MDPI

Article

SMGen: A generator of synthetic models of biochemical reaction networks

Simone G. Riva [1]*, Paolo Cazzaniga [2,7,8], Marco S. Nobile [3,7,8], Simone Spolaor [4,7], Leonardo Rundo [5,6], Daniela Besozzi [4,7,8], and Andrea Tangherloni [2]*

1 Weatherall Institute of Molecular Medicine, Radcliffe Department of Medicine, University of Oxford, OX3 9DU, Oxford, UK
2 Department of Human and Social Sciences, University of Bergamo, 24129, Bergamo, Italy
3 Department of Industrial Engineering & Innovation Sciences, Eindhoven University of Technology, 5612 AZ, Eindhoven, The Netherlands
4 Department of Informatics, Systems and Communication, University of Milano-Bicocca, 20126, Milan, Italy
5 Department of Radiology, University of Cambridge, CB2 0QQ, Cambridge, United Kingdom
6 Cancer Research UK Cambridge Centre, University of Cambridge, CB2 0RE, Cambridge, United Kingdom
7 Bicocca Bioinformatics Biostatistics and Bioimaging Centre (B4), 20854, Vedano al Lambro, Italy
8 SYSBIO/ISBE.IT Centre for Systems Biology, 20126, Milan, Italy
* Correspondence: simone.riva@imm.ox.ac.uk (S.G.R.); andrea.tangherloni@unibg.it (A.T.)

Abstract: Several software tools for the simulation and analysis of biochemical reaction networks have been developed in the last decades; however, assessing and comparing their computational performance in executing the typical tasks of Computational Systems Biology can be limited by the lack of a standardized benchmarking approach. To overcome these limitations, we propose here a novel tool, named SMGen, designed to automatically generate synthetic models of reaction networks that, by construction, are characterized by both features (e.g., system connectivity, reaction discreteness) and non trivial emergent dynamics of real biochemical networks. The generation of synthetic models in SMGen is based on the definition of an undirected graph consisting of a single connected component, which generally results in a computationally demanding task. To avoid any burden in the execution time, SMGen exploits a Main-Worker paradigm to speed up the overall process. SMGen is also provided with a user-friendly Graphical User Interface that allows the user to easily set up all the parameters required to generate a set of synthetic models with any user-defined number of reactions and species. We analysed the computational performance of SMGen by generating batches of symmetric and asymmetric Reaction-based Models (RBMs) of increasing size, showing how a different number of reactions and/or species affects the generation time. Our results show that when the number of reactions is higher than the number of species, SMGen has to identify and correct high numbers of errors during the creation process of the RBMs, a circumstance that increases the overall running time. Still, SMGen can create synthetic models with 512 species and reactions in less than 7 seconds. The open-source code of SMGen is available on GitLab: https://gitlab.com/sgr34/smgen.

Keywords: Synthetic Models; Reaction-based Models; Biochemical Networks; Systems Biology

Citation: Simone G. Riva, Paolo Cazzaniga, Marco S. Nobile, Simone Spolaor, Leonardo Rundo, Daniela Besozzi and Andrea Tangherloni SMGen: A generator of synthetic models of biochemical reaction networks. *Symmetry* **2021**, *1*, 0. https://doi.org/

Received:
Accepted:
Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2021 by the authors. Submitted to *Symmetry* for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/ 4.0/).

1. Introduction

Systems Biology is a multidisciplinary research field that combines mathematical, computational, and experimental expertise to understand and predict the behavior of complex biological systems [1,2]. Among the different formalisms that can be used to describe intracellular processes, Reaction-based Models (RBMs) [3–5] are the most suitable for obtaining a detailed comprehension of the mechanisms that control the emergent behavior of the system under analysis [4]. The analysis of RBMs can be used to drive the design of focused lab experiments; to this aim, computational tasks such as Parameter Estimation (PE), Sensitivity Analysis (SA), and Parameter Sweep Analysis

(PSA) are generally applied [1,5–7]. Unfortunately, these computational tasks require the execution of huge amounts of simulations, so that the capabilities of biochemical simulators running on Central Processing Units (CPUs) (see, e.g., [8–10]) can be easily overtaken. Thus, several simulators exploiting Graphics Processing Units (GPUs) have been lately introduced to reduce the running times (see, e.g., [11–19]).

A crucial point, whenever new simulators are designed and implemented, regards the evaluation of their computational performance and their efficiency in executing the aforementioned demanding tasks. In this context, RBMs represent a key means as they can be exploited to run both stochastic simulation algorithms and (deterministic) numerical integration methods. Although there are more than $1,000$ models of real biochemical systems publicly available on BioModels [20,21], with some of them characterized by hundreds of species and reactions, the majority of these models do not follow the law of mass-action. In addition, these models do not have the structural characteristics (e.g., the number of species and the number of reactions) necessary to perform a fair comparison among the simulators under different scenarios. Though, only a limited number of RBMs is present in the literature (e.g., signal transduction pathways [22–24] or metabolic pathways [25]). Thus, the lack of detailed RBMs, especially those characterized by hundreds or thousands of reactions and molecular species, thus hampers the possibility of performing a thorough analysis of the performance of these simulators.

The computational performance of several GPU-powered tools have been already assessed using randomly generated synthetic RBMs [13,18,19]. However, only a few generators of biochemical models have been proposed so far, hindering the possibility of having a common and well-defined benchmarking approach. For instance, Komarov *et al.* [13,14] developed a tool to generate synthetic networks, which was then used to test the performance of their GPU-based simulators. Given the number of reactants, the type of reactions to be included in the RBM, and the total number of reactions, they generated synthetic RBMs by exploiting a hash table to avoid duplicates. The tool was then modified by randomly sampling the values of the initial concentrations of the species from a uniform distribution and the kinetic constants from a logarithmic distribution [18]. Another known and established model generator is the Reaction Mechanism Generator (RMG) [26], which was specifically developed to create synthetic chemical processes. RMG exploits an extensible set of 45 reaction families to generate elementary reactions from chemical species, while the reaction rates are estimated using a database of known rate rules and reaction templates. RMG relies on graphs to represent the chemical structures, and trees to represent thermodynamic and kinetic data. Due to its peculiarities, RMG was used to, e.g., automatically create kinetic models for the conversion of bio-oil to syngas through gasification [27]. Finally, other tools, such as Moleculizer [28], were introduced for the generation of reaction systems to obtain a deeper understanding of transduction networks.

Despite the efforts done to automatically define synthetic models, all these generators share a common drawback, that is, they have a limited flexibility and can generate only a restricted set of biochemical networks and processes. As a matter of fact, the existing methods generally do not perform any check on the generated models; while some of them can be only used to create a restricted set of models (e.g., RMG can generate only elementary reactions). Moreover, in the case of the approach proposed by Komarov *et al.*, even thought it can generate first- and second-order reactions, including degradation and production reactions—which are the most common reactions in real biochemical networks—the tool is not publicly available.

Considering the impelling necessity of defining a common benchmarking approach that allows for fairly evaluating and comparing different simulation approaches [29], we propose here a novel tool, named SMGen, designed to automatically generate synthetic biological networks codified as RBMs, which allow to obtain non trivial dynamics. It is worth mentioning that we are mainly interested in the generation of RBMs suitable for the analysis and comparison of the performance of the biochemical simulators. An

example of such RMBs is a model composed of plausible reactions (e.g., transformation, production and degradation reactions, which are common in real biochemical networks) leading to non trivial dynamics. For instance, in the case of a dynamics that instantly exhaust all reactants, some of the most advanced integration algorithms are able to simulate such stable and/or flat dynamics in just one computation step, thus hampering the possibility of a fair comparison among the different simulation approaches.

In addition, SMGen overcomes the existing approaches and tools, which generally do not allow for using different probability distributions for the initialization of the species amounts and the kinetic constants. As a matter of fact, the possibility of exploiting different probability distributions for the initialization of these parameters is an important features offered by SMGen, since the effort required for the simulation of the same model can vary according to the model parameterizations (i.e., species amounts and the kinetic constants).

SMGen adheres to well-defined structural characteristics based on graph theory and linear algebra properties, in particular, it exploits the definition of an undirected graph with a single connected component, which makes the whole generation process a computationally demanding task. To overcome this limitation, on the one hand, SMGen internally codifies all data structures by means of sparse matrices as well as *ad-hoc* structures specifically designed to avoid worthless values that would increase the running time required to generate RBMs. On the other hand, SMGen is able to drastically reduce the computational time by exploiting a Main-Worker paradigm used to distribute the overall generation process of RBMs onto multi-core CPUs. We show that SMGen can create, in less than 7 seconds, synthetic RBMs with hundreds of chemical species and molecular reactions, whose dynamics exhibit non trivial characteristics of real biochemical networks. Among the different features provided by SMGen, it allows for easily generating both symmetric and asymmetric RBMs: symmetric RBMs are composed of a number of species equal to the number of reactions, while in asymmetric RBMs the number of species can be lower than the number of reactions or vice-versa. From a computational point of view, the concept of symmetry is crucial in the analysis of complex networks to measure their information and entropy [30]. In addition, considering that every RBM can be converted into the corresponding system of coupled Ordinary Differential Equations (ODEs), studying the symmetries of this system of ODEs can reveal the intrinsic properties of the system of interest [31]. As a matter of fact, Ohlsson *et al.* pointed out that an alternative analysis of the system of ODEs can be carried out by considering the symmetries of the system solutions, aiming at formalizing the structures and behavior of the underlying dynamics of biological systems. Moreover, the possibility of evaluating GPU-powered simulators using symmetric and asymmetric RBMs is fundamental to understand their performance under different conditions. Indeed, a fair comparison would allow the user to select the best simulator based on characteristics of the RBM that has to be analysed. Thanks to its features and efficiency, SMGen was used to generate the synthetic RBMs necessary to realize a thorough comparison of the performance of different meta-heuristics in solving the PE problem of biochemical networks [3,5], which is one of the most common and difficult computational issues in Systems Biology. The outcome of such analyses showed that some well-known and widely used meta-heuristics, generally able to outperform all competitors in the optimization of benchmark functions, obtained poor performance when used to solve the PE problem. Moreover, SMGen was exploited to generate symmetric and asymmetric RBMs required for the in-depth analyses and comparisons of the computational performance of different biochemical simulators and to highlight their peculiarities [19]. In particular, these analyses allowed for determining the best simulator to employ under specific conditions, such as the size of the RBM and the total number of simulations to perform. SMGen can also be used for the generation of biochemical networks to investigate the performance of approaches specifically designed to tackle other well-known and computationally demanding tasks in Systems Biology (e.g., SA and PSA [1,7]).

