented in the ALOHANET and ARPANET computer networks, which constitute the basis for today's Internet, were designed and analyzed using queueing theory. In order to address the complex nature of such scenarios, an extensive arsenal of mathematical tools has been devised, most of them general enough to be utilized in an assortment of different problems.

The simplest queueing scheme—that of a single service facility—is described in Fig. 1. In this scheme, customers requiring service enter the queueing system and join the queue, waiting to be served. From time to time, a customer standing in the queue is selected for service according to some predefined policy. The required service is then performed, after which the customer leaves the queueing system. Note that, given the wide variety of problems handled by queueing theory, it should be clear that the terms used—“customers”, “service” and so on—are general and metaphoric in their essence.

Usually, the arrival of customers and their service are considered to be the outcome of *stochastic* processes. Making specific assumptions regarding the distribution functions of these processes enables one to build mathematical probabilistic models for the analysis of the queueing system. These, in turn, allow for the derivation of certain characteristics, such as the average queue length, the average waiting time of a customer, the average time required to clear the system out of customers, and so on.

Queueing theory mainly focuses on the *steady state* of the inspected system. That is, it is assumed that after a sufficient time, the queueing system stabilizes, and its state (usually defined as the number of customers in it) becomes essentially independent of its initial conditions. Note that this does not mean that the queueing system reaches a fixed state; rather, it obtains a fixed (or stationary) *distribution function* describing its state—a distribution that does not change over time.

Below is considered a series of representative examples, demonstrating common models explored by queueing theory, as well as some of the results obtained; these examples will also assist us in establishing our ideas, further below. We open with two models of single queues, and continue to discuss networks of queues.




Fig. 1. A single service facility. Customers arriving from outside the system join a queue formed in front of the local server. Customers are picked for service, after which they leave the system. Both the interarrival times and the service times are usually characterized by distribution functions. L_q denotes the length of the queue. A more conventional schematic representation appears in the lower part of the figure.

2.2. The M/G/1 model

The title of this section lends us an opportunity to present a standard notation used in regard to queueing models: the $A/B/C$ notation. A supplies information about the arrival process; B describes the service times; and C designates the number of servers present in the service facility. In our example, the arrival process is Poisson (and thus Markovian, represented by the letter “ M ”); the distribution of the service times is not specifically defined (“ G ” stands for “general”); and the number of servers is 1.

Let λ denote the rate of arrivals. We assume that the service times are all identically distributed, having the same distribution as a certain random variable, B . No assumptions are made regarding this distribution; however, it is required that the averages of B and B^2 , denoted by $E(B)$ and $E(B^2)$, are finite. Let $\rho = \lambda E(B)$. This size is of extreme importance in queueing theory: it can be interpreted as the average amount of work (that is, required service) flowing into the system during each time unit, since λ is the average number of customers arriving during that period, and $E(B)$ is the average amount of work required by each customer. Assuming that the server can handle a single unit of work during each time unit, one can intuitively see that a necessary condition for the stability of the system is that $\rho < 1$. This condition will also arise from the analytical results presented below.

In order to analyze this model, one can use the fact that the arrival process is Poisson, and define a Markov chain embedded in the moments of service completion. Let X_n be the number of customers present in the system at the moment the n th customer completes his service; and let ζ_n be the number of customers arriving to the system during the service of the n th customer. It follows that

$$X_{n+1} = X_n - 1 + \zeta_{n+1} \quad \text{if } X_n > 0$$

and

$$X_{n+1} = \zeta_{n+1} \quad \text{if } X_n = 0 .$$

The reason for this is as follows: If the system was not empty after the departure of the n th customer (that is, $X_n > 0$), then the number of customers in it after the next service completion, is the number of customers present when that service began, minus 1 (the $(n + 1)$ th served customer), plus the customers arriving during that service. If, on the other hand, the n th customer left no customers behind him, then the $(n + 1)$ th customer reaches an empty system, receives service, and leaves behind him only the customers arriving during his service.

Using this recursive rule, it is possible to derive the Laplace–Stieltjes transform of X , defined as $\tilde{X}(s) = E[e^{sX}]$. Relying again on the Poisson nature of the arrival process, one can show that the distribution of X , the number of customers in the system at the moments of service completion, is the same as the distribution of L , the number of customers in the system at an *arbitrary* moment. One can then proceed to deduce the mean of L , given by the so-called Pollaczek-Khintchine formula:

$$E(L) = \frac{\lambda^2 E(B^2)}{2(1 - \rho)} + \rho . \tag{1}$$

We can see that indeed, as $\rho \rightarrow 1$, $E(L) \rightarrow \infty$, and the system loses its stability. A less intuitive result is the dependence of L on the second moment (and hence, the variance) of the service times distribution. This means that it is not enough to rely on a calculus of averages alone to describe the behaviour of a queueing system; variances must be considered as well.

There is a wealth of other results that can be found for this model. However, we will settle with what was presented, and proceed to the next example.

2.3. The M/M/1 model

Here we look at a scenario identical to the previous one, but further assume that the distribution of the service times is exponential, with a rate denoted by μ (thus, the events of service completion are generated by a Markovian process. This explains the second “M” in the notation used for this case). This additional assumption enables the derivation of more detailed results, most notably an explicit expression of the distribution function of L . However, this assumption may be difficult to justify in some cases.

Again, we set $\rho = \lambda E(B)$. Here, $E(B) = 1/\mu$, so $\rho = \lambda/\mu$.

The analysis of the $M/M/1$ system is conducted in a different manner than that of the $M/G/1$ model. The system is described using a birth-and-death process, where a birth stands for the arrival of a customer, and a death for a departure of a customer (after service completion). One can then show that the stationary probability of having k customers in the system is equal to

$$P(L = k) = (1 - \rho)\rho^k \quad (2)$$

(conditioned that $\rho < 1$).

In particular, $P(L = 0) = 1 - \rho$, or, alternatively, $P(L > 0) = \rho$. This gives another interpretation to the meaning of ρ —the proportion of time during which the server is busy serving customers.

Some of the additional results obtainable here are the average number of customers in the system, and the average time a single customer spends in the system (denoted by W):

$$E(L) = \frac{\rho}{1 - \rho} = \frac{\lambda}{\mu - \lambda}, \quad (3)$$

$$E(W) = \frac{1}{\mu - \lambda}. \quad (4)$$

Furthermore, one can show that the probability distribution function of W is exponential with a rate equal to $(\mu - \lambda)$, namely $P(W \leq t) = 1 - e^{-(\mu - \lambda)t}$.

2.4. Networks of queues and product form solutions

Queueing theory analysis is not limited to the case of a single service facility; one possible complication is the consideration of a *network* of queues. Such a network is composed of multiple service facilities, each with its own queue and server.¹ Customers may generally arrive from outside the system to any of these facilities. After a customer is served, he can either leave the system or move to a different facility in the network, joining the queue there. The possible movements between the queues²—usually described by transition probabilities, specified separately for each queue—define the structure of the network.

Queueing networks are rather difficult to explore, and in order to enable some sort of analysis to be conducted, one must usually assume that all of the distributions involved—both of the interarrival times and the service times—are exponential. Further assuming that the rates of these distributions are constant, results in a family of models

¹ Note that it is possible to derive analytical results in the case where multiple servers are present in a service facility, as well. However, we will restrict our discussion here to the case of a single server in each facility.

² That is, movements between service facilities. In what follows, the term “queue” will be used to denote both the service facility, and the queue in it—in accordance with the conventional terminology used in relation to queueing networks.




Fig. 2. Two queues in a Jackson network. The figure displays the possible transitions, together with their rates and probabilities. Each queue is denoted by a circle and each transition is denoted by an arrow. The service rates are written inside the circles; the arrival rates, in front of arrows; probabilities, beside arrows.

named *Jackson networks*. In these networks, it is possible to derive an explicit expression of the stationary *joint* distribution function of the queue lengths—that is, one can precisely figure the probability that in the first queue there will be k_1 customers, in the second queue k_2 customers, etc. Furthermore, this distribution function has a particular form, called *product form*: it can be written as the product of n components, n being the number of queues in the network, where each component complies with the marginal distribution function of a single queue. This allows not only for the exploration of the joint behaviour of the network, but also for the separate analysis of each single queue. We will briefly show all this now.

Let γ_i be the rate of arrivals to the i th queue, from *outside* the system; and let μ_i be the service rate at facility i . When a customer completes his service there, he moves to queue j with probability p_{ij} , or leaves the system with probability $d_i = 1 - \sum_{j=1}^n p_{ij}$. This scheme is presented in Fig. 2.

