Chapter 10
From Artificial Chemistries to Systems Biology: Software for Chemical Organization Theory

Christoph Kaleta

Artificial Chemistries abstract from real-world chemistries by reducing them to systems of interacting and reacting molecules. They have been used to study phenomena in a wide array of fields like social and ecological modelling, evolution or chemical computing. Artificial Chemistries are inherently difficult to study and, thus, methods have been proposed to analyze their complexity. This chapter outlines how the concept of chemical organization and software dedicated at their analysis can help to ease this task. The chemical organizations of a reaction network correspond to sets of molecules that can coexist over long periods of (simulation-) time. Thus, they can be used to study the full dynamic behavior a system can exhibit without the need to simulate it in every detail. Due to this appealing property, Chemical Organization Theory has been used in the study of a wide array of systems ranging from Artificial Chemistries to real-world chemistries and biological systems. Especially the analysis of biological systems motivated an integration of the tools dedicated to the study of chemical organizations into an application framework from Systems Biology. The benefit of this integration is that tools from Systems Biology can be used without much effort along with the tools for the computation of chemical organizations and vice versa. Thus, software for the analysis of chemical organizations seamlessly integrates into a framework covering almost any aspect of network design and analysis.

10.1 Introduction

The quest of understanding the emergence and development of life on Earth has a long-standing history in research. While the concept of Darwinian Evolution readily explains the appearance of more and more complex organisms through inheritance, mutation and selection, it is still not fully understood how the first organism for selection to act upon came to life. Artificial Life

and, in particular, Artificial Chemistry are trying to shed light on this question by taking an approach that is in some parts opposed to the one taken by biology. These fields try to understand real living systems not by considering them in mechanistic detail, but from an abstract point of view. Thus, they study how concepts like self-replication or information processing and storage might have evolved. An abstraction borrowed from real chemistries that is common to all systems in Artificial Chemistries is that of colliding molecules. Thus, [9] defined an Artificial Chemistry as "[..] a man-made system which is similar to a real chemical system." Formally, an Artificial Chemistry is defined by the triple (S, R, A), with S denoting the set of possible species, R denoting the set of collision rules among the species, commonly also referred to as reactions, and A denoting an algorithm describing the reaction vessel and how the collision rules are applied to the species [9]. An example for a simple Artificial Chemistry can be found in Fig. 10.1.

Artificial Chemistries have seen a wide array of applications. Those applications can be aligned on three broad axes: modelling, information processing and optimization [9]. In the first direction they have been used to create and analyze models of ecological and social systems as well as evolution. Prominent fields of Artificial Chemistry in information processing are, for example, chemical computing or DNA computing. Due to their ability to create evolutionary behavior, they have also been used for optimization in the context of evolutionary algorithms. More details on some of these applications in connection to Chemical Organization Theory are outlined in Section 10.3.3. For a comprehensive overview on the field of Artificial Chemistry the reader is referred to [9].

Due to the inherent complexity of systems in Artificial Chemistry, methods to analyze their behavior are needed. One of these is Chemical Organization Theory [8], which allows to study the dynamic behavior of chemical reaction networks, a particular type of Artificial Chemistry. In contrast to concepts that explicitly take into account space (e.g., dissipative particle dynamics, Chapter 11), a chemical reaction network in the context of Chemical Organization Theory is best visualized as a well-stirred reaction tank containing a finite set of species reacting with each other through a finite set of reactions. Thus, the concentrations of the species are assumed to be homogeneous. In consequence, the system can be described by a set of species, reactions among the species and the underlying dynamics of the reactions alone (see Fig. 10.1 for an example). Chemical Organization Theory focuses on the structure of a reaction network (i.e., the set of species and reactions). From this structure it tries to infer constraints on the dynamics of the reaction network. The set of chemical organizations of a reaction network consists of sets of species that are likely to coexist over a long period of simulation-time of the dynamics. In consequence, the dynamics can be mapped to a movement in the set of organizations [8].

The relationship between the structure and the dynamics of a reaction network can be derived by imposing two conditions on a set of species to be

an organization. Both serve to guarantee the stability of this set. First, the species set is required to be *closed*. This condition is fulfilled if there is no reaction in the network that could produce a species not within the set. Thus, the dynamic of a reaction network will be restricted to the set of species in an organization if they were present at the beginning of the simulation (with some species possibly vanishing). The second condition, *self-maintenance*, ensures that there exists a flux through the reactions of the network (i.e., an instance of the dynamics, such that no species vanishes). Even though the existence of such a flux does not guarantee that the species of an organization will persist during a simulation, it can be shown in the other direction that if such a species set is encountered during a simulation, it corresponds to an organization. A formal definition of these concepts is given in Section 10.3.1.