SMGen allows also for exporting the generated RBMs into the Systems Biology Markup Language (SBML) [32], Version 4 Level 2, and into the BioSimWare standard [33], which is used by different GPU-powered simulators. Thus, we designed and developed SMGen to be a unifying, user-friendly, and standalone tool freely accessible to the Systems Biology community. The RBMs can be easily generated by using the provided user-friendly Graphical User Interface (GUI), which is designed to help the users in setting all the parameters required to generate the desired RBMs.

The manuscript is structured as follows. Section 2 describes the mathematical formalism of RBMs, as well as the structural characteristics that must be complied to generate synthetic biological networks. In addition, we provide all the algorithms and details at the basis of SMGen. Section 3 shows the experimental results achieved by SMGen. Finally, a discussion and conclusive remarks are provided in Section 4.

2. Materials and Methods

2.1. Reaction-based Models

An RBM is defined by specifying the set $\mathcal{S} = \{S_1, \ldots, S_N\}$ of N molecular species, and the set $\mathcal{R} = \{R_1, \ldots, R_M\}$ of M biochemical reactions that describe the interactions among the species appearing in $\mathcal{S}$. Each reaction R_i, with $i = 1, \ldots, M$, is defined as:

$$R_i : \sum_{j=1}^{N} a_{ij} S_j \xrightarrow{k_i} \sum_{j=1}^{N} b_{ij} S_j, \tag{1}$$

where a_{ij} and $b_{ij} \in \mathbb{N}$ are the stoichiometric coefficients, and $k_i \in \mathbb{R}^+$ is the kinetic constant associated with R_i. The stoichiometric coefficients specify how many molecules of species S_j, with $j = 1, \ldots, N$, appear either as reactants or products in reaction R_i. Note that some species might not appear in a reaction, so that the corresponding stoichiometric coefficient will be equal to 0. The order of a reaction is equal to the total number of molecules (of the same or different species) that appear as reactants in that reaction.

Each RBM can be written in the compact matrix-vector form $\mathbf{AS} \xrightarrow{\mathbf{K}} \mathbf{BS}$, where $\mathbf{S} = [S_1 \cdots S_N]^\top$ is the N-dimensional column vector of the molecular species, $\mathbf{K} = [k_1 \cdots k_M]^\top$ is the M-dimensional column vector of the kinetic constants, while $\mathbf{A}, \mathbf{B} \in \mathbb{N}^{M \times N}$ are the stoichiometric matrices, whose non-negative elements $[A]_{i,j}$ and $[B]_{i,j}$ correspond to the stoichiometric coefficients a_{ij} and b_{ij} of the reactants and products of the reactions, respectively.

Starting from an RBM and assuming the law of mass-action [34–36], the system of coupled ODEs corresponding to the RBM can be derived as follows:

$$\frac{d\mathbf{X}}{dt} = (\mathbf{B} - \mathbf{A})^T [\mathbf{K} \circ \mathbf{X}^{\mathbf{A}}], \tag{2}$$

where each ODE describes the variation in time of a species' concentration. In Equation 2 the N-dimensional vector $\mathbf{X} = [X_1 \cdots X_N]$ represents the concentration values of species $S_1, \ldots, S_N$, while $\mathbf{X}^{\mathbf{A}}$ is the vector-matrix exponentiation form [34]; the symbol $\circ$ denotes the entry-by-entry matrix multiplication (Hadamard product).

2.2. SMGen

In order to generate synthetic models of biochemical networks, SMGen complies with specific structural characteristics that the RBMs have to satisfy, that is:

- *System connectivity*: a biochemical network can be represented as an undirected graph with a single connected component, where the nodes represent the molecular species and the edges correspond to the species interactions. This representation easily allows for ensuring the system connectivity, a property that is strictly required to ensure that each species $S_j \in \mathcal{S}$, with $j = 1, \ldots, N$, will be involved in at least one reaction $R_i \in \mathcal{R}$, with $i = 1, \ldots, M$.

To be more precise, as a first step, an undirected graph with a single connected component is built. This undirected graph is randomly generated by using $N-1$ edges and by taking into account the maximum number of reactants and products, obtaining a connected biochemical reaction network. It is worth noting that a graph with a single connected component can be built if and only if the following condition is met:

$$(max_{num_r} + max_{num_p}) \times M > N,$$

where max_{num_r} and max_{num_p} are the maximum number of reactants and products, respectively. As a second step, starting from the initial undirected graph, the stoichiometric matrices, which can be viewed and treated as a Petri net [37,38] that, in turn, can be considered as a bipartite graph, are generated. Then, the stoichiometric matrices are randomly updated by adding and removing connections among the species, always taking into consideration the maximum number of reactants and products. Note that the initial connections, which correspond to the edges of the initial undirected graph, are never removed to ensure that all the species are involved in at least a reaction, maintaining the whole biochemical network connected. We designed an algorithm (see Algorithm 2) that builds the graph, composed of a single connected component, with the minimum number of edges that are needed to connect all the nodes.

- *Maximum number of reactants and products*: for each reaction $R_i \in \mathcal{R}$, with $i = 1, \ldots, M$, the number of reactants and the number of products cannot be arbitrarily large, but has to be lower than or equal to a user-defined values (i.e., max_{num_r} and max_{num_p}). Stated otherwise, the maximum order of the generated reactions should be fixed. It is worth mentioning that SMGen does not explicitly account for conservation conditions during the generation of the RBMs. However, it can generate reactions that are akin, for example, to protein modifications or conformational changes, thus resulting in two biochemical species whose sum is constant during the simulation.
- *Linear independence*: to ensure that each reaction R_i, with $i = 1, \ldots, M$, is endowed with plausible characteristics of a real biochemical reaction, the vectors of the stoichiometric coefficients of the reactants and products involved in R_i must be linearly independent. The linear independence condition allows for avoiding unrealistic reactions. An example of an unrealistic reaction consists in increasing or decreasing the amount of a species starting from one molecule of the species itself, e.g., $R_i : S_j \xrightarrow{k_i} \alpha S_j$, with $\alpha > 1$. Since the linear independence is evaluated between the reactants and the products of each reaction, it is related to the number of species and reactions. Thus, it could happen that when the number of species is much higher than the number of reactions it is not possible to build a graph with a single connected component. In such a case, duplicated reactions or linear dependent vectors between reactants and products of some reactions will occur. On the contrary, it is always possible to build a graph, composed of a single connected component, when the number of reactions is much higher than the number of species.
- *Reaction discreteness*: each reaction R_i, with $i = 1, \ldots, M$, must appear only once in the network, that is, duplicated reactions are not allowed.

SMGen is provided with a user-friendly GUI (see Figure 1) that allows the user to easily set up all the parameters required to generate the desired synthetic RBMs:

- the number of species N and the number of reactions M;
- the maximum number of reactants and products max_{num_r} and max_{num_p} that might appear in any reaction;
- the probability distribution $\mathcal{D}_s$ that is used to initialize the species amounts (to be chosen among uniform, normal, logarithmic or log-normal distributions). It is

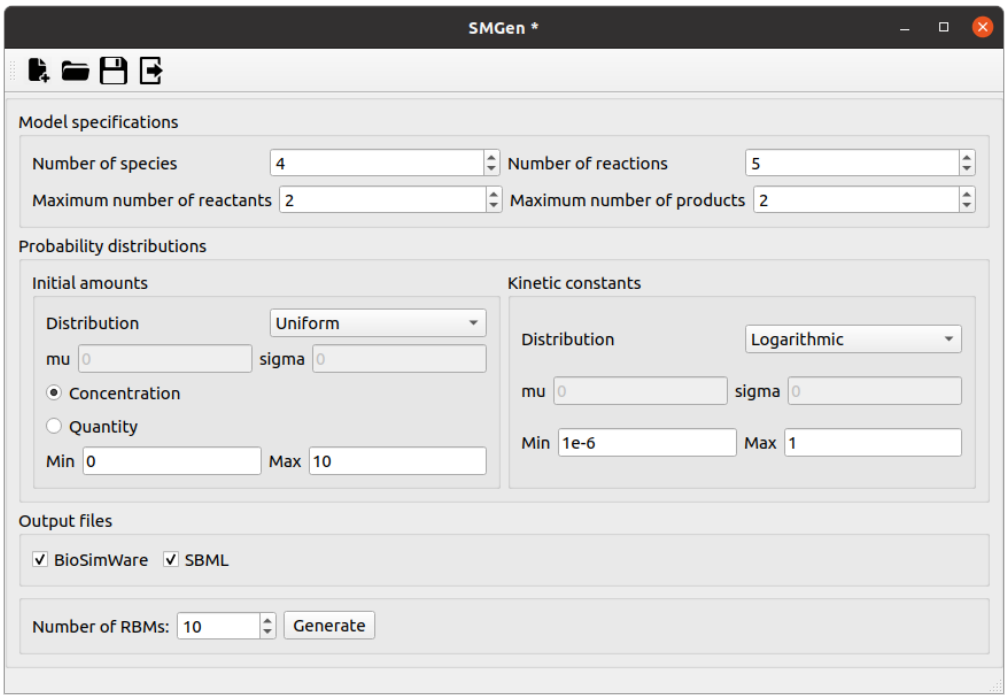

Figure 1. Graphical User Interface of SMGen. The user can set all the parameters to generate the desired RBMs, i.e., number of species and reactions, maximum number of reactants and products, probability distribution for the initial amounts and kinetic constants, and the output format file (i.e., BioSimWare, SBML).

worth mentioning that all species amounts are initialized using the same distribution probability;

- the minimum and maximum values min_s and max_s for the initial species amounts (to be specified either as number of molecules or concentrations);
- the probability distribution $\mathcal{D}_r$ that is used to set the values of the kinetic constants (to be chosen among uniform, normal, logarithmic or log-normal distributions). Note that all kinetic constants are generated using the same distribution probability;
- the minimum and maximum values min_r and max_r for the kinetic constants;
- the total number of RBMs that the user wants to generate;
- the output format file to export the generated RBMs (i.e., BioSimWare [33] and SBML [32]);
- the mean and standard deviation values μ_s and σ_s for the initial amounts—as well as the mean and standard deviation values μ_r and σ_r for the kinetic constants—must also be provided if the normal or log-normal distributions are selected.