We further define the *aggregate* input rates in each queue—that is, the rate of arrivals to the queue, both from the outside and from other queues in the network. The aggregate arrival rate to queue i is denoted by λ_i . It follows that:

$$\lambda_i = \gamma_i + \sum_{j=1}^n \lambda_j p_{ji} = \gamma_i + \sum_{j=1}^n \rho_j \mu_j p_{ji}, \quad i = 1, \dots, n, \tag{5}$$

where $\rho_i = \lambda_i / \mu_i$. These n linear equations are called the *traffic equations* of the network. Note that as in the $M/M/1$ and $M/G/1$ models, ρ_j represents the proportion of time during which the server in queue j is busy. Thus, $\rho_j \mu_j$ is the *effective* service rate at queue j , and $\rho_j \mu_j p_{ji}$ is the effective passage rate from queue j to queue i .

The so-called Jackson Theorem states that the stationary joint probability distribution of the system is given by

$$P(L_1 = k_1, \dots, L_n = k_n) = \prod_{i=1}^n (1 - \rho_i) \rho_i^{k_i}. \tag{6}$$

The condition for stability is $\rho_i < 1, \forall i$.

Indeed, one can see that the joint probability distribution is presented as the product of n components, each depending on the parameters of a single queue only (once the aggregate rates have been computed, that is). This allows for the conclusion of the

marginal probability distribution function of each queue, as well as its average length:

$$P(L_i = k_i) = (1 - \rho_i)\rho_i^{k_i}, \quad (7)$$

$$E(L_i) = \frac{\rho_i}{1 - \rho_i}. \quad (8)$$

Comparing these results to those obtained in the $M/M/1$ model, it is possible to see that once the aggregate rates have been computed (taking into account the interdependencies of queues in the network), each queue can be statistically treated as if it were a separate $M/M/1$ system, standing on its own.

2.5. G -Networks

G -networks (the “ G ” stands for “generalized”), first presented by Gelenbe in the early 1990s [18], are an extension of Jackson networks. Here, two types of elements move around in the network: *Customers*, which are identical to the customers in regular queueing networks; and *signals*, which induce some effect on the queues of the network, and then leave the system. The exact nature of this effect differs in the various types of G -networks: signals can cause a single customer to leave the system (in which case they are sometimes called “negative customers”); they can trigger the movement of a customer from one queue to another; or they can cause the deletion of a random amount of work from a queue—these are but a few of the signals types studied so far (for surveys of papers exploring G -networks, see Refs. [19,20]). Signals can appear from outside the network; or they can be a result of a “metamorphosis” occurring to a regular customer.

A common result recurring in the study of G -networks, is that of product form stationary solutions—as with Jackson networks. However, in G -networks, the traffic equations, used here for the description of the aggregate rates of customer arrivals as well as those of signal arrivals, are generally *nonlinear*; this poses questions regarding the existence of such solutions (note that uniqueness is derived automatically from existence, since we are dealing here with the normalized stationary solution of a system of Chapman-Kolmogorov equations—see also Ref. [19]).

Most of the applications of G -networks explored to date refer to computer networks, where cancellation of work, as well as the movement of work from one server to another, is indeed possible—for example, as part of load balancing schemes. Another application—the one that initially motivated the use of G -networks—is the modeling of *neural networks*. In this context, a service facility represents a neuron; the length of the queue represents the activation level of the neuron; the movement of a customer from one service facility to another represents the excitation of the latter neuron by the former; and the movement of a signal, which has the effect of removing a single customer from the target queue, represents inhibition. For further information, see Refs. [19,21,22].

Thus, G -networks constitute a convenient framework for modeling rather general phenomena—that of stochastic processes involving the coupled increase and decrease of some studied amounts (which can be represented by the queue lengths), organized in some predefined structure. This framework can be applied to a wide range of cases.

However, at least to our knowledge, G -networks were not yet used for the description of biochemical or genetic networks.

2.6. Approximations of queueing networks

Finding an explicit expression of the joint distribution function of a queueing network is usually possible only in the simplest cases, e.g. when one considers the *stationary* distribution and assumes *fixed rates* for the stochastic processes involved. For more complicated scenarios, it is still possible to write an integro-differential equation describing the time evolution of the network's transition probabilities: this is the well-known *Kolmogorov forward integro-differential equation*, also called the *master equation* of the network [23]. Solving this equation supplies the time-dependent joint probability distribution function of the network; however, this is generally a difficult task.

An alternative way to write the master equation is to use the Kramers–Moyal expansion [24,25]. Here, the time derivatives of the transition probabilities are expressed using an infinite series involving the moments of the distribution function of the jump process. This does not necessarily simplify the solution of the equations, but it readily leads to an approximation scheme, namely keeping only the leading terms in the series.

Taking into account only the first term in the series results in what is known as the *fluid approximation*, which approximates the *average* behaviour of the system over time, employing a *deterministic* description. Such an approximation is usually valid only when the changes induced by the movement of a single customer are relatively negligible compared to the overall network's state. In addition, if the fluid approximation features more than one stable steady state, its validity is limited to a certain time interval.

When one keeps the first *two* terms in the Kramers–Moyal expansion, the result is a *diffusion equation*. Its solution approximates the transition probability density function of the system, so that the statistical properties of the system can be assessed.

3. Genetic networks

3.1. Biological motivation

The main functions in a living cell are carried out by proteins, which are synthesized from the information kept in its genes. Though all the cells in an organism contain the same DNA (with few exceptions), different cells express different genes and produce different proteins, therefore exhibiting different behaviours. Thus, in order to understand the cell's functioning, one cannot settle for acquiring knowledge of the DNA sequences alone, but must also become acquainted with the processes regulating protein synthesis, determining which protein will be produced and when, and at what level [6].

A key feature of these regulatory processes is the fact that *they themselves* involve the utilization of proteins. For example, proteins may act as enhancers or repressors of the expression of specific genes (possibly the genes which are responsible to their own

production), controlling the amount of mRNA transcribed from them; and the synthesis of these regulatory proteins themselves may be controlled by yet other proteins. This gives rise to an intricate *network of regulatory interactions* between DNA, RNA, proteins and small molecules, sometimes referred to as a *genetic network*. Through these networks, the organism can implement complicate logical “circuits”, which enable it to respond appropriately to various environmental conditions. A well-studied example is the relatively simple network, which enables *E. coli* to produce the metabolic enzymes required for the digestion of lactose, only in the absence of the more favourable glucose [5,26].

Since the interactions comprising a genetic network may be fairly complex, including interlocking positive and negative feedback loops, mere intuition usually does not suffice for understanding their behaviour. Rather, formal mathematical methods and computer tools are required for their modeling, analysis and simulation, as means of gaining some insight into their functioning.

More specifically, such approaches are useful for several purposes. First, there are certain regulatory circuits, extensively studied over the years, for which most elements participating in the regulatory processes, as well as the interactions between them, have been identified. For these cases, detailed models can be built and studied, shedding new light on the possible dynamics that can arise, their robustness to changes in certain conditions, etc. The results of such an analysis may pose new questions and point to areas where more experimental work is needed, which can lead to more accurate and specific models, and so forth.

A second use for modeling and analysis is the study of general structural motifs, which are recurrent in several regulatory circuits (for example, negative feedback loops). The exploration of general motifs may result in a more profound understanding of regulatory networks in general (see Ref. [27] for an example of this approach).

Third, recent years have seen the rise of such laboratory techniques as microarray technology, which allow for the simultaneous measurement of the expression levels of many genes. One possible use for such data is the reverse-engineering of the underlying genetic network, through the use of statistical inference methods (see, for example, Ref. [28]). However, these methods may result in several candidate networks. Incorporating the knowledge gained from the study of regulatory circuits dynamics may enhance such a research process, helping to disqualify unsuitable networks.

Over the past 40 years, various approaches have been suggested for studying regulatory networks, each with its own strengths and weaknesses. Most of these focus on the first point of regulation, i.e., that of gene expression (protein synthesis is regulated at each of its steps—gene expression, RNA processing and transport, RNA translation, and posttranslational modifications of proteins). Here we discuss some of these approaches, in particular those which are more closely related to the one suggested by us. We base our discussion on Refs. [6,7], which contain more elaborate surveys.