Chemical Organization Theory has seen a wide array of applications in both Artificial Chemistry and Systems Biology. In Artificial Chemistry it has been used to study systems for chemical computing [23, 25], to serve as a design principle in chemical computing [7] and to study chemical evolution [24]. In Systems Biology it has been used to study models of HIV infection [22], the diauxic shift in *Escherichia coli* [4, 18], as well as the properties of a genome-scale network of the same organism [5] and effects of gene knock-outs [18]. Some of these applications are presented in more detail in Section 10.3.3.

Due to a close relationship of the concept of a chemical reaction network to the understanding of a metabolic network in Systems Biology, many tools from this field can be applied to the study of such networks. In the other direction, Chemical Organization Theory can be used to complement methods developed in Systems Biology. Thus, the tools for the study of chemical organizations were developed with the aim to allow an easy integration into a whole array of existing tools from Systems Biology.

This integration is achieved on two different levels. First, chemical reaction networks are represented in the *Systems Biology Markup Language* (SBML) format [16]. Even though species and reactions can be stored in a simple text-file, different kinds of analysis of such networks focus on different aspects. Thus, a common representation serves, on the one hand, to circumvent the cumbersome way of interconvert reaction networks between different data formats. On the other hand, an integrated representation of all aspects that might be of interest is achieved. One analysis might focus, for example, on the set of reactions and species, while another analysis necessitates information on the kinetics of the reactions. If a reaction network has been designed with the former kind of analysis in mind, the second kind can be applied by "just" adding information about the kinetics to the already available set of reactions.

Second, an integration on the level of applications is achieved. Using SBML as a representation already makes available a whole array of tools that use this standard (see [32] for a comprehensive list of applications using SBML). Furthermore, the *Systems Biology Workbench* (SBW) [31] offers an even deeper integration of applications that not only encompasses the exchange of reaction networks between two applications but also the possibility to call specific

functions offered by each application that interface SBW. Thus, one application can be used to design a reaction network. If one wants to study the dynamic of this network under a particular condition for a certain time span, a function from another application offering this *service* through SBW can be called. Thus, it is, on the one hand, not necessary to develop such a tool anew or integrate the corresponding code into the own project. On the other hand, the cumbersome process of saving the network, initializing the simulation tool, loading the network again and specifying the parameters of the simulation can be circumvented and all functionality is available from the initial application. A more comprehensive overview on SBML and SBW is given in Section 10.2.

This chapter is organized as follows. A short overview on the tools from Systems Biology that can be used with the application for the analysis of chemical organizations is outlined in Section 10.2. A comprehensive introduction into the concept of chemical organizations is given in Section 10.3.1. Details on the algorithms for the computation of chemical organizations and applications of Chemical Organization Theory in the field of Artificial Chemistry are presented in Sections 10.3.2 and 10.3.3. Finally, applications for the computation and analysis of chemical organizations in reaction networks are outlined in Section 10.4.

10.2 Using Tools from Systems Biology: SBML and SBW

Even though Chemical Organization Theory has been introduced in the field of Artificial Chemistries, it has seen a wide array of applications in Systems Biology. Consequently, it makes use of some tools commonly used in this field of research.

10.2.1 Systems Biology Markup Language

Reaction networks can be stored using simple text files. However, this simplicity can result in every research group using proprietary formats that are most appropriate to their needs. In consequence, it is difficult to interconvert reaction networks between different formats. A standardized representation of reaction networks that incorporates the most commonly used aspects of their representation allows a simpler interchange of models and reduces mistakes made during interconversion.

The Systems Biology Markup Language (SBML) [15, 16] is proposed as a standard for the representation of reaction networks in the field of Systems Biology. Through its growing acceptance, SBML is extended to meet

the needs of its users. In its current version, *level 2*, it allows the definition of compartments, species, reactions, units, rules and events. Through the definition of compartments and relations between the compartments, some basic spatial differentiation seen in living systems can be integrated into a reaction network. Reactions can be described by rate laws. Since these laws are formulated using *MathML*, arbitrary functions can be described. Using rules, concentrations and parameters can be constrained. Events offer the possibility to perturbate networks at given time-points by, for example, setting the concentration of a species or changing a specific parameter. Additionally it is possible to include user-defined tags that enable, for example, the display of information in HTML format or the incorporation of positional information (e.g., for the display in network editors). Since SBML is used by a growing array of applications virtually covering any aspect of network analysis, it offers a sophisticated base that can also be useful in the field of Artificial Chemistry.