Figure 2 shows a high-level scheme of the proposed implementation of SMGen, which exploits the Main-Worker paradigm to speed up the generation of the RBMs [39]. The user can specify the number of processes P—otherwise automatically set to the minimum value 3—which are used as follows:

- Proc_1 manages the GUI;
- Proc_2 is the Main process that orchestrates the computation;
- Proc_p, with $p = 3, \ldots, P$, are the Worker processes.

The whole functioning of SMGen can be summarized as follows:

- the user interacts with the GUI, managed by Proc_1, to fill in all the required values for the parameters necessary to create the RBMs;
- Proc_1 sends the values of all parameters to the Main process (Proc_2), which allocates the resources and distributes the work to the Workers (Proc_p, with $p = 3, \ldots, P$);

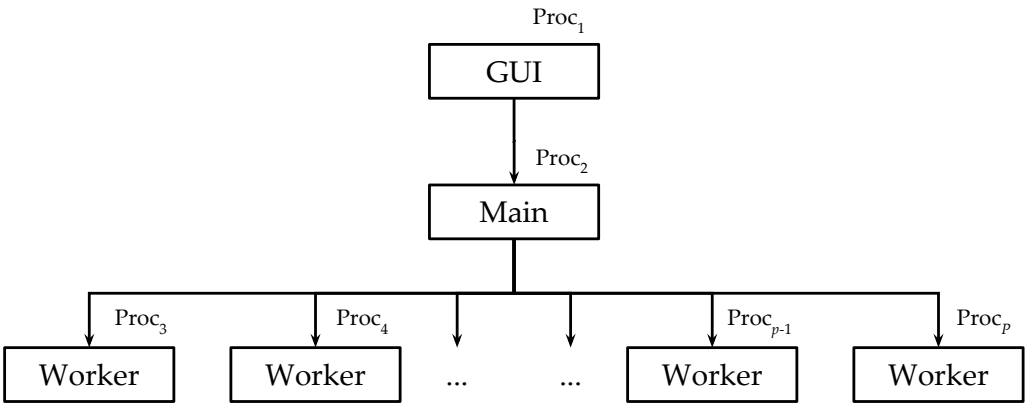

Figure 2. Scheme of the Main-Worker implementation of SMGen. The Main process (Proc_2) orchestrates all the available Workers (Proc_p, with $p = 3, \ldots, P$), which generate the RBMs in a distributed computing fashion.

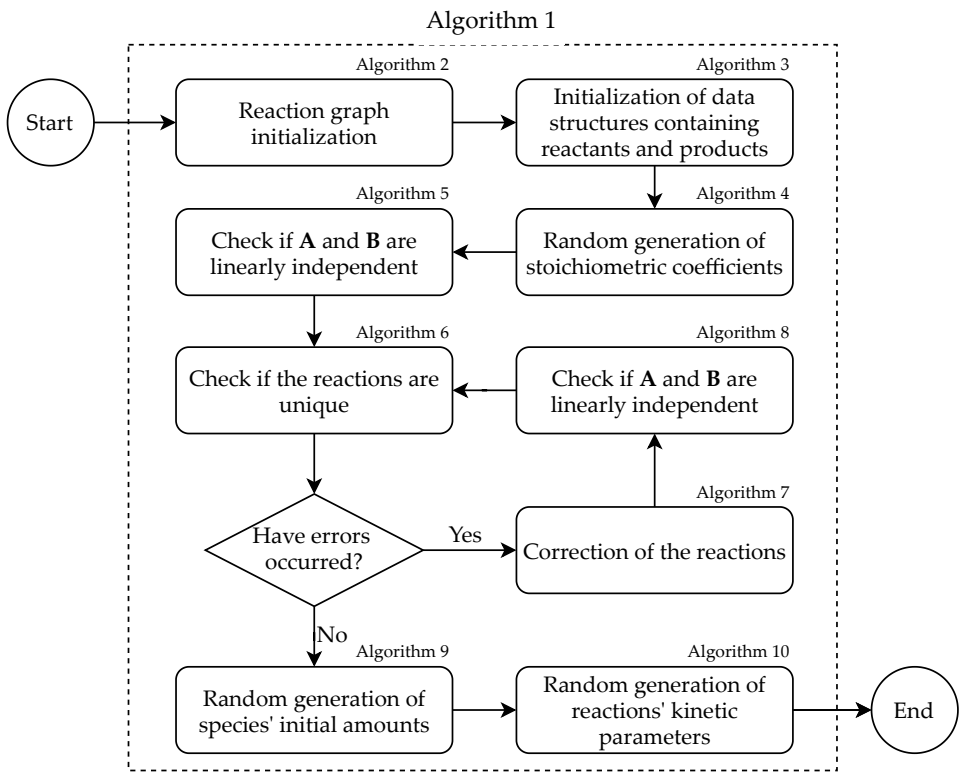

Figure 3. Workflow of a single Worker execution. First, the graph of the reactions is randomly initialized, and then converted into the data structures used to store the reactants and products. Second, the stoichiometric coefficients are randomly generated and the consistency of the reactants and products is verified. Third, the initial amounts of the species and the kinetic parameters of the reactions are randomly generated using the probability distributions specified by the user.

- each Worker (Proc_p, with $p = 3, \ldots, P$) generates an RBM. As soon as a Worker terminates its execution, it communicates to the Main process that the RBM has been created. If necessary, the Main process assigns the generation of other RBMs to idle Workers. When all required RBMs are obtained, the Workers enter in the death state, while the Main process waits for further instructions from Proc_1.

The workflow of each Worker consists in 9 different phases, in which a specific algorithm is executed (see Figure 3).

Algorithm 1 SMGen: workflow of a single Worker execution.

1: **function** GENERATOR($M, N, max_{num_p}, max_{num_r}, \mathcal{D}_s, \mathcal{D}_r, min_s, max_s, min_r, max_r, \mu_s, \sigma_s, \mu_r, \sigma_r$)
2: ## Algorithm 2
3: $\mathbf{G} \leftarrow$ GRAPH_GEN(N)
4: ## Algorithm 3
5: $\mathbf{A}, \mathbf{B} \leftarrow$ STOICH_MATRICES_GEN($M, N, \mathbf{G}$)
6: $AIJ[\cdot], BIJ[\cdot] \leftarrow [\,]$ ▷
7: **for** $i = 1$ to M **do** ▷
8: **for** $j = 1$ to N **do** ▷
9: **if** $\mathbf{A}[i, j] == 1$ **then** ▷
10: $AIJ \leftarrow AIJ \odot \langle i, j \rangle$ ▷
11: **if** $\mathbf{B}[i, j] == 1$ **then** ▷
12: $BIJ \leftarrow BIJ \odot \langle i, j \rangle$ ▷
13: ## Algorithm 4
14: $\mathbf{A}, \mathbf{B} \leftarrow$ STOICH_COEFFICIENTS_GEN($\mathbf{A}, \mathbf{B}, M, N, max_{num_r}, max_{num_p}, AIJ, BIJ$)
15: ## Algorithm 5
16: $err_{LinDep} \leftarrow$LINEAR_INDEPENDENCE_MATRIX($\mathbf{A}, \mathbf{B}, M$)
17: ## Algorithm 6
18: $err_{Repeat} \leftarrow$UNIQUE_REACTIONS($\mathbf{A}, \mathbf{B}, M$)
19: **while** $err_{LinDep} \wedge err_{Repeat}$ are not empty **do**
20: $rows_{Err} \leftarrow unique(err_{LinDep} \odot err_{Repeat})$
21: ## Algorithm 7
22: $\mathbf{A}, \mathbf{B} \leftarrow$ CORRECTION_REACTIONS($\mathbf{A}, \mathbf{B}, rows_{Err}, AIJ, BIJ, max_{num_r}, max_{num_p}$)
23: ## Algorithm 8
24: $err_{LinDep} \leftarrow$ LINEAR_INDEPENDENCE_REACTION($\mathbf{A}, \mathbf{B}, rows_{Err}$)
25: ## Algorithm 6
26: $err_{Repeat} \leftarrow$UNIQUE_REACTIONS($\mathbf{A}, \mathbf{B}, M$)
27: ## Algorithm 9
28: $\mathbf{M_0} \leftarrow$ AMOUNTS_GEN($N, \mathcal{D}_s, min_s, max_s, \mu_s, \sigma_s$)
29: ## Algorithm 10
30: $\mathbf{K} \leftarrow$ KINETIC_CONSTANTS_GEN($M, \mathcal{D}_r, min_r, max_r, \mu_r, \sigma_r$)

▷ $\rightarrow$ the instructions shown in lines 6–12 are required to build the structure of the initial graph of the reactions.

The pseudo-code reported in Algorithm 1 briefly summarizes all the steps required to generate a single RBM; the pseudo-code of the procedures invoked within Algorithm 1 are reported in Appendix A. For the sake of clarity, Table 1 lists the symbols used in the following description and in the pseudo-codes.