3.2. Boolean networks

Boolean networks were first suggested as a model of genetic regulatory systems by Kauffman, in 1969 [10]. In this approach, it is assumed that in order to effectively

describe the behaviour of such systems, it is sufficient to consider explicitly only the expression levels of genes, without specifying the concentrations of mRNA molecules, proteins and other participants in the modeled processes. The state of a gene is approximated by a Boolean variable—that is, a gene can be either “on” (currently expressed) or “off” (currently silenced). The regulatory control of the expression of each gene is represented by a logical function, depending on the states of other genes in the network. For example, such a function may specify that gene A will be “on” in the next time step if currently gene B is “on” and gene C is “off”. As implied in this example, the states of the genes in the network are updated in discrete time steps; moreover, in each time step the state of *every* gene in the network is updated. Thus, all network elements are assumed to change *synchronously*.

Since the number of elements in a Boolean network is finite, and since each element has a discrete number of possible states, the number of possible network states (defined as the vector of the states of the individual elements) is also finite; thus, a specific run of a network is bound to return to a state it already encountered. Also, since all state transitions are deterministic, the system will then reconstruct all its steps since the first appearance of the recurrent state. Thus, the network is said to have reached a *cycle*. If this cycle consists of a single state, it is termed a *steady state* of the system.

At the price of making somewhat radical assumptions regarding the nature of gene expression and its regulation, the Boolean networks approach allows for the study of extremely large genetic networks, due to its computational simplicity: already in the late 1960s, Kauffman was able to inspect networks containing up to 10,000 elements. Such large networks are way beyond the reach of any other existing modeling technique, even today.

The study of Boolean networks is performed using computer simulations. These allow for the identification of steady states and cycles, and of their domains of attraction (that is, states which lead to these attractors). One research direction concerns studying the implications of local properties and structural motifs, such as the average connectivity of each element in the network or the logical functions used, on the global behaviour of the network (see Refs. [29–33] for reviews).

For example, Kauffman [10] showed that the average connectivity affects the length and stability of cycles in the network.

The simplicity of Boolean networks, and the relatively low number of parameters involved in their definition, make them attractive models for the purpose of reverse-engineering the structure of a genetic network out of actual gene-expression data (cf. Ref. [34]). However, here their deterministic nature may pose a problem, due to the noise inherent in real-life gene-expression data [35,36].

Almost all of the assumptions made by the Boolean networks approach draw criticism [8]. First, though it is acceptable, in many cases, to regard gene expression as having a binary nature, as the activation profile of genes commonly has the shape of a sigmoid curve, it is not always appropriate to do so; there are cases where intermediate levels of gene expression cannot be neglected, as these have a different effect than very low or very high expression levels. Also, the description of gene expression regulation using logical functions has been pointed out to be an oversimplification. Third, the assumption that all elements change synchronously is problematic, as this is not usually the case,

and as it prevents from certain behaviours from happening [6,7]. In addition, the fact that Boolean networks lack specific modeling of the role played by mRNA molecules, proteins and other molecules in gene regulation, makes them unsuitable for describing various phenomena [7]. Finally, the assumption of determinism is challenged; this is discussed below.

3.3. The “continuous” approach

Under this title we group several works, all using differential equations to describe the time evolution of gene expression levels and the amounts or concentrations of mRNA, proteins and other molecules. All of the involved variables are non-negative and continuous. The regulatory interactions are expressed using functional and differential relations between the variables.

Such methods have been widely used for the modeling of genetic regulatory systems, starting from 1963 [9]; see Refs. [6,7] for a partial list of works. Models were built to describe either specific regulatory circuits, or for the study of general motifs and characteristics recurring in a large number of known regulatory systems, such as negative feedback loops. The interactions considered can be either linear or nonlinear (the latter more realistic and widespread). Some models take into account the spatial variation, using partial differential equations; others use time delays for the same purpose. Several types of behaviour were described using continuous models: steady states, limit cycles, chaos, bistability and multistability, and more.

Due to the nonlinearity of most models, numerical simulations are used for their analysis, rather than analytical methods. In addition, bifurcation analysis is commonly used for the investigation of the sensitivity of steady states and limit cycles to changes in parameter values.

The main advantage of this approach is the ability to accurately model the quantities described and their interactions. Also, it makes available to the researcher the extensive knowledge and methodologies of dynamical systems theory.

The main disadvantage is the high computational cost involved in running the numerical simulations required. This considerably limits the size of the systems that can be explored using this approach. In addition, building a complete model requires the accurate knowledge of the functional interactions between the described molecules, as well as the precise assessment of the involved parameter values. The absence of such data may seriously cripple the model, as the derived behaviours are usually quite sensitive to these specifications.

3.4. Probabilistic models

The continuous approach assumes that the modeled amounts vary *continuously* and *deterministically*. However, in the context of gene regulation, these assumptions may be inappropriate, for two reasons: first, the number of molecule instances participating in such processes may be quite small—few tens of molecules in some cases; and second, some of the chemical reactions involved occur at slow rates. The first point suggests that *discrete* variables may be more fitting for the representation of molecules amounts;

the second implies the existence of *stochasticity* in genetic regulatory systems. Indeed, for the latter claim, experimental evidence exists for some time (see Refs. [37,12] for a more elaborate discussion). For example, in phage λ , the choice between the lytic and lysogenic outcome is probabilistic in its essence, and was successfully modeled as such in Ref. [11]. Thus, the fate of the entire organism may lie on a single roll of a dice.

This leads us to the discrete, probabilistic modeling approach. Here, the state of the system is defined as the number of molecules of each type, and is expressed using a vector of integer, non-negative variables. The possible chemical reactions are determined, together with the probability of each reaction to occur (these probabilities may be state-dependent). From these definitions stems implicitly the joint probability distribution of the system, which specifies the probability of a certain state to occur at a certain time. The time evolution of the joint distribution is governed by the *master equation* of the system; solving it produces the explicit functional form of this distribution, thus giving a complete description of the system's stochastic behaviour over time.

Unfortunately, in most cases it is quite difficult to derive the exact form of the master equation, let alone solve it. Thus, most researchers resort to the use of simulations of the explicit molecular interactions occurring in the regulation process, disregarding the master equation altogether. Even so, some analytic work on the subject exists (see below).

Stochastic simulations use the framework suggested by Gillespie in 1977 [38]. In each step, two choices are made randomly, according to the state-dependent probabilities specified by the model: the next reaction to occur, and the time on which it will take place. The state of the system is then updated accordingly, and the simulation proceeds to the next step. This framework was later modified and improved by other authors, for example in Ref. [39].

A single run of a stochastic simulation will produce an arbitrary trajectory in the phase space. In order to derive general results concerning the stochastic behaviour of the system, several such runs must be performed. It is then possible to get an estimate of the joint probability distribution, as well as the average trajectory and the dispersal around it.

The main advantage of the stochastic simulations approach is the ability to construct accurate models for the molecular processes involved in gene expression regulation, taking into account all the details held relevant while not neglecting such effects as stochasticity. However, exaggerating in this perspective will limit the ability to make useful interpretations of the results.

From the requirement to execute numerous simulation runs, and from the detailed description of the chemical reactions involved, usually incorporated in such models, stems the main disadvantage of this approach—its high computational cost. In addition, as with the continuous approach, an exact knowledge of the reaction rates and indeed, of the reaction mechanisms themselves, is required; however, such knowledge is often unavailable.

As mentioned above, some analytical results were obtained for probabilistic models of genetic networks, mostly in recent years. In some of these works, it was

possible to derive, and even solve, the master equation itself; others use approximation schemes, based mainly on the Kramers–Moyal expansion. Following is a sample of these works.

Peccoud and Ycart [13] suggested the use of Markovian birth-and-death processes to model a single gene and the protein it produces. The gene can be either “on” or “off”, while the amount of protein molecules is described using an integer, non-negative variable. The possible transitions in this model are the activation and deactivation of the gene, the production of a protein molecule (if the gene is active), and the degradation of a protein molecule (if the number of these molecules is greater than 0). The rate of degradation is linear in the number of protein molecules; all other rates are constant. The model does not include an explicit description of the regulation of the gene (that is, it is not related to the protein molecules). The authors were able to derive the time-dependent mean and variance of the number of protein molecules in the system, as well as those in steady state.

Kepler and Elston [12] considered a similar model, but examined additionally a gene enhancing its own production, and a system composed of two mutually repressing genes. Again, all reactions are considered to be Markovian, with a linear rate describing the degradation of the protein, and fixed rates for all other interactions. The activation and repression of the gene result from an interaction of two protein molecules, so their rate is quadratic in the number of molecules. In the case of the self-enhancing gene, the authors were able to write the diffusion approximation of the master equation, and to derive from it an approximation of the stationary distribution function of the number of protein molecules. They then used this distribution function to generate a bifurcation diagram, depicting regions of qualitatively different steady state distributions, as a function of the model parameter values; in addition, they explored the first-passage times between alternative stable steady states. In the case of the two mutually repressing genes, the authors conducted a similar analysis, but instead of using a diffusion approximation, they approximated the number of protein molecules using a set of deterministic ODEs, and generated a bifurcation diagram using these. The analytical results were verified using numerical simulations.

Thattai and van Oudenaarden [14] presented a general framework for the treatment of regulatory *networks*, consisting of an arbitrary number of genes. Each gene is represented by a triplet of non-negative integers, specifying the numbers of DNA, mRNA and protein molecules (the number of DNA molecules can be used to represent the gene activation level). The state of the entire network is the collection of all these triplets. The possible stochastic transitions—all Markovian—are an increment or a decrement in the amount of a single molecule. The state-dependent rates of these transitions are assumed to be *linear combinations* of the numbers of molecules present. The authors managed to obtain the explicit form of the master equation, and from it derived a differential equation for the moment generating function of the joint distribution of the network. From the latter one can deduce, solving linear algebraic equations, the stationary mean and variance of the network state. These analytical results were further verified using stochastic simulations. As the authors note, such a framework can be used for the exploration of the stochastic behaviour around stable steady states, where the linearization of transition rates can be considered to be valid.

Gonze et al. [15] looked at a different scheme: They started with a set of deterministic differential equations, which can be used for the description of circadian oscillations. They then considered a stochastic system, where the transition rates are chosen such that the *average* behaviour of the stochastic system can be approximated by the limit cycle produced by the original deterministic equations; the stochastic system evolves around this cycle, displaying noisy oscillations. More precisely, their model consists of three variables, describing the amount of mRNA molecules, and the amounts of the protein produced from it, in its two forms—cytosolic and nuclear. The possible transitions are the increment and decrement of these amounts; some of the transitions involve simultaneous changes (that is, the increment of one number and the decrement of another). Here also, these transitions are taken to be Markovian; however, in order to reconstruct the desired oscillatory behaviour, their rates are nontrivial functions of the molecule amounts. The authors also introduced a new parameter Ω , designating the size of the stochastic network; dividing the numbers of molecules by it, they obtained the concentrations of each molecule type. These concentrations serve as the variables of the model. Note also, that the larger Ω is, the smaller the relative change in a concentration due to a stochastic transition is; thus, for large values of Ω , the system is expected to behave regularly, maintaining a close trajectory around the approximated average, while in the opposite case, significant random fluctuations are to be observed.

The authors employed the diffusion approximation of the master equation (that is, a Fokker–Planck equation). From it, they obtained several analytical results, including the probability density of the system around the limit cycle, the probability density of the first return time of one of the concentrations to its average value (this is an indication of the period of the noisy oscillations), and the time-dependent autocorrelation functions of the chemical concentrations. They showed that the smaller Ω is, the larger the variance in the above-mentioned distribution functions, and the less correlated are successive oscillations. All these analytical results were verified using stochastic simulations.

3.5. Other approaches

In this section we will briefly mention some modifications and enhancements made in the reviewed approaches, in light of their weaknesses, pointed out above.

As mentioned in the previous section, Boolean networks are, due to their relative simplicity, handy for the purpose of inferring the structure of a genetic network out of experimental data. However, their deterministic nature poses two problems in this context: First, since such data is noisy, and since the model itself is obviously an oversimplification of reality, a perfect match between the model and the experimental data may not be possible. Second, if the data is scarce (which is usually the case), there can be several models matching it.

In order to cope with the first problem, Akutsu et al. proposed [35] a modification for the basic model, which they termed *noisy Boolean networks*. Their model contains an additional parameter, p_{noise} ; this is an upper bound on the probability that the outcome of a logical function in the network will be inverted (that is, ‘0’ instead of ‘1’ and

vice-versa). They then presented an algorithm for inferring a noisy Boolean network given experimental gene expression patterns.

Shmulevich et al. [36] suggested a wider framework, termed *probabilistic Boolean networks* (PBN). Here, instead of defining a single logical function for each gene, it is assigned a *distribution* over the possible logical functions, specifying a probability of each function to appear. The authors showed that the dynamics of such networks can be studied within the context of Markov chains.

Generalized logical networks [6,40] are another extension of Boolean networks, allowing for a more detailed representation of gene expression levels (using integer rather than binary variables for their quantification), and for an asynchronous change of network elements. Note that incorporating asynchronicity into the model is still insufficient to ensure an accurate ordering of gene activations [7].

It is possible to insert stochasticity into continuous models, in the form of an additional noise term added to each differential equation; this results in the so-called Langevin equations. Under some conditions, these equations constitute a good approximation of the master equation; see Ref. [41] for a detailed discussion.

In the *hybrid Boolean-Continuous approach* (see for example Ref. [42]), genetic networks are modeled using “circuits” containing both types of elements: whenever a Boolean representation is suitable for the description of the activation of a gene, it is used; otherwise, the gene is modeled using continuous equations. The computation time required for simulating the resulting model is reported to be much smaller than that of ordinary continuous models.

4. Applying queuing theory to the modeling of genetic networks

In this section we show how queueing theory can be used to model genetic regulatory systems, in accordance with the ideas depicted in the introduction. We do this by presenting an example for such an application: using formulations and methods borrowed from queueing theory, a model for a genetic network of *arbitrary* size and structure is constructed and analyzed. Among other things, this model demonstrates how a common result in queueing theory—the existence of product form solutions—can be established for genetic networks.

We start by defining the model, and proceed to supply some of the results that can be derived for it. It is then employed to describe a specific regulatory network—that of the *lac* operon. We conclude with a brief discussion regarding the main features of the proposed model.

Note that the model considered here is not an adaptation of an existing queueing model to the biological problem at hand (although a queueing model with similar mathematical attributes, unrelated to a biological problem, was suggested by Gelenbe in Ref. [43]). Rather, it demonstrates how the analytical tools of queueing theory can be put to work in the biological domain. Indeed, this model exhibits a feature usually not found in queueing networks: the main events in it involve customers waiting at the queue, rather than customers completing service.




Fig. 3. Some of the possible stochastic transitions and their rates and probabilities, occurring at queue i . A plus sign is written beside transitions that increase the length of the target queue; a minus sign denotes transitions that decrease the length of the target queue. Note that these transitions apply to each queue in the network; as an example, one of the transitions occurring at queue m is depicted (the dashed arrow).

4.1. The model

We first present our model using queueing theory terms, and then make the analogy to regulatory networks.

Consider a network of n service facilities, each with its own queue and server. In every facility, there are 3 Poisson processes (see Fig. 3): (a) An arrival process, through which customers arrive to the queue from outside the system. We denote the rate of this process in queue i ($i = 1, \dots, n$) by γ_i ; (b) A service process, in which a customer is picked from the queue by the local server and receives the required service. The rate of this process is denoted by μ_i . After a customer receives service, he leaves the system; (c) A signal process, with rate ν_i . Whenever a signal reaches queue i , it induces one of the following 4 effects (once a signal acts, it leaves the system): (i) With probability p_{ij}^+ ($j = 1, \dots, n$), a customer will leave queue i and join queue j . If queue i is currently empty, nothing will happen. (ii) With probability p_{ij}^- , a customer will leave queue i , travel to queue j , pick a customer standing there, and they will *both* leave the system (that is, both queues i and j will lose a single customer). If queue i is currently empty, nothing will happen; if queue i is not empty but queue j is, only queue i will lose a customer; (iii) With probability p_{ijm}^+ , a customer will leave queue i , travel to queue j , cause a customer standing there to move to a third queue, m , and then leave the system (that is, queue i and queue j will both lose a single customer, queue m will gain a single customer). If queue i is currently empty, nothing will happen; if queue j is currently empty and queue i isn't, queue i will lose a single customer, and queue j and queue m will remain unchanged; (iv) With probability $d(i) = 1 - \sum_{j=1}^n p_{ij}^+ - \sum_{j=1}^n p_{ij}^- - \sum_{j=1}^n \sum_{m=1}^n p_{ijm}^+$, the customer will leave queue i and exit the system, without influencing any other queue. Of course, practically speaking, there's no point in allowing some of these probabilities to be positive; for

example, having $p_{ii}^+ > 0$ for some i implies an event that does not change the state of the system. However, from a pure mathematical perspective, no such restrictions are required.