The steps performed by each Worker to generate an RBM are the following:

1. Given the parameters provided by the user, the graph representing the species and their interactions is randomly initialized (line 3 of Algorithm 1, see Algorithm 2).
2. The adjacency matrix of the graph generated in Step 1 is converted into the stoichiometric matrices **A** and **B** (line 5 of Algorithm 1, see Algorithm 3). Note that the instructions in lines 6–17 of Algorithm 1 are required to build the data structure of the initial graph, which is then modified.
3. The stoichiometric coefficients are randomly generated (line 19 of Algorithm 1, see Algorithm 4).
4. For each reaction R_i, with $i = 1, \ldots, M$, the linear independence between the reactants and products is verified (line 21 of Algorithm 1, see Algorithm 5).
5. The uniqueness of each reaction in the RBM is verified (line 23 of Algorithm 1, see Algorithm 6).
6. Any error in the RBM identified in the previous steps is corrected (line 27 of Algorithm 1, see Algorithm 7); the linear independence and the uniqueness of the

Symbol	Description
M	Number of reactions composing the RBM
N	Number of species involved in the RBM
max_{num_r}	Maximum number of the reactants
max_{num_p}	Maximum number of the products
$\mathbf{M_0}$	Array of the initial amounts
$\mathbf{K}$	Array of the kinetic constants
$\mathbf{A}$	Stoichiometric matrix of the reagents
$\mathbf{B}$	Stoichiometric matrix of the products
$\mathbf{G}$	Adjacency matrix of the graph of the reactions
$\mathcal{D}_s$	Probability distribution for the initial amounts
min_s	Minimum value of the initial amounts
max_s	Maximum value of the initial amounts
μ_s	Mean of the normal and log-normal distributions for the initial amounts
σ_s	Standard deviation of the normal and log-normal distributions for the initial amounts
$\mathcal{D}_r$	Probability distribution for the kinetic constants
min_r	Minimum value of the kinetic constants
max_r	Maximum value of the kinetic constants
μ_r	Mean of the normal and log-normal distributions for the initial amounts
σ_r	Standard deviation of the normal and log-normal distributions for the kinetic constants amounts
$\odot$	The concatenation operator

Table 1: List of symbols used in the pseudo-code of algorithms at the basis of SMGen.

reactions in the modified RBM are iteratively verified (lines 29 and 31 of Algorithm 1, see Algorithms 8 and 6, respectively).

7. The initial amounts of the species are generated according to the chosen probability distribution (line 33 of Algorithm 1, see Algorithm 9). If a species appears only as a reactant in the whole RBM, its amount is set to remain unaltered. The rationale behind this is double: on the one hand, we avoid the possibility of creating reactions that could be applied at most once, which is a highly improbable situation in biological systems; on the other hand, we mimic the non-limiting availability of some biochemical resources, for instance, it might be used to reproduce the execution of *in vitro* experiments where some species are continually introduced in the systems to keep their amount constant [40].
8. The kinetic constants of the reactions are generated according to the chosen probability distribution (line 35 of Algorithm 1, see Algorithm 10).

SMGen was developed using the Python programming language and exploiting `mpi4py` [41], which provides bindings of the Message Passing Interface (MPI) specifications for Python to leverage multi-core CPUs [42]. The open-source code of SMGen is available on GitLab (https://gitlab.com/sgr34/smgen) under the GNU GPL-3 license. In addition, it can be easily installed using the Python package installer pip (https://pypi.org/project/smgenerator).

3. Results

We analyzed the performance of SMGen regarding both its capability of creating RBMs characterized by non trivial dynamics, and the computational time required to generate sets of RBMs of increasing size. All tests were executed on a workstation equipped with an Intel Core i7-8750H CPU (clock 4.1 GHz), 16 GB of RAM and a

No.	Reagents	Products	Constant
R_1	$X_0 + X_4$	X_3	$4.295 \cdot 10^{-5}$
R_2	X_4	$X_1 + 2X_2$	$2.207 \cdot 10^{-2}$
R_3	X_4	$X_2 + X_4$	$7.070 \cdot 10^{-4}$
R_4	$X_1 + X_4$	$X_2 + X_3$	$4.613 \cdot 10^{-2}$

Table 2: List of the reactions of an RBM with by 4 reactions and 5 species generated by SMGen.

Species	Initial amount
X_0	4
X_1	8
X_2	7
X_3	8
X_4	1

Table 3: Initial molecular amounts of the RBM generated by SMGen shown in Table 2.

Samsung 970 EVO solid-state drive NVMe PCIe (up to 3400 MB/s and 1500 MB/s read and write speed, respectively), running on Ubuntu 20.04 LTS.

As a first batch of tests, we generated 100 synthetic RBMs characterized by a limited number of reactions and species (4 and 5, respectively), and we analysed their characteristics and dynamics. We set to 3 both the maximum number of reactants max_{num_r} and products max_{num_p}. We sampled the initial amounts of species from a normal distribution with mean $\mu_s = 5$ and standard deviation $\sigma_s = 5$, considering a minimum value $min_s = 0$ and maximum value $max_s = 10$. The kinetic constants were instead sampled from a logarithmic distribution with minimum value $min_r = 10^{-16}$ and maximum value $max_r = 10$.

Table 2 shows the list of reactions along with the kinetic constants of one of these 100 synthetic RBMs. Since the species X_0 appears only as a reactant, its amount will be kept constant during the simulation. The initial molecular amounts of all species—given as number of molecules—are listed in Table 3. This small RBM includes the basic "cascade of reactions" structure typically observed in signaling pathways, starting from the source represented by species X_0 and X_4, toward species X_2 and X_3.

We simulated the dynamics of this RBM for 50 time steps (arbitrary units), and the achieved dynamics are shown in Figure 4. These plots evidence that, although the RBM was randomly generated by SMGen, the simulated behavior is non trivial (e.g., the reactants are not instantly exhausted, resulting in a flat dynamics). Note that currently SMGen can be only used to generate RBMs and does not include any simulation tools yet. The simulation of the dynamics of this RBM was carried out using FiCoS [19]. It is worth mentioning that obtaining synthetic RBMs exhibiting non trivial dynamics is fundamental to perform in-depth computational analyses and comparisons among the existing and the novel simulators. Indeed, in the case of stable or flat dynamics, or when the overall behavior of the network is extremely fast and instantly exhausts all the reactants, the most advanced integration algorithms are able to simulate the emergent dynamics in just one computation step [19]. In such a case, the computational performance of the simulation tools is only partially assessed, thus hindering a fair comparison among the tools.

As a second batch of tests, we evaluated the computational performance of SMGen exploiting the Main-Worker paradigm running on 4 distinct cores of the CPU. First, we considered the generation of symmetric RBMs with an increasing number of species and reactions (i.e., $M = N = 2^x$, with $x = 2, \ldots, 9$). The initial amounts and kinetic constants were randomly sampled from a uniform distribution with minimum values $min_s = min_r = 0$ and maximum values $max_s = max_r = 10$. We also varied the maximum

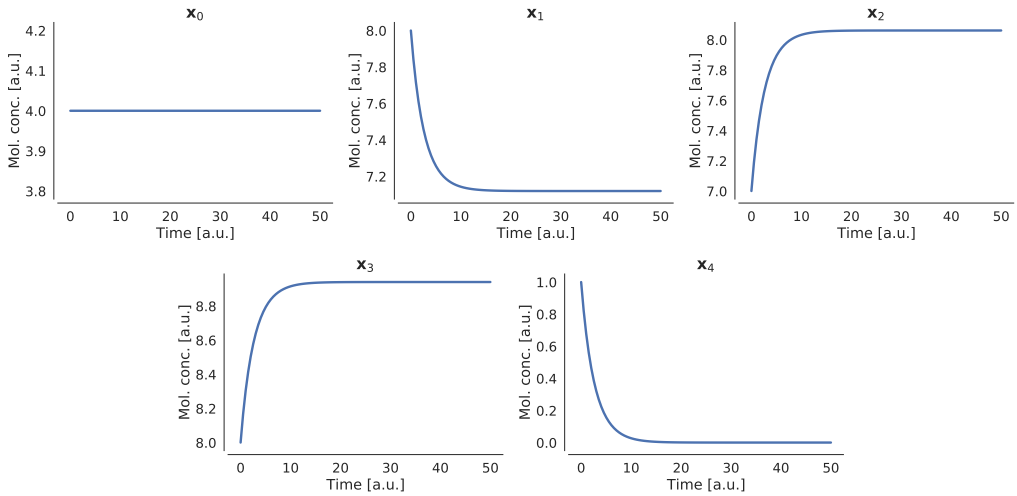

Figure 4. Dynamics of the species of the synthetic RBM generated by SMGen shown in Table 2.

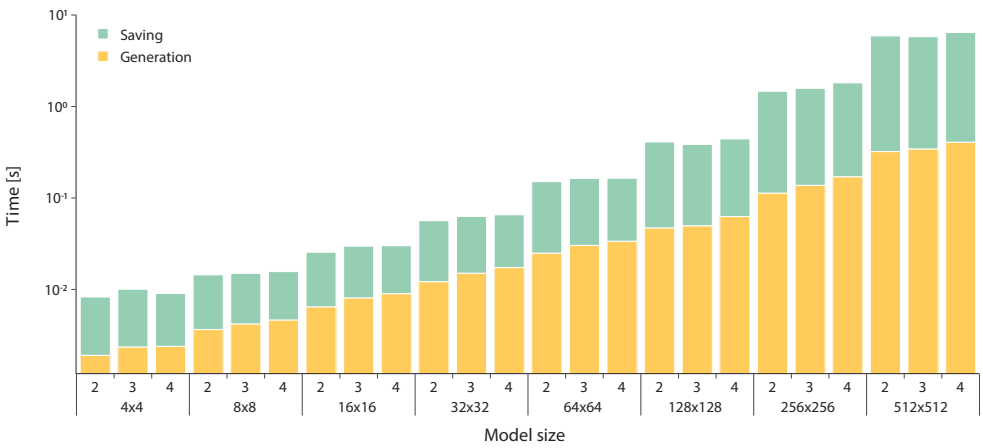

Figure 5. Stacked bar plot showing the average generation time (yellow bars) and the average saving time (green bars) required by SMGen to generate a symmetric RBM. Note that the y-axis is in logarithmic scale.

numbers of reactants and products considering the set of values $\{2, 3, 4\}$ and setting $max_{num_r} = max_{num_p}$. For each of the resulting 24 parameters combinations, we created 100 RBMs to collect statistically sound results about the performance of SMGen. As described in Section 2, two kinds of error can occur during the generation of an RBM: a linear dependence between reactants and products, and duplicated reactions. Since the correction of these errors is one of the most time-consuming phases of SMGen, we separately measured the generation time, which indicates the running time spent by SMGen to generate an RBM, and the saving time, which refers to the writing operations on the solid-state drive. Figure 5 shows the average running time required by SMGen to generate and save an RBM. As expected, both the generation and saving time increase along with the number of species and reactions of the RBM. Moreover, we observe that the maximum number of reactants and products have a slight impact on both the generation and the saving time; in most of the cases, increasing these values results in a higher running time.