One can see that unlike the case of Jackson networks, here the movement of customers between queues does not follow a completion of service, but is rather caused by the arrivals of signals, which act on customers waiting at the queue. The possible movements are defined by the sets of probabilities $\{p_{ij}^+\}$, $\{p_{ij}^-\}$ and $\{p_{ijm}^+\}$; thus, it is these probabilities that characterize the system as a *network*, with predefined connections between its elements.

Note that generally, all the mentioned rates and probabilities are allowed to be any functions of the current state of the network; this, in turn, is defined by the vector (k_1, k_2, \dots, k_n) , where k_i is the number of customers currently standing at queue i (including the one being served).

We can now discuss a possible analogy between this system and a genetic regulatory network. Each service facility represents a single, specific, type of molecule (be it an RNA molecule of some gene, a protein or a small molecule) or gene. The length of the queue can be considered to be the number of instances of that molecule type, or, if the service facility represents a gene—its expression level. Arrivals of customers from outside the system denote a rise in the number of molecules or expression level due to some external condition change; service, which causes a decrease in the queue length, can be considered as degradation, not related to any other quantity described by the model apart from the decreased quantity itself. The passage of a customer from queue i to queue j (the first type of signal effect) can be thought of as an enhancement of quantity j by quantity i , at the cost of some work performed by i ; for example, if queue i represents the expression level of some gene, then queue j may denote the number of mRNA molecules transcribed from that gene, and the referred change in quantities will represent a transcription of a single mRNA molecule. The loss of a customer in j due to a single arrival in i (the second type of effect induced by signals) can be looked at as a repression of i by j . For instance, if the length of queue i is the number of molecules of some protein, and queue j is the expression level of some gene, the mentioned effect can represent the repression of the transcription of gene j by protein i . The third type of effect—that of i and j teaming together to increase m —can represent the formation of complexes from simpler molecules. The final possibility—that of i simply losing a customer due to the arrival of a signal—can represent a reduction in expression level or number of molecules due to some external event or condition.

Thus, one can use the special case of queueing network presented above, for the description of a regulatory network, designating directly amounts of molecules or gene expression levels. The model is able to capture the chemical interactions between these elements, maintaining their discrete and stochastic nature. Since the rates and probabilities can be any function of the state of the whole network, the interactions described can be quite complicated (of course, this affects the ability to derive analytical results for the model, as discussed below). Interactions which cannot be represented explicitly by the model, can still be described indirectly, at least to some degree of success. For example, the formation of a complex from 3 types of molecules, can be described using an intermediate “queue”, representing the complex formed by 2 of these molecules;

the third molecule will then join the elements of that “queue” (a tight timing of these events can be obtained by assigning a relatively high rate to the signals affecting the intermediate queue). We note also that the possible events included in the model are sufficient for the representation, by analogy, of *any* logical function. This is explained further below.

4.2. Results

In this section we present the results that can be derived when studying a special case of the model suggested above—namely, when the rates of all the stochastic transitions involved are taken to be *constant*. Making this assumption allows for the derivation of several results, most notably an explicit expression of the joint probability distribution function of the network. We also give formulas for the marginal distributions of each element in the network, as well as their average values. Note the similarity of these results to those obtained for Jackson networks (see Section 2).

Let us start by presenting the notion of aggregate rates of “positive” customers and “negative” customers arriving to each queue. A “positive” customer arrival is a stochastic event that has the effect of adding a customer to the queue; a “negative” customer arrival is a stochastic event that has the effect of removing a customer from the queue (if the queue is not empty). By the word “aggregate” we mean customers arriving to the queue not only from outside the system, but also from other queues, due to events occurring there. The aggregate arrival rate of “positive” customers in queue i is denoted by λ_i^+ ; similarly, λ_i^- represents the aggregate rate of “negative” customers arriving to that queue.

It is possible to prove [44] that if the network reaches a steady state, then there exist numbers ρ_i , $0 < \rho_i < 1$, $i=1, \dots, n$, such that the probability of having (simultaneously) k_1 customers in the first queue (including the one begin served), k_2 customers in the second queue, and generally k_i customers in the i th queue ($i = 1, \dots, n$), is given by the formula

$$P(L_1 = k_1, L_2 = k_2, \dots, L_n = k_n) = \prod_{i=1}^n (1 - \rho_i) \rho_i^{k_i} . \tag{9}$$

The numbers ρ_i depend on the aggregate rates defined above, and obey the equations

$$\rho_i = \frac{\lambda_i^+}{\mu_i + \lambda_i^-}, \quad i = 1, \dots, n . \tag{10}$$

The aggregate rates, in turn, are found by solving the traffic equations of the network:

$$\lambda_i^+ = \gamma_i + \sum_{j=1}^n \rho_j \nu_j P_{ji}^+ + \sum_{j=1}^n \sum_{m=1}^n \rho_j \rho_m \nu_j P_{jmi}^+ , \tag{11}$$

$$\lambda_i^- = \nu_i + \sum_{j=1}^n \rho_j \nu_j P_{ji}^- + \sum_{j=1}^n \sum_{m=1}^n \rho_j \nu_j P_{jim}^+ . \tag{12}$$

To see that these are indeed the aggregate rates, one should keep in mind that here too, as in the Jackson networks, ρ_j is the proportion of time during which queue j is not

empty (we show this below explicitly). Thus, for example, $\rho_j v_j$ is the *effective* rate in which signals arriving to queue j induce some change on the system (remember that a signal arriving to an empty queue has no effect), and $\rho_j v_j \rho_{ji}^+$ is the effective rate in which customers move to queue i from queue j due to signals arriving to queue j . Using similar considerations completes the construction of the aggregate arrival rates for each queue.

As mentioned previously, the special form of the joint distribution function—a product of n factors, each depending on a single queue only—is called *product form*, and is relatively common in queueing networks. It allows us to easily derive several features regarding each queue separately, without specifying the state of the other queues in the network. For example, the probability that there are k_i customers in the i th queue is

$$P(L_i = k_i) = (1 - \rho_i) \rho_i^{k_i}. \quad (13)$$

In particular, the probability of having an empty queue is $1 - \rho_i$; that is, ρ_i is indeed the proportion of time during which the i th queue is not empty.

Another result readily available, due to the product form, is the average length of queue i , given by

$$E(L_i) = \frac{\rho_i}{1 - \rho_i}. \quad (14)$$

4.3. Application: Modeling the *lac* operon

We will now use the suggested model to describe a specific, well-studied regulatory circuit, which is responsible for controlling the expression of the *lac* operon in *E. coli* (see, for example, Ref. [26]). The proteins coded for by this operon are required for the transportation and breakdown of lactose. The regulation circuit ensures that these proteins will be synthesized only when lactose is present in the cell and a more favorable carbon source, namely glucose, is absent.

More specifically, the expression of the *lac* operon is regulated through a combination of a positive and a negative transcriptional controls; both are required for an effective transcription to occur. The positive regulation is accomplished through the CAP protein, which binds to the DNA near the promoter of the *lac* operon, and enhances its transcription considerably. In order for this to occur, CAP must also bind to cyclic AMP molecules; these, in turn, are abundantly found in the cell only when the level of glucose is low. Thus, *lac* is expressed only when glucose levels are low.

The negative control is implemented through the *lac* repressor protein. When present, this protein binds to an operator on the promoter of the *lac* operon, and prevents the RNA polymerase from starting the transcription. The presence of lactose causes the removal of this protein from the DNA, thus enabling the expression of *lac*.

The general model introduced in the previous section is now employed in representing this mechanism (see Fig. 4). Eight queues are used, designating glucose, lactose, cyclic AMP, CAP (not bound to cyclic AMP), the *lac* repressor protein, the *lac* operon, the mRNA produced by it and the synthesized proteins (the *lac* operon codes for a single mRNA molecule, which is translated into several proteins; a single queue is used to represent all of these proteins). We recall that the proposed general model allowed




Fig. 4. A queuing network model for the regulatory circuit of the *lac* operon. Each circle represents a biological element in the circuit (a “queue”). Arrows denote the possible transitions of “customers” between “queues”; a plus sign beside an arrow represents an increase in the “queue length” at the target of the transition; a minus sign implies a decrease. For each of the elements, a degradation (“service”) rate is also defined (not presented in the figure).

for signals inducing 4 types of effects: (i) The enhancement of one element by another; (ii) the repression of one element by another; (iii) the enhancement of an element by the joint effort of two other elements; (iv) the repression of an element due to an external condition. In this application, only the first 3 types of effects are employed. Effects of the first type involve the *lac* operon and its mRNA, and this mRNA and the lac proteins; Effects of the second type occur between glucose and cyclic AMP, between lactose and the lac repressor protein, and between the lac repressor protein and the *lac* operon. A single effect of the third kind is defined here—the enhancement of the *lac* operon by the combination of cyclic AMP and CAP. Service rates are defined for each of the elements in the model, representing degradation. Arrival rates are defined for each element whose molecules are not generated by other elements in the system.