Finally, we exploited SMGen for the creation of asymmetric RBMs, to evaluate how a different number of species and reactions affects the running time. As in the case

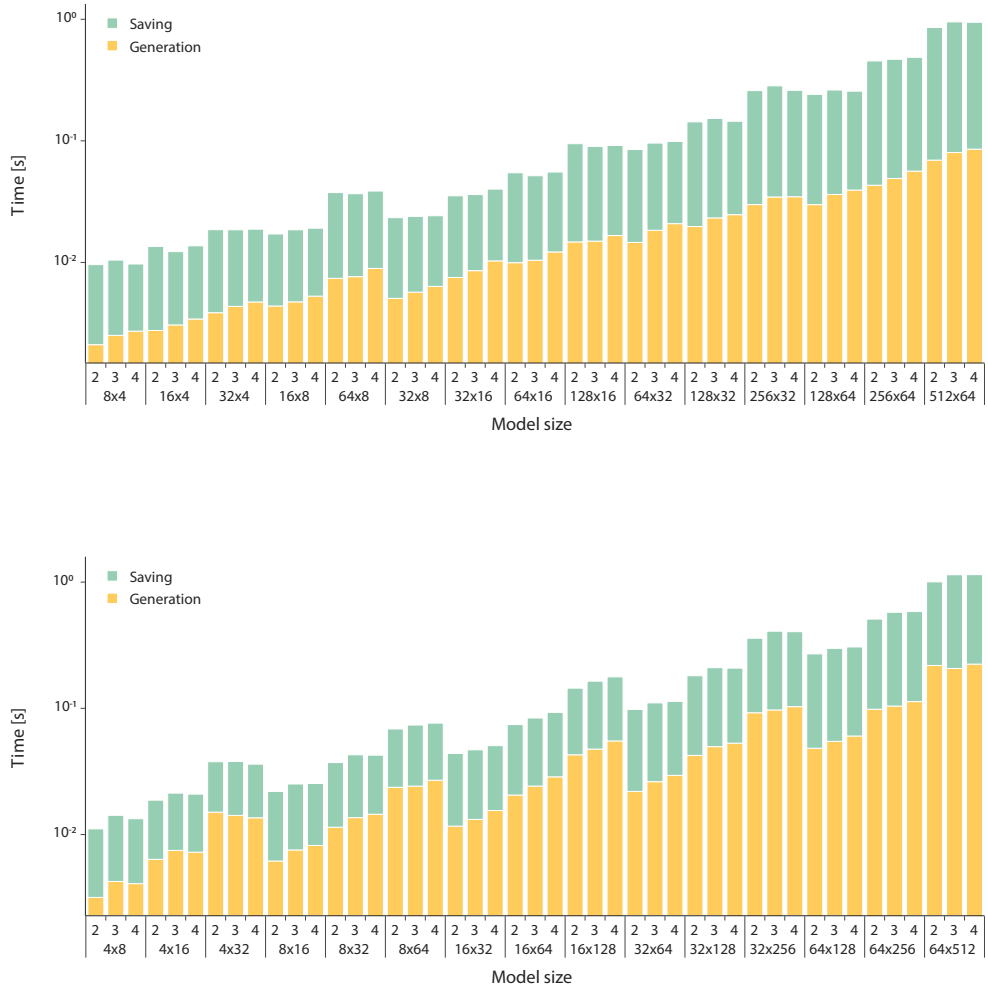

Figure 6. Stacked bar plots showing the average generation time (yellow bars) and the average saving time (green bars) required by SMGen to generate an asymmetric RBM with more species than reactions (top) and with more reactions than species (bottom). Note that the y-axes are in logarithmic scale.

of symmetric RBMs, we measured both the generation time and the saving time. The asymmetric RBMs were created as follows:

- we set the number of species $N \in \{4, 8, 16, 32, 64\}$, and then we varied the number of reactions $M \in \{2N, 4N, 8N\}$;
- we set the number of reactions $M \in \{4, 8, 16, 32, 64\}$, and then we varied the number of species $N \in \{2M, 4M, 8M\}$;
- we varied both the maximum numbers of reactants max_{num_r} and products max_{num_p} in $\{2, 3, 4\}$.

In such a way, we obtained a total of 90 different combinations of the parameters (i.e., number of species, number of reactions, and maximum number of reactants and products) to be tested; as in the previous tests, for each combination we generated 100 RBMs to collect statistically sound results. Figure 6 shows the average running time required to create RBMs with dimensions $N \times M$, highlighting once again that both the generation time and the saving time increase along with the size of the RBMs. As in the case of symmetric RBMs, we observed the same effect due to the maximum number of reactants and products allowed in the reactions. As expected, when there are more

reactions than species (bottom panel in Figure 6) the generation times are higher than the opposite situation (top panel in Figure 6). This circumstance is due to the potential higher number of errors that SMGen has to identify and correct. Indeed, when $M \gg N$, the probability that repeated reactions are randomly generated is higher than the case when $N \gg M$, because the number of admissible reactions strictly depends on the number of species.

4. Conclusions

In this work we presented SMGen, a generator of synthetic RBMs displaying the characteristics of real biochemical networks, which can be exploited to create benchmarks for the evaluation of novel and existing simulators. SMGen is particularly suitable to create the RBMs necessary to assess the performance of GPU-based simulators. As a matter of fact, the performance of GPU-powered simulators can drastically change with the number of chemical species and reactions composing an RBM. Considering that every RBM can be converted into the corresponding system of coupled ODEs, the resolution of this system of ODEs can be performed in a parallel fashion, where each ODE is resolved by a thread. Since each ODE corresponds to a specific chemical species, a higher number of species generally lead to a higher parallelization, increasing the computational performance of the simulator. On the contrary, considering that the number of the reactions composing the biological system is roughly related to the length of each ODE, in terms of the mathematical complexity, the higher the number of reactions the higher the number of operations that must be performed by each thread, leading to a higher running time [18,19].

SMGen was developed in Python and was designed to be a unifying, user-friendly, and standalone tool. In addition, SMGen exploits the Main-Worker paradigm to speed up the generation of RBMs; this was implemented using the `mpi4py` [41] library, where the first process manages the GUI, the second one is the Main process, and all the other processes are the Workers that generate the RBMs in a distributed computing fashion. Thanks to the GUI of SMGen, the user can easily set up all the parameters characterizing the required RBMs, e.g., the number of species and reactions, the maximum number of reactants and products per reaction, the probability distributions (uniform, normal, logarithmic, log-normal) to generate the initial amounts of the species and the values of the kinetic constants associated with the reactions, the output file format to save the RBMs. It is noteworthy that the performance of SMGen is not affected by the choice of different distributions for the sampling of the species amounts and kinetic constants. On the contrary, different distributions can drastically affect the dynamics of the model (see, e.g., the Brusselator model [38,43]). The analysis of the features of the generated models, according to different distributions, will be addressed in a future work.

We assessed the capabilities of SMGen for the creation of RBMs characterized by a non trivial behavior, and we presented an example of a synthetic RBM together with the simulated dynamics. We also tested the computational performance of SMGen by generating batches of symmetric and asymmetric RBMs of increasing size, showing the impact of the number of reactions and species, and of the number of reactants and products per reaction, on the generation times. We observed that when the number of reactions is higher than the number of species, SMGen generally identifies and corrects high numbers of errors during the creation process of the RBMs, a circumstance that inevitably increases the overall running time.

As a future extension of this work, we plan to develop an Application Programming Interface (API), so that SMGen can be seamlessly integrated into other processing pipelines, tools, and simulators. We are also developing a well-documented command-line interface, which will be released together with the API, to increase the user experience as well as the integrability with the existing approaches. Another feature that could be introduced is a module for the visualization of the networks. However, even though the visualization of the networks might help the users, plotting RBMs composed

of thousands of species and reactions is quite a hard task. Moreover, the visualization of such large RBMs might not add any relevant information. We will also introduce a new feature specifically developed to generate feedback loops in synthetic RBMs, exploiting the theory of Petri nets [37,38]. Feedback loops are fundamental elements of biological processes that lead to the establishment of oscillatory regimes and non-linear dynamics [22]. Moreover, we plan to develop a function to rescale the kinetic constants by taking into account the number of reactants involved in the reactions. As a matter of fact, both the order of magnitude and the value of each kinetic constant are related to the number of reactants composing the corresponding reaction. In the current version of SMGen, we assume that all reactions follow the law of mass-action; however, we will implement additional kinetics (e.g., Michaelis-Menten and Hill kinetics [44,45]) in the future.

SMGen relies on graph theory and linear algebra properties to comply with specific structural characteristics that are essential to create synthetic but real RBMs. Nevertheless, the current version of SMGen does not exploit network graphlets and motifs, which characterize families of real biological networks. Graphlets are induced sub-graphs appearing at any frequency because they are independent of the network's null model (i.e., a random graph model defined by a probability distribution), while network motifs are repeated sub-graphs appearing with a frequency higher than in random graphs and depending on the network's null model [46]. Graphlets have been used to analyze local network structures, and to cluster different network types [46]. Probabilistic graphlets have been also developed to analyze the local wiring patterns of probabilistic networks. When applied to study biological networks, probabilistic graphlets resulted in a robust tool, able to capture the biological information underlying the biological networks, thanks to the capabilities of managing the low signal topology information [47]. Graphlets can also be used to measure the structural similarity among large networks, calculating the graphlet degree distribution [48]. Motifs can be used to study autoregulation, single input module, dense overlapping regulons and feedback loops, as well as to reveal answers to many important biological questions [49,50]. The frequencies of the motifs were also exploited as classifiers for the selection of the network models [51]. Studying the motifs is also fundamental to understand the stability and robustness of the biological networks to small perturbations [52]. As a matter of fact, Prill *et al.* showed that the stability and robustness of a network is strictly related to the relative abundance of the motifs [52]. This result suggests that the structural organization of a biological network can be highly related to the dynamic properties of the small network motifs. We plan to extend SMGen to incorporate network graphlets and motifs during the generation of the RBMs, to obtain synthetic models with an improved biological significance.