The results obtained from this model (assuming fixed transition rates), both analytical and by using stochastic simulations, are presented in Fig. 5. Four scenarios were tested:




Fig. 5. Results for a model of *lac* regulation. The graphs to the left depict the distribution of the number of *lac* proteins molecules. For the other quantities modeled, only the averages are shown—in the graphs to the right (the “CAP” bar refers to CAP molecules not bound to cyclic AMP). Note that the rightmost bar in these graphs depicts the average of the distribution shown to the left. The figures display both predicted values (light shades) and values occurring in stochastic simulations (dark shades). The scenarios presented in each graph: (a) and (b)—no glucose, no lactose; (c) and (d)—glucose present, lactose absent; (e) and (f)—both glucose and lactose present; (g) and (h)—only lactose present. One can see that indeed, only in the last scenario the *lac* proteins are produced in high levels. Note that some of the quantities measured display a rather high variance in their amounts; this can account for the differences between predicted and encountered values, and reaffirms the necessity to study statistical properties other than the mere averages alone.

(i) low glucose level, low lactose level; (ii) high glucose level, low lactose level; (iii) high glucose level, high lactose level and (iv) low glucose level, high lactose level. The levels of glucose and lactose were controlled through the respective arrival rates. As required, the proteins serving for the utilization of lactose were produced in high quantities only in the fourth case.

4.4. Discussion

The main feature of the suggested model is the ability to obtain, for a genetic network of *any* size and configuration, the exact distribution function. Analytical results regarding arbitrary probabilistic genetic networks are quite rare (in fact, we are familiar only with Ref. [14]). We haven't encountered a previous derivation of the explicit joint distribution function of such a network.

On the other hand, these results come at the price of making a quite restricting assumption—that of fixed rates. The biological implication of this assumption is that the rate in which an element of the regulatory network—for example, some protein—propagates its effect on other elements in the network, is independent of its amount; once molecules of this element are present (or, in the case the element is a gene, once it is active), the rate of its effect is constant. Note that this does not mean that the elements in the network function as *Boolean* components: the larger the amount there is of a protein, the more time it will induce an effect on the network, and the more enzyme molecules, for example, will be required to turn it off. This model is fit to describe regulation phenomena involving more than two levels of activation; however, the timing of events, as well as their ordering, in some cases, may be erroneous. Thus, such a model should be probably best used in cases where general activity patterns are sought, rather than the exact time evolution of a regulatory system.

We note also that the possible chemical interactions included in the model are sufficient for the imitation, by analogy, of *any* logical function. The repression of one element by the other can be thought of as representing a *NOT* gate; the combination of two elements for the activation of a third element embodies an *AND* gate. Recalling results obtained in the theory of Boolean logic, we conclude that any logical function is attainable by combining these two elementary operators. Thus, the presented model possesses a considerable computational power.

The analytical results derived here, most notably the product form solution of the network, and the methods used for their deduction, are considered to be common practice in terms of queueing theory. Thus, by merely adopting the notion that the genetic network can also be viewed as a queueing network, one can readily earn a wealth of imported insights.

5. Queues and dynamic behaviours

As was mentioned in Section 2, most of the research conducted in queueing theory focuses on systems that have reached a steady state. Some attention has been given to the analysis of transient behaviours of queues (see Ref. [45]). However, there is almost

no work dealing with specific dynamic behaviours, such as oscillations or chaos (for examples of the few exceptions, see Refs. [46,47]; note, however, that these differ from the work presented here, both in the modeling scheme and in the analysis methods applied).

On the other hand, the mathematical field of dynamic systems is well established, with a vast body of knowledge accumulated over centuries. In addition, it has extensive applications in physics, chemistry and biology, as well as other areas of science (see, for example, Refs. [48,49]).

These facts come to mind, when one considers the idea that bridging queueing theory and computational biology may be profitable to *both* fields. Queueing theory can only gain, if one is able to import into it some of the insights arising in the study of dynamic systems in general, and regulatory circuits in particular.

One possible way to obtain this is through the fluid approximation of queueing models. As explained above, this scheme approximates the time evolution of the average behaviour of the stochastic system. One can start with a specific dynamic, deterministic, system, and then find a matching stochastic counterpart, such that the average behaviour of the latter is approximated by the former. This can be accomplished by correctly defining the rates of the stochastic transitions, as functions of the system state.

We now demonstrate this point. Consider a network of queues, where customers arrive to service facilities, leave them or move between them due to some stochastic events, generated by Poisson processes. In each queue, some of these events result in the arrival of customers to the queue, while others cause the departure of customers from the queue. Let us inspect the j th queue. Denote the rates of the former type of events (those increasing the length of the queue) by $\{r_{j1}(\mathbf{L}), \dots, r_{jR}(\mathbf{L})\}$, and those of the latter type by $\{q_{j1}(\mathbf{L}), \dots, q_{jQ}(\mathbf{L})\}$, where $\mathbf{L} = (L_1, \dots, L_n)$ is the vector of queue lengths, and R and Q are the numbers of rates of each type. That is, the rates are considered to be some functions of the state of the system. For the simplicity of this demonstration, it is assumed that each arrival or departure changes the queue length by a single customer only.

Since all the inspected events result from independent Poisson processes, the probability of each of them to occur during an infinitesimally short interval of time, $(t, t+h]$, is equal to the product of its rate and h , plus a small quantity $o(h)$ [23].³ In addition, up to one event can occur in this interval. Thus, the probability of an arrival of a customer to the queue during this interval is given by

$$P_j^+(\mathbf{L}) = \sum_{i=1}^R r_{ji}(\mathbf{L})h + o(h) \quad (15)$$

and the probability of a departure of a customer from the queue is equal to

$$P_j^-(\mathbf{L}) = \sum_{i=1}^Q q_{ji}(\mathbf{L})h + o(h). \quad (16)$$

³ $o(h)$ is defined as being of a smaller order of magnitude than h . That is, when h tends to 0, $o(h)/h$ vanishes.

It follows that the length of the queue at time $t + h$, given its length at time t , is

$$L_j(t + h) = L_j(t) + \xi_j(\mathbf{L}), \tag{17}$$

where ξ_j is the random variable designating the change in the queue’s length due to an event occurring in the interval $(t, t + h]$. ξ_j can be either 0, -1 or 1 , with probabilities depending on $P_j^+(\mathbf{L})$ and $P_j^-(\mathbf{L})$; hence, ξ_j depends indeed on the lengths of the other queues.

It may be desired to estimate the average behaviour of this system. To this end, the fluid approximation $\hat{\mathbf{L}}$ is presented. This size, which approximates $E(\mathbf{L})$, satisfies the equations

$$\begin{aligned} \hat{L}_j(t + h) &= \hat{L}_j(t) + 1 \cdot P_j^+(\hat{\mathbf{L}}) + (-1) \cdot P_j^-(\hat{\mathbf{L}}) + o(h) \\ &= \hat{L}_j(t) + h \left(\sum_{i=1}^R r_{ji}(\hat{\mathbf{L}}) - \sum_{i=1}^Q q_{ji}(\hat{\mathbf{L}}) \right) + o(h). \end{aligned} \tag{18}$$

Note that this is indeed only an approximation of the time evolution of $E(L_j)$. To see this, consider the case where one of the rates is given by the product of two queue lengths, L_i and L_k . In Eq. (18), this rate will be replaced by $\hat{L}_i \cdot \hat{L}_k$, representing its average value; however, it is not generally true that $E(L_i \cdot L_k) = E(L_i) \cdot E(L_k)$. Thus, this is a mere approximation of the actual average behaviour of $E(L_j)$.