In addition, we will modify the random generation of the kinetic constants to take into account the fact that different biological processes can operate on different time scales. In order to introduce this modification, as a first step the reactions will be grouped into different families based on the biological process that they describe (e.g., protein or mRNA degradation, transcription rates, translation rates). Then, for each family of reactions, a different probability distribution will be used to sample the kinetic constants to reflect the time scale of the described biological process. We are also investigating the possibility of generating models characterized by specific emergent dynamical aspects (e.g., stable oscillatory regimes, state switching, multistability [16]). We derived a set of heuristics to detect these conditions from the time-series, based on the combination of SMGen with GPU-powered stochastic simulators [16]. Still, regardless of the acceleration used, robustly detecting these phenomena is computationally challenging because some parameterizations might lead to very stiff models which, in turn, might require a very long running time. Thus, we are investigating alternative approaches to guide the generation of synthetic models having such features. In particular, we are considering the possibility of providing a set of "functional modules", i.e, a group of already parameterized reactions implementing

specific behaviors. These modules could be coerced into the models generated with SMGen, possibly leading to the emergence of the desired phenomena. Finally, we plan to include an initial check of the parameter values set by the user, based on some heuristics, to verify whether the RBMs can be actually generated as requested. This initial step will allow for avoiding worthless calculations and to suggest useful modifications of parameters to the user.

Author Contributions: S.G.R., P.C., M.S.N., and A.T. conceived and designed the tool. S.G.R. developed the tool. S.G.R., P.C., and A.T. conceived and designed the analyses. S.G.R. performed the analyses. P.C., D.B., and A.T. wrote the paper. S.G.R., M.S.N., S.S., and L.R. reviewed the paper. D.B., P.C., M.S.N., and A.T. supervised the whole work. All authors have read and agreed to published this version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Algorithms

We report here all the algorithms referring to the functions called by Algorithm 1, which represents the workflow of each Worker process.

Algorithm 2 Random initialization of the graph of reactions

```
function GRAPH_GEN(N)
    G[·,·], v[·] ← 0
    for j = 1 to N do
        v[j] ← j
    ind_1 ← random(1, len(v))
    ind_2 ← random(1, len(v))
    n_1, n_2 ← v[ind_1], v[ind_2]
    v ← delete(v[ind_1])
    v ← delete(v[ind_2])
    G[n_1, n_2] ← 1
    while v is not empty do
        if random ∈ {0,1} == 0 then
            ind ← random(1, len(v))
            k ← v[ind]
            v ← delete(v[ind])
            G[n_1, k] ← 1
            n_2 ← k
        else
            ind ← random(1, len(v))
            k ← v[ind]
            v ← delete(v[ind])
            G[k, n_2] ← 1
            n_1 ← k
    return G
```

Algorithm 3 Conversion of the adjacency matrix **G** into the stoichiometric matrices **A** and **B**

```
function STOICH_MATRICES_GEN(M, N, G)
    i ← 1
    for n_1 = 1 to N do
        for n_2 = 1 to N do
            if G[n_1, n_2] == 1 then
                A[i, n_1] ← 1
                B[i, n_2] ← 1
                if i == M then
                    i ← 1
                else
                    i ← i + 1
    return A, B
```

Algorithm 4 Generation of the random stoichiometric coefficients

```
function STOICH_COEFFICIENTS_GEN(A, B, M, N, max_num_r, max_num_p, AIJ, BIJ)
    for i = 1 to M do
        for k = 0 to max_num_r do
            coef ← random(0, max_num_r)
            j ← random(1, N)
            if coef ≠ 0 & A[i, ·] + coef − A[i, j] ≤ max_num_r then
                A[i, j] ← coef
            else if coef == 0 & ⟨i, j⟩ ∉ AIJ then
                A[i, j] ← coef
    for i = 1 to M do
        for k = 0 to max_num_p do
            coef ← random(0, max_num_p)
            j ← random(1, N)
            if coef ≠ 0 & B[i, ·] + coef − B[i, j] ≤ max_num_p then
                B[i, j] ← coef
            else if coef == 0 & ⟨i, j⟩ ∉ BIJ then
                B[i, j] ← coef
    return A, B
```

Algorithm 5 Checking the linear independence between $\mathbf{A}[i,\cdot]$ and $\mathbf{B}[i,\cdot]$

```
function LINEAR_INDEPENDENCE_MATRIX(A, B, M)
    err_LinDep ← [ ]
    for i = 1 to M do
        if A[i, ·] ∧ B[i, ·] are linearly dependent then
            err_LinDep ← err_LinDep ⊙ i
    return err_LinDep
```

Algorithm 6 Checking if the generated reactions are unique

```
function UNIQUE_REACTIONS(A, B, M)
    AB, ABs ← [ ]
    err_Repeat ← [ ]
    for i = 1 to M do
        AB[i] ← A[i, ·] ⊙ B[i, ·]
    for i = 1 to M do
        if AB[i] is in ABs then
            err_Repeat ← err_Repeat ⊙ i
        else
            ABs ← ABs ⊙ AB[i]
    return err_Repeat
```

Algorithm 7 Random correction of the repeated reactions

```
function CORRECTION_REACTIONS(A, B, rows_Err, AIJ, BIJ, max_num_r, max_num_p)
    for i = 1 to len(rows_Err) do
        A[rows_Err[i], ·] ← 0
        B[rows_Err[i], ·] ← 0
        if rows_Err[i] ∈ AIJ[·, 1] then
            for k = 1 to len(AIJ) do
                if AIJ[k, 1] == rows_Err[i] then
                    A[AIJ[k, 1], AIJ[k, 2]] ← 1
            for c = 0 to max_num_r do
                coef ← random(0, max_num_r)
                col ← random(1, N)
                if coef ≠ 0 & A[rows_Err[i], ·] + coef − A[rows_Err[i], col] ≤ max_num_r then
                    A[rows_Err[i], col] ← coef
                else if coef == 0 & ⟨rows_Err[i], col⟩ ∉ AIJ then
                    A[rows_Err[i], col] ← coef
        if rows_Err[i] ∈ BIJ[·, 1] then
            for k = 1 to len(BIJ) do
                if BIJ[k, 1] == rows_Err[i] then
                    B[BIJ[k, 1], BIJ[k, 2]] ← 1
            for c = 0 to max_num_p do
                coef ← random(0, max_num_p)
                col ← random(1, N)
                if coef ≠ 0 & B[rows_Err[i], ·] + coef − B[rows_Err[i], col] ≤ max_num_p then
                    B[rows_Err[i], col] ← coef
                else if coef == 0 & ⟨rows_Err[i], col⟩ ∉ BIJ then
                    B[rows_Err[i], col] ← coef
    return A, B
```

Algorithm 8 Checking of the linear independence between $\mathbf{A}[i, \cdot]$ and $\mathbf{B}[i, \cdot]$

```
function LINEAR_INDEPENDENCE_REACTION(A, B, rows_Err)
    err_LinDep ← [ ]
    for i = 1 to len(rows_Err) do
        if A[rows_Err[i], ·] ∧ B[rows_Err[i], ·] are linearly dependent then
            err_LinDep ← err_LinDep ⊙ rows_Err[i]
    return err_LinDep
```

Algorithm 9 Random initialization of the amounts of the species

```
function AMOUNTS_GEN(N, D_s, min_s, max_s, μ_s, σ_s)
    M_0[·] ← 0
    if dist is Uniform or Logarithmic then
        for j = 1 to N do
            M_0[j] ← random(D_s, min_s, max_s)
    else
        for j = 1 to N do
            M_0[j] ← random(D_s, min_s, max_s, μ_s, σ_s)
    return M_0
```

Algorithm 10 Random generation of kinetic constants of the reactions

```
function KINETIC_CONSTANTS_GEN(M, D_r, min_r, max_r, μ_r, σ_r)
    K[·] ← 0
    if dist is Uniform or Logarithmic then
        for i = 1 to M do
            K[i] ← random(D_r, min_r, max_r)
    else
        for i = 1 to M do
            K[i] ← random(D_r, min_r, max_r, μ_r, σ_r)
    return K
```

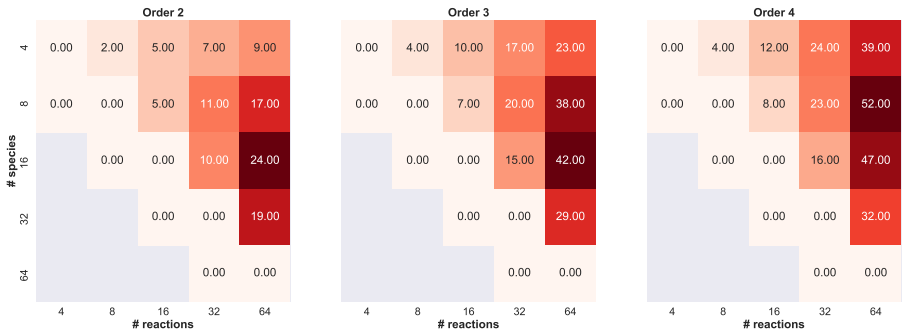

Figure A1. Heatmaps showing the median value of the network deficiency calculated on RBMs generated by varying the number of species and reactions, and the maximum number of reactants and products.

Appendix B. Analysis of the properties of the RBMs generated with SMGen

We analyzed the properties of the RBMs generated with SMGen by considering the network deficiency [53], the scale-freeness of the generated network [54], and the number of zero complexes, i.e., the null species used to denote degradation reactions or the influx of chemicals.

Network deficiency is an important structural attribute of a reaction network; indeed, this measure gives an indication of the independence of the reaction vectors. Specifically, the network deficiency is a non-negative integer index calculated considering three coefficients: the number distinct complexes n, the number of linkage classes l, and the rank of the network r. In this notation, the complexes of a network are the entities that appear at the heads and tails of each reaction (examples of distinct complexes are A, $2A$, $A + B$, etc.), while the linkage classes are the sets of complexes in the various "parts" composing the network, where complexes in a set are linked to each other, directly or indirectly, but they are not linked to any other complex in the network. Finally, the rank of the network is exactly the rank of its set of reaction vectors. To be more precise, the network has a rank r if there is a subset containing r linearly independent reaction vectors but there is not any subset containing $r + 1$ linearly independent reaction vectors. The network deficiency is equal to $n - l - r$. It is worth noting that networks having a deficiency equal to zero are the ones where the reaction vectors are as independent as the partition of complexes into linkage classes will allow.

To perform such analysis, as a first step, we generated symmetric and asymmetric RBMs characterized by a number of species N and reactions M ranging from 4 to 64, by also varying max_{num_r} and max_{num_p} into $\{2, 3, 4\}$. Specifically, for each condition (i.e., N, M, max_{num_r} and max_{num_p}), 100 RBMs were randomly created to collect statistically sound results. We limited this analysis to the models' size that can be generated with a

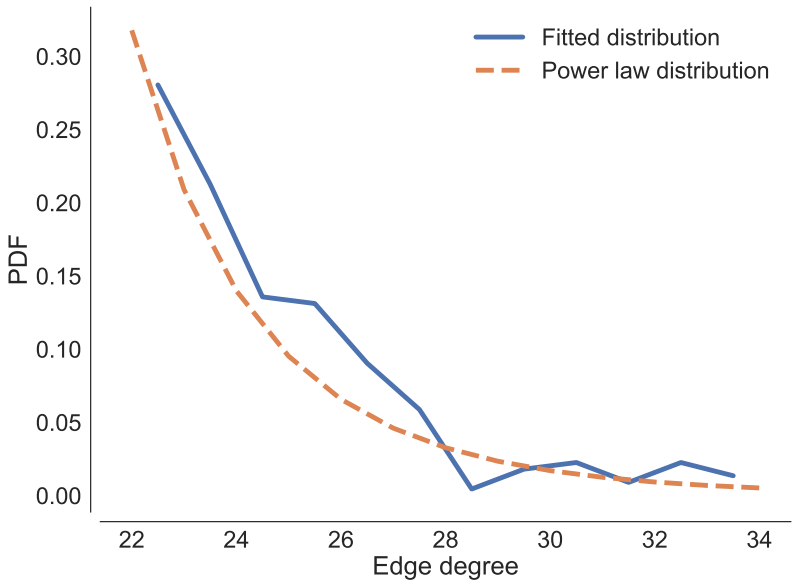

Figure A2. Degree distribution of an interaction network induced by an RBM randomly generated using SMGen. The discrete fitted distribution is shown using linearly spaced bins.

single connected component (independently from the maximum number of reactants and products); i.e., with the following constraint:

$$(max_{num_r} + max_{num_p}) \times M > N.$$

Then, we calculated the network deficiency for each generated RBM; finally, for each condition, we calculated the median of the network deficiency. The heatmaps depicted in Figure A1 clearly show that the network deficiency increases when $M > N$. Moreover, the maximum number of reactants and products affects the network deficiency: indeed, the higher the maximum number of reactants and products, the higher the network deficiency. The achieved results are coherent with the definition of network deficiency itself. As a matter of fact, the number of possible distinct complexes and the number of possible distinct linkage classes increase along with N and the maximum number of reactants and products, while the number of linearly independent reaction vectors is bounded by $\min(M, N)$. Finally, the results obtained from this analysis present values congruent with those of real biochemical models, such as the Brussellator [38,43]) (5 species and 4 reactions) with deficiency $= 0$, the prokaryotes gene expression model [55] (5 species and 8 reactions) with deficiency $= 1$, the heat shock response in eukaryotes model (10 species and 17 reactions) with deficiency $= 5$, the Ras/cAMP/PKA model [22] (33 species and 39 reactions) with deficiency $= 7$, and the human intracellular metabolic pathway in red blood cells [56] (114 species and 226 reactions) with deficiency $= 13$.

Biological networks, in particular metabolic networks [57], often show scale-freeness properties [54]. In order to prove that SMGen produces networks that are characterized by this property, we converted the bipartite graph defined by the stoichiometric matrices to a common interaction graph, where nodes represent the chemical species and edges represent reactions involving the nodes as reactants or products. Then, we analyzed the structural properties of the generated network and, in particular, we investigated whether the degree distribution follows a power law distribution, which is characteristic of scale-free graphs. Since the scale-free properties emerge only for large scale networks, we generated and analyzed the interaction network induced by an RBM with by $10,000$ chemicals species involved in $10,000$ reactions. According to our results, the degree distribution seems to fit well with a power law (see Figure A2), confirming that SMGen

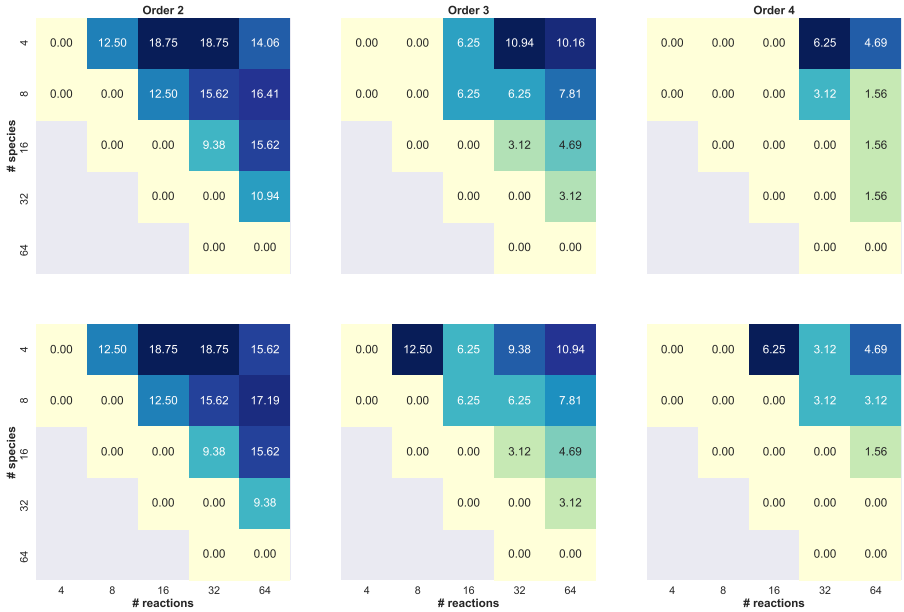

Figure A3. Heatmaps showing the median percentage of reactions where the zero complex appears as reactant (top) or product (bottom), calculated on RBMs generated by varying the number of species and reactions, and the maximum number of reactants and products

can generate scale-free networks. It is worth noting that curated models of this size are seldom published, because both the initial amounts of the chemical species and the kinetic reaction rates would be difficult to collect, and in any case they would represent a challenge for biochemical simulation.

Finally, in order to evaluate the abundance of the zero complex appearing as reactant or product in a reaction, we used the same RBMs generated to calculate the network deficiency. As a first step, for each RBM, we calculated the percentage of reactions where the zero complex appears either as a reactant or as a product. The heatmaps reported in Figure A3 clearly show that this value increases when $M >> N$. It is worth noting that the higher the maximum number of reactants and products for each reaction, the lower the percentage of reactions involving the zero complex. Overall, we observe that at most 2 out of 10 species are involved in one of such reactions. We investigated the abundance of the zero complex in real models. In the prokaryotes gene expression model [55], the zero complex appears in 25% of the reactions, while it is present in 5.88% of the reactions composing the heat shock response model [58]. Thus, the values calculated on synthetic RBMs generated by SMGen are in agreement with those obtained by analyzing models representing real biochemical networks.

References

1. Aldridge, B.B.; Burke, J.M.; Lauffenburger, D.A.; Sorger, P.K. Physicochemical modelling of cell signalling pathways. *Nat. Cell Biol.* **2006**, *8*, 1195–1203.
2. Szallasi, Z.; Stelling, J.; Periwal, V. *System Modeling in Cellular Biology: From Concepts to Nuts and Bolts*; The MIT Press, 2006.
3. Nobile, M.S.; Tangherloni, A.; Rundo, L.; Spolaor, S.; Besozzi, D.; Mauri, G.; Cazzaniga, P. Computational Intelligence for Parameter Estimation of Biochemical Systems. Proc. Congress on Evolutionary Computation (CEC). IEEE, 2018, pp. 1–8.
4. Munsky, B.; Hlavacek, W.S.; Tsimring, L.S. *Quantitative biology: theory, computational methods, and models*; MIT Press, 2018.
5. Tangherloni, A.; Spolaor, S.; Cazzaniga, P.; Besozzi, D.; Rundo, L.; Mauri, G.; Nobile, M.S. Biochemical parameter estimation vs. benchmark functions: a comparative study of optimization performance and representation design. *Appl. Soft Comput.* **2019**, *81*, 105494.