Rearranging Eq. (18), dividing by h and taking the limit as $h \rightarrow 0$, leads to the next differential equation, approximating the time evolution of the average queue length:

$$\frac{d\hat{L}_j}{dt} = \sum_{i=1}^R r_{ji}(\hat{\mathbf{L}}) - \sum_{i=1}^Q q_{ji}(\hat{\mathbf{L}}). \tag{19}$$

Since no constraints are imposed on the functional form of the involved rates (provided that they all remain positive), the resulting differential equation can take a wide range of possible forms. Thus, an abundance of types of complex dynamic behaviours can be integrated into the queueing network.

In the analysis conducted here, it is desired to inspect the effect of the workload in the system on the divergence from the approximated average behaviour. Alternatively, one can consider the effect of the size of a change induced by a single stochastic event. Instead of assuming that each arrival increments the queue length by 1, and each departure decrements it by 1, we will set these changes to be $\pm \varepsilon$. By varying the size of ε , it is now possible to study the approximated average behaviour of the queueing system at different orders of magnitude of workloads. This leads to the following equation:

$$\frac{d\hat{L}_j}{dt} = \varepsilon \left[\sum_{i=1}^R r_{ji}(\hat{\mathbf{L}}) - \sum_{i=1}^Q q_{ji}(\hat{\mathbf{L}}) \right]. \tag{20}$$

The multiplication by ε implies a mere scaling of the time axis; thus, the approximated average behaviour itself does not qualitatively change—only the time it takes for its manifestation. However, the divergence of the actual behaviour from this approximated average does depend on ε .

The immediate benefit from such an analysis is the ability to recognize the local transition rates required to produce a desired system-wise behaviour, and by this gain, maybe, a more general insight regarding the relation between local interactions and global dynamics, in the context of queues. However, one needs not settle for this: the full strength of dynamic systems theory can be made available for the exploration of queueing systems, allowing for the utilization of such tools as phase space investigation, stability analysis, bifurcation analysis, and so on. Moreover, using tools developed in the study of the approximation schemes of stochastic models, it is possible to establish further results concerning the behaviour of stochastic queueing systems, such as the dispersion around the approximated average trajectory, the distribution of the period time (if the matching dynamic system displays cycles), the distribution of the duration of time such a system stays in one steady state before jumping to another (if multistability exists), and so forth.


An example of such a work, done in relation to a biochemical system, can be found in Ref. [15] (mentioned in Section 3). The authors there started with a set of differential equations, generating an oscillatory behaviour, which can be regarded as a simplified model of circadian oscillations. They then presented a stochastic model, describing the same biological sizes, and defined the transition rates in it so that indeed, its average behaviour can be approximated by the oscillations predicted by the deterministic model. They derived results similar to those suggested above. In particular, they discussed the validity of the fluid approximation, as a function of the system's size: the larger the system is, the more closer its behaviour to that of its deterministic counterpart.

Here we suggest to perform a similar analysis, but to further interpret the functional form of the transition rates in terms relevant to queueing systems. This is demonstrated in the following example.

Consider the next set of differential equations, consisting of the well-known Brusselator model, studied in chemistry [50]:

$$\begin{aligned}\frac{dX}{dt} &= A - (B + 1)X + X^2Y, \\ \frac{dY}{dt} &= BX - X^2Y.\end{aligned}\tag{21}$$

These equations describe the next (hypothetical) chemical reactions:



We can now consider a specific queueing network (a Jackson network with varying rates, as it is), whose average behaviour can be approximated by the dynamics of the Brusselator model. The network will consist of two queues. The length of the first queue matches X , while the length of the second queue coincides with Y . The respective fluid approximations are denoted by \hat{X} and \hat{Y} . Denote the external arrival rates to the queues by γ_1 and γ_2 , and the service rates by μ_1 and μ_2 . Let $p_{1,2}$ be the

probability that a customer completing service in the first queue moves to the second queue, and $d_1 = 1 - p_{1,2}$ the probability that the customer leaves the system. Similarly, define $p_{2,1}$ and d_2 .

We now follow the procedure outlined above. The analysis here is limited to the case where both queues are not empty. Since, in order for the presented approximation to be valid, it is assumed that the system is overloaded (that is, ε is close to 0), this is the common state of the network. However, to complete the picture presented here, the cases where at least one of the queues is empty should be investigated as well.

Consider the first queue. The stochastic events resulting in the increment of the length of this queue are the arrival of a customer from outside the system, with rate γ_1 , and the completion of service in the second queue, provided that the served customer moves to the first queue; the rate of this event is given by the product $\mu_2 p_{2,1}$ (there is no need to further multiply this size by the proportion of time during which the second queue is not empty, since, as stated above, we assume that both queues currently contain customers). A loss of a customer in the queue is caused solely by a completion of service, with rate μ_1 . Hence, the resulting differential equation is

$$\frac{d\hat{X}}{dt} = \varepsilon[(\gamma_1 + \mu_2 p_{2,1}) - \mu_1]. \quad (22)$$

A similar analysis of the second queue produces the equation

$$\frac{d\hat{Y}}{dt} = \varepsilon[(\gamma_2 + \mu_1 p_{1,2}) - \mu_2]. \quad (23)$$

In these equations, the rates are written in terms of \hat{X} and \hat{Y} , instead of X and Y . Comparing these equations to the Brusselator equations, one can see that the following choice of parameters results in an identical set of equations (up to a multiplication by ε):

$$\begin{aligned} \gamma_1 &= A, & \gamma_2 &= 0 \\ \mu_1 &= (B + 1)X, & \mu_2 &= X^2Y \\ p_{1,2} &= \frac{B}{B + 1}, & p_{2,1} &= 1 \\ d_1 &= \frac{1}{B + 1}, & d_2 &= 0 \end{aligned}$$

That is, customers reach the system from the outside at a constant rate, A , and join the first queue (the second queue does not have an external inflow of customers). These customers are served in a rate equal to $(B + 1)X$. After his service is completed, a customer leaves the system with probability $1/(B + 1)$, or joins the second queue with probability $B/(B + 1)$. The service rate in the second queue is given by X^2Y . A customer served in the second queue, returns to the first queue with probability 1. This scenario is depicted in Fig. 6.

As mentioned above, the distance of the *actual* behaviour from this approximated average depends on the workload in the system or, equivalently, on the size of the change induced by a single stochastic transition, ε : the smaller the change, the more regular the global behaviour of the system. This is demonstrated in Fig. 7, displaying




Fig. 6. A Jackson network displaying, approximately, an average behaviour of oscillations. Each circle represents a single queue. Customers arrive from the outside to queue 1 only, at a constant rate A . The service rates are written inside the circles; transition probabilities appear next to the arrows.

examples of actual trajectories produced by this network, for different sizes of ε . The resulting queueing system exhibits, roughly, a periodic behaviour, rarely encountered in queueing theory models and studies.

We will now suggest an interpretation to the functional form of the rates appearing in this model, focusing our attention on the rates which can be considered to be “peculiar”, in queueing theory terms—the service rates. The linear service rate in the first queue poses the lesser problem, since such rates do occur in queues: when a service facility has an infinite number of servers (or, more realistically, a number of servers which is practically larger than any number of waiting customers), the total service rate in the facility increases by a constant value (representing the service rate of a single server) with every new customer joining the system. So, explaining the form of this rate does not require us to employ unorthodox ideas.

The service rate of the second queue, however, is a different story. Here it follows that the longer the first queue is, the more efficient is the server in the second facility. To explain this odd situation, we suggest the notion of *customers acting as servers*: customers waiting in the first queue participate, in the meantime, in the service of customers standing in the second queue, enhancing the rate of service there. The fact that this rate increases with X^2 , rather than X , implies that this enhancement of service results from connections, or a cooperation, forming between pairs of customers standing in the first queue; the “computation” performed by the “network” of these customers is done by its “edges”, rather than its “nodes”. This scheme of computation is similar to the one encountered in neural networks. To complete the picture, the linear dependence on Y of this rate suggests that a large number of servers are present in the second facility as well, and that these servers are responsible for the utilization of customers waiting in the first queue.

The idea of waiting customers, which in the meantime supply some sort of service, is not so farfetched as it may first seem. In fact, such a scheme already exists for some time: by this, we refer to the SETI@home project, managed by the Space Sciences Laboratory at the University of California, Berkeley. This venture, which aims at searching for extraterrestrial intelligence, utilizes idle personal computers, connected to the Internet, using their computational power to process its data (recordings made by




Fig. 7. The occurrence of periodic behaviour in queuing networks. The queuing network defined in the text was simulated, using different sizes of ε , the size of a change induced by a single transition. The figures on the left column depict trajectories in the phase space of the system, in which the horizontal axis designates the length of the first queue (X) and the vertical axis the length of the second queue (Y); the figures on the right column show the change in time in the length of each queue (the first queue is shown in dark shades, the second in light shades). (a) and (b) describe the deterministic system, which approximates the average behaviour of the queuing network. The other Figures correspond to different sizes of ε : in (c) and (d), $\varepsilon = 0.01$; in (e) and (f), $\varepsilon = 0.05$; and in (g) and (h), $\varepsilon = 0.1$. As can be seen, the larger ε is, the further the system deviates from its approximated average behaviour. Notice that changing ε also scales the time axis.

radio telescopes). This is accomplished using a designated screen saver, installed on the participating PCs: when such a computer is waiting for user input long enough, this screen saver is activated, and the processing of SETI data takes place. Hence, the PC waiting for user response can be viewed as a customer waiting on a queue, supplying some service in the meantime.

Generally speaking, researching queuing networks which display complex dynamic behaviours will usually involve transitions rates depending on various components of the system, and in particular, the lengths of other queues in the network. Such

dependencies can be expected to have nontrivial, nonlinear forms. Thus, notions such as customers supplying service, or networks of customers performing some computation, will probably be recurring themes in such models. To these one must add the request that the queueing network is overloaded, so that the fluid approximation can be considered to be valid.

6. Conclusions

It is becoming increasingly acknowledged that genetic regulatory systems have discrete and stochastic aspects to their behaviour, and should be modeled accordingly. Thus, the study of these systems can benefit from the integration of knowledge acquired in existing disciplines, which specialize in modeling systems displaying such aspects. Queueing theory is an example of such a discipline.

We have demonstrated here the application of modeling and analysis techniques, borrowed from queueing theory, to the description of an arbitrary genetic network. This allowed for the derivation of the probability distribution function of the network. It should be clear, however, that there is more to accomplish from the interaction of the two disciplines. Queueing theory literature is extensive and diverse, and further delving into it, while still keeping in mind the biological problems at hand, can probably yield several other useful results.

In addition, we have suggested that queueing theory may benefit as well from this interdisciplinary dialogue. The employment of methodologies and insights gained in the study of dynamical systems in general, and the dynamics of genetic regulatory networks in particular, can promote the integration of complex behaviours into queueing models, enriching and diversifying the existing results.

Acknowledgements

The authors would like to thank Zeev Schuss for his helpful comments in preparing this paper. We would also like to thank Hila Rosenberg, for the illustration appearing in Fig. 1. E.B.J. would like to thank Mark Ptashne, for insightful conversations regarding genetic networks.

References

- [1] F.R. Blattner, *Science* 277 (1997) 1453.
- [2] S.T. Cole, *Nature* 393 (1998) 537.
- [3] The C. Elegans Sequencing Consortium, *Science* 282 (1998) 2012.
- [4] F.S. Collins, A. Patrinos, E. Jordan, A. Chakravati, R. Gesteland, L. Walters, *Science* 282 (1998) 682.
- [5] M. Ptashne, A. Gann, *Genes & Signals*, Cold Spring Harbor Laboratory Press, New York, 2002.
- [6] H. De Jong, *J. Comput. Biol.* 9 (2002) 67.
- [7] P. Smolen, D.A. Baxter, J.H. Byrne, *Bull. Math. Biol.* 62 (2000) 247.
- [8] H.H. McAdams, A. Arkin, *Annu. Rev. Biophys. Biomol. Struct.* 27 (1998) 199.
- [9] B.C. Goodwin, *Temporal Organization in Cells*, Academic Press, New York, 1963.

- [10] S.A. Kauffman, *J. Theor. Biol.* 22 (1969) 437.
- [11] A. Arkin, J. Ross, H.H. McAdams, *Genetics* 149 (1998) 1633.
- [12] T.B. Kepler, T.C. Elston, *Biophys. J.* 81 (2001) 3116.
- [13] J. Peccoud, B. Ycart, *Theor. Popul. Biol.* 48 (1995) 222.
- [14] M. Thattai, A. van Oudenaarden, *Proc. Natl. Acad. Sci. USA* 98 (2001) 8614.
- [15] D. Gonze, J. Halloy, P. Gaspard, *J. Chem. Phys.* 116 (2002) 10997.
- [16] L. Kleinrock, *Queueing System, Vol. 1*, John Wiley & Sons, New York, 1975.
- [17] T.G. Robertazzi, *Computer Networks and Systems: Queueing Theory and Performance Evaluation*, Springer, New York, 1990.
- [18] E. Gelenbe, *J. Appl. Probab.* 28 (1991) 656.
- [19] J.R. Artalejo, *Eur. J. Oper. Res.* 126 (2000) 233.
- [20] E. Gelenbe, *Lect. Notes Comput. Sci.* 2459 (2002) 1.
- [21] E. Gelenbe, *Neural Comput.* 1 (1989) 502.
- [22] H. Bakircioglu, T. Kocak, *Eur. J. Oper. Res.* 126 (2000) 319.
- [23] W. Feller, *An Introduction to Probability Theory and its Application*, Wiley, New York, 1968.
- [24] H.A. Kramers, *Physica* 7 (1940) 284.
- [25] J.E. Moyal, *J. Roy. Stat. Soc.* 11 (1949) 150.
- [26] B. Alberts, D. Bray, J. Lewis, M. Raff, K. Roberts, J.D. Watson, *Molecular Biology of the Cell*, 3rd Edition, Garland Publishing, New York, 1994.
- [27] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon, *Science* 298 (2002) 824.
- [28] N. Friedman, M. Linial, I. Nachman, D. Pe'er, *J. Comput. Biol.* 7 (2000) 601.
- [29] S.A. Kauffman, *Sci. Am.* 265 (1991) 78.
- [30] S.A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, New York, 1993.
- [31] R. Somogyi, C.A. Sniegoski, *Complexity* 1 (1996) 45.
- [32] Z. Szallasi, S. Liang, in: R.B. Altman, A.K. Dunker, L. Hunter, T.E. Klein (Eds.), *Proceedings of the Pacific Symposium on Biocomputing (PSB'98)*, Vol. 3, World Scientific Publishing, Singapore, 1998, pp. 66–76.
- [33] G. Weisbuch, *J. Theor. Biol.* 121 (1996) 255.
- [34] S. Liang, S. Fuhrman, R. Somogyi, in: R.B. Altman, A.K. Dunker, L. Hunter, T.E. Klein (Eds.), *Proceedings of Pacific Symposium on Biocomputing (PSB'98)*, Vol. 3, World Scientific Publishing, Singapore, 1998, pp. 66–76.
- [35] T. Akutsu, S. Miyano, S. Kuhara, *Bioinformatics* 16 (2000) 727.
- [36] I. Shmulevich, E.R. Dougherty, S. Kim, W. Zhang, *Bioinformatics* 18 (2002) 261.
- [37] H.H. McAdams, A. Arkin, *Trends Genet.* 15 (1999) 65.
- [38] D.T. Gillespie, *J. Phys. Chem.* 81 (1977) 2340.
- [39] M.A. Gibson, J. Bruck, *J. Phys. Chem. A* 104 (2000) 1876.
- [40] R. Thomas, *J. Theor. Biol.* 153 (1991) 1.
- [41] D.T. Gillespie, *J. Chem. Phys.* 113 (2000) 297.
- [42] H. McAdams, L. Shapiro, *Science* 269 (1995) 650.
- [43] E. Gelenbe, *J. Appl. Probab.* 30 (1993) 742.
- [44] A. Arazi, M.Sc. Dissertation, Tel-Aviv University, 2003.
- [45] H. Chen, D.D. Yao, *Fundamentals of Queueing Networks*, Springer, Berlin, 2001.
- [46] C. Chase, J. Serrano, P.J. Ramadge, *IEEE Trans. Automat. Control* 38 (1993) 70.
- [47] C.M. Rump, S. Stidham Jr., *Manage Sci.* 44 (1998) 246.
- [48] J.M.T. Thompson, H.B. Stewart, *Nonlinear Dynamics and Chaos*, Wiley, New York, 1986.
- [49] H. Haken, *Synergetics*, Springer, Berlin, 1977.
- [50] G. Nicolis, I. Prigogine, *Self-Organization in Non-Equilibrium Systems. From Dissipative Structures to Order Through Fluctuations*, Wiley, New York, 1977.