6. Kitano, H. Systems biology: a brief overview. *Science* **2002**, *295*, 1662–1664.
7. Chou, I.C.; Voit, E.O. Recent developments in parameter estimation and structure identification of biochemical and genomic systems. *Math. Biosci.* **2009**, *219*, 57–83.
8. Somogyi, E.T.; Bouteiller, J.M.; Glazier, J.A.; König, M.; Medley, J.K.; Swat, M.H.; Sauro, H.M. libRoadRunner: a high performance SBML simulation and analysis library. *Bioinformatics* **2015**, *31*, 3315–3321.
9. Hoops, S.; Sahle, S.; Gauges, R.; Lee, C.; Pahle, J.; Simus, N.; Singhal, M.; Xu, L.; Mendes, P.; Kummer, U. COPASI—a complex pathway simulator. *Bioinformatics* **2006**, *22*, 3067–3074.
10. Moraru, I.I.; Schaff, J.C.; Slepchenko, B.M.; Blinov, M.; Morgan, F.; Lakshminarayana, A.; Gao, F.; Li, Y.; Loew, L.M. Virtual Cell modelling and simulation software environment. *IET Syst. Biol.* **2008**, *2*, 352–362.
11. Li, H.; Petzold, L. Efficient parallelization of the stochastic simulation algorithm for chemically reacting systems on the graphics processing unit. *Int. J. High Perform. Comput. Appl.* **2010**, *24*, 107–116.
12. Zhou, Y.; Liepe, J.; Sheng, X.; Stumpf, M.P.; Barnes, C. GPU accelerated biochemical network simulation. *Bioinformatics* **2011**, *27*, 874–876.
13. Komarov, I.; D'Souza, R.M.; Tapia, J. Accelerating the Gillespie τ-leaping method using graphics processing units. *PLoS ONE* **2012**, *7*, e37370.
14. Komarov, I.; D'Souza, R.M. Accelerating the Gillespie exact stochastic simulation algorithm using hybrid parallel execution on graphics processing units. *PLoS ONE* **2012**, *7*, e46693.
15. Nobile, M.S.; Cazzaniga, P.; Besozzi, D.; Mauri, G. GPU-accelerated simulations of mass-action kinetics models with cupSODA. *J. Supercomput.* **2014**, *69*, 17–24.
16. Nobile, M.S.; Cazzaniga, P.; Besozzi, D.; Pescini, D.; Mauri, G. cuTauLeaping: A GPU-powered tau-leaping stochastic simulator for massive parallel analyses of biological systems. *PLoS ONE* **2014**, *9*, e91963.
17. Sumiyoshi, K.; Hirata, K.; Hiroi, N.; Funahashi, A. Acceleration of discrete stochastic biochemical simulation using GPGPU. *Front. Physiol.* **2015**, *6*, 42.
18. Tangherloni, A.; Nobile, M.; Besozzi, D.; Mauri, G.; Cazzaniga, P. LASSIE: simulating large-scale models of biochemical systems on GPUs. *BMC Bioinform.* **2017**, *18*, 246.
19. Tangherloni, A.; Nobile, M.S.; Cazzaniga, P.; Capitoli, G.; Spolaor, S.; Rundo, L.; Mauri, G.; Besozzi, D. FiCoS: a fine-and coarse-grained GPU-powered deterministic simulator for biochemical networks. *bioRxiv* **2021**. Accepted by PLoS Computational Biology.
20. Glont, M.; Nguyen, T.V.N.; Graesslin, M.; Hälke, R.; Ali, R.; Schramm, J.; Wimalaratne, S.M.; Kothamachu, V.B.; Rodriguez, N.; Swat, M.J.; Eils, J.; Eils, R.; Laibe, C.; Malik-Sheriff, R.S.; Chelliah, V.; Le Novère, N.; Hermjakob, H. BioModels: expanding horizons to include more modelling approaches and formats. *Nucl. Acids Res.* **2018**, *46*, D1248–D1253. doi:10.1093/nar/gkx1023.
21. Malik-Sheriff, R.S.; Glont, M.; Nguyen, T.V.N.; Tiwari, K.; Roberts, M.G.; Xavier, A.; Vu, M.T.; Men, J.; Maire, M.; Kananathan, S.; Fairbanks, E.L.; Meyer, J.P.; Arankalle, C.; Varusai, T.M.; Knight-Schrijver, V.; Li, L.; Dueñas-Roca, C.; Dass, G.; Keating, S.M.; Park, Y.M.; Buso, N.; Rodriguez, N.; Hucka, M.; Hermjakob, H. BioModels — 15 years of sharing computational models in life science. *Nucl. Acids Res.* **2020**, *48*, D407–D415. doi:10.1093/nar/gkz1055.
22. Besozzi, D.; Cazzaniga, P.; Pescini, D.; Mauri, G.; Colombo, S.; Martegani, E. The role of feedback control mechanisms on the establishment of oscillatory regimes in the Ras/cAMP/PKA pathway in *S. cerevisiae*. *EURASIP J. Bioinform. Syst. Biol.* **2012**, *2012*.
23. Cazzaniga, P.; Nobile, M.S.; Besozzi, D.; Bellini, M.; Mauri, G. Massive exploration of perturbed conditions of the blood coagulation cascade through GPU parallelization. *BioMed Res. Int.* **2014**, *2014*. Article ID 863298.
24. Pescini, D.; Cazzaniga, P.; Besozzi, D.; Mauri, G.; Amigoni, L.; Colombo, S.; Martegani, E. Simulation of the Ras/cAMP/PKA pathway in budding yeast highlights the establishment of stable oscillatory states. *Biotechnol. Adv.* **2012**, *30*, 99–107.
25. Renz, A.; Widerspick, L.; Dräger, A. Genome-Scale Metabolic Model of Infection with SARS-CoV-2 Mutants Confirms Guanylate Kinase as Robust Potential Antiviral Target. *Genes* **2021**, *12*, 796.
26. Gao, C.W.; Allen, J.W.; Green, W.H.; West, R.H. Reaction Mechanism Generator: Automatic construction of chemical kinetic mechanisms. *Comput. Phys. Commun.* **2016**, *203*, 212–225.
27. Khanshan, F.S.; West, R.H. Developing detailed kinetic models of syngas production from bio-oil gasification using Reaction Mechanism Generator (RMG). *Fuel* **2016**, *163*, 25–33.
28. Lok, L.; Brent, R. Automatic generation of cellular reaction networks with Moleculizer 1.0. *Nature Biotechnol.* **2005**, *23*, 131–136.
29. Städter, P.; Schälte, Y.; Schmiester, L.; Hasenauer, J.; Stapor, P.L. Benchmarking of numerical integration methods for ODE models of biological systems. *Sci. Rep.* **2021**, *11*, 1–11.
30. Garrido, A. Symmetry in complex networks. *Symmetry* **2011**, *3*, 1–15.
31. Ohlsson, F.; Borgqvist, J.; Cvijovic, M. Symmetry structures in dynamic models of biochemical systems. *J. R. Soc. Interface* **2020**, *17*, 20200204.
32. Keating, S.M.; Waltemath, D.; König, M.; Zhang, F.; Dräger, A.; Chaouiya, C.; Bergmann, F.T.; Finney, A.; Gillespie, C.S.; Helikar, T.; others. SBML Level 3: an extensible format for the exchange and reuse of biological models. *Mol. Syst. Biol.* **2020**, *16*, e9110.
33. Besozzi, D.; Cazzaniga, P.; Mauri, G.; Pescini, D. BioSimWare: a software for the modeling, simulation and analysis of biological systems. International Conference on Membrane Computing. Springer, 2010, pp. 119–143.
34. Chellaboina, V.; Bhat, S.P.; Haddad, W.M.; Bernstein, D.S. Modeling and analysis of mass-action kinetics. *IEEE Control Syst.* **2009**, *29*, 60–78.

35. Voit, E.O.; Martens, H.A.; W., O.S. 150 years of the mass action law. *PLoS Comput Biol* **2015**, *11*, e1004012.
36. Nelson, D.L.; Lehninger, A.L.; Cox, M.M. *Lehninger principles of biochemistry*, 5 ed.; Macmillan: London, UK, 2008.
37. Chaouiya, C.; Remy, E.; Thieffry, D. Petri net modelling of biological regulatory networks. *J. Discr. Alg.* **2008**, *6*, 165–177.
38. Davidrajuh, R. Detecting Existence of Cycles in Petri Nets. Proc. of International Joint Conference SOCO'16-CISIS'16-ICEUTE'16. Springer, 2016, pp. 376–385.
39. Tangherloni, A.; Spolaor, S.; Rundo, L.; Nobile, M.S.; Cazzaniga, P.; Mauri, G.; Liò, P.; Merelli, I.; Besozzi, D. GenHap: a novel computational method based on genetic algorithms for haplotype assembly. *BMC Bioinform.* **2019**, *20*, 172.
40. Cho, Y.J.; Ramakrishnan, N.; Cao, Y. Reconstructing chemical reaction networks: data mining meets system identification. Proc. 14th ACM International Conference on Knowledge Discovery and Data Mining. ACM, 2008, pp. 142–150.
41. Dalcín, L.; Paz, R.; Storti, M. MPI for Python. *J. Parallel Distrib. Comput.* **2005**, *65*, 1108–1115.
42. Gropp, W.D.; Gropp, W.; Lusk, E.; Skjellum, A. *Using MPI: Portable Parallel Programming With the Message-Passing Interface*; Vol. 1, MIT press: Cambridge, MA, USA, 1999.
43. Vaidyanathan, S. Dynamics and control of Brusselator chemical reaction. *Int. J. ChemTech. Res.* **2015**, *8*, 740–749.
44. Cornish-Bowden, A. One hundred years of Michaelis–Menten kinetics. *Perspect. Sci.* **2015**, *4*, 3–9.
45. Hill, A.V. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *J. Physiol.* **1910**, *40*, 4–7.
46. Dimitrova, T.; Petrovski, K.; Kocarev, L. Graphlets in multiplex networks. *Sci. Rep.* **2020**, *10*, 1–13.
47. Doria-Belenguer, S.; Youssef, M.K.; Böttcher, R.; Malod-Dognin, N.; Pržulj, N. Probabilistic graphlets capture biological function in probabilistic molecular networks. *Bioinformatics* **2020**, *36*, i804–i812.
48. Pržulj, N. Biological network comparison using graphlet degree distribution. *Bioinformatics* **2007**, *23*, e177–e183.
49. Wong, E.; Baur, B.; Quader, S.; Huang, C.H. Biological network motif detection: principles and practice. *Brief. Bioinform.* **2012**, *13*, 202–215.
50. Stone, L.; Simberloff, D.; Artzy-Randrup, Y. Network motifs and their origins. *PLoS Comput. Biol.* **2019**, *15*, e1006749.
51. Kim, W.; Li, M.; Wang, J.; Pan, Y. Biological network motif detection and evaluation. *BMC Syst. Biol.* **2011**, *5*, 1–13.
52. Prill, R.J.; Iglesias, P.A.; Levchenko, A. Dynamic properties of network motifs contribute to biological network organization. *PLoS Biol.* **2005**, *3*, e343.
53. Feinberg, M. *Foundations of chemical reaction network theory*; Springer, 2019.
54. Arita, M. Scale-freeness and biological networks. *J. Biochem.* **2005**, *138*, 1–4.
55. Wang, Y.; Christley, S.; Mjolsness, E.; Xie, X. Parameter inference for discretely observed stochastic kinetic models using stochastic gradient descent. *BMC Syst. Biol.* **2010**, *4*, 99.
56. Totis, N.; Tangherloni, A.; Beccuti, M.; Cazzaniga, P.; Nobile, M.S.; Besozzi, D.; Pennisi, M.; Pappalardo, F. Efficient and settings-free calibration of detailed kinetic metabolic models with enzyme isoforms characterization. Proc. of International Meeting on Computational Intelligence Methods for Bioinformatics and Biostatistics. Springer, 2018, pp. 187–202.
57. Ravasz, E.; Somera, A.L.; Mongru, D.A.; Oltvai, Z.N.; Barabási, A.L. Hierarchical organization of modularity in metabolic networks. *Science* **2002**, *297*, 1551–1555.
58. Kline, M.P.; Morimoto, R.I. Repression of the heat shock factor 1 transcriptional activation domain is modulated by constitutive phosphorylation. *Mol. Cell. Biol.* **1997**, *17*, 2107–2115.