10.2.2 Systems Biology Workbench

Using the Systems Biology Workbench (SBW) [31] the unification of reaction networks based on SBML is extended to the level of applications. SBW allows the integration of some commonly used steps in network design and analysis. SBW basically represents a message passing architecture. Its central element is a broker to which applications, written in a wide variety of programming languages, can interface. Messages can be passed from one application to another allowing applications that interface to SBW to communicate between each other. The aim of SBW is to offer a framework that allows the reuse of software written in heterogeneous programming languages. Each tool interfacing to SBW offers a set of methods that can be called by other applications that interface to SBW. One of the most basic methods is the *doAnalysis()* method that is offered by most SBW-compliant applications. This method is called with a string representing a reaction network in SBML format. This application then loads the network, which can subsequently be examined with the analysis methods offered. In connection with SBW interfacing network editing and design tools like *JDesigner* [31] and *CellDesigner* [20], tools for deterministic and stochastic simulation like *Jarnac* [31] and *Dizzy* [30] as well as analysis tools like *METATOOL* [27] and *Bifurcation Discovery Tool* [6], the applications for the analysis of chemical organizations are integrated into one easily expandable framework.

10.3 Chemical Organization Theory

This section introduces formally the concept of chemical organizations and gives some details about the algorithms that can be used for their computation. This is followed by an overview on recent applications of this concept.

10.3.1 Background

Extending ideas by Fontana and Buss [12], the theory of chemical organizations [8] provides a new method to analyze complex reaction networks. It allows one to predict sets of species that are likely to coexist over long periods of simulation time (i.e, the potential phenotypes of the reaction network). A *reaction network* $\langle \mathcal{M}, \mathcal{R} \rangle$ is described by a set of species $\mathcal{M}$ and a set of reactions $\mathcal{R} \subseteq \mathcal{P}_M(\mathcal{M}) \times \mathcal{P}_M(\mathcal{M})$ among these species, with $\mathcal{P}_M(\mathcal{M})$ denoting the set of all multisets with elements from $\mathcal{M}$. Furthermore, each reaction network $\langle \mathcal{M}, \mathcal{R} \rangle$ implies an $m \times n$ stoichiometric matrix $\mathbf{S} = (s_{i,j})$, where $s_{i,j}$ is the number of molecules of species $i \in \mathcal{M}$ that is produced in reaction $j \in \mathcal{R}$ (i.e., right-hand side minus left-hand side). An example for a reaction network and the associated stoichiometric matrix is given in Fig. 10.1.

To be an organization, a species set has to fulfill two criteria: closure and self-maintenance. The closure condition ensures that no new species appear through the reactions among the species of an organization. We call a set of species M *closed* if there exists no reaction $r = M_1 \times M_2$ with $M_1 \subseteq M$ and $M_2 \nsubseteq M$.

The self-maintenance condition ensures that there exists a flux through the reactions, such that no species in M will vanish over time. The dynamics of the reaction network can be written using the ordinary differential equation

$$\dot{\mathbf{x}}(t) = \mathbf{S}\mathbf{v}(\mathbf{x}(t)),$$

with $\mathbf{x}(t)$ being the concentration values of the species in $\mathcal{M}$ at time t and $\mathbf{v}(\mathbf{x}(t))$ being the flux through the reactions as a function of $\mathbf{x}(t)$. The self-maintenance condition requires the existence of a flux vector $\mathbf{v} \in \mathbb{R}^n$ containing positive entries for all reactions among the species in M and zero entries for the remainder such that the concentration change given as $\mathbf{Sv}$ is nonnegative. Formally, we call a set M *self-maintaining* if:

1. For every reaction $(\mathcal{A} \to \mathcal{B}) \in \mathcal{R}$ with $\mathcal{A} \in \mathcal{P}_M(M)$, its corresponding flux is $\mathbf{v}_{\mathcal{A}\to\mathcal{B}} > 0$.
2. For every reaction $(\mathcal{A} \to \mathcal{B}) \in \mathcal{R}$ with $\mathcal{A} \notin \mathcal{P}_M(M)$, its corresponding flux is $\mathbf{v}_{\mathcal{A}\to\mathcal{B}} = 0$.
3. For every species $i \in M$, its concentration change is nonnegative: $(\mathbf{Sv})_i \geq 0$.

Please note that this condition also includes the steady state condition that is a commonly used constraint in the stoichiometric analysis of metabolic networks like flux balance analysis [34] or elementary mode analysis [33].

$$
\begin{aligned}
\mathcal{R} = \{ & \\
& 1\,a \rightarrow 2\,a \\
1\,a + 1\,b & \rightarrow 1\,a + 2\,b \\
1\,d + 1\,c & \rightarrow 1\,a + 2\,c \\
& 1\,d \rightarrow 2\,d \\
& 1\,a \rightarrow \\
& 1\,b \rightarrow \\
& 1\,c \rightarrow \\
\} &
\end{aligned}
\qquad
\mathbf{S} = \begin{pmatrix} 1 & 0 & 1 & 0 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 1 & 0 & 0 & 0 \end{pmatrix}
$$

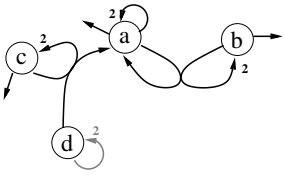

Fig. 10.1 Example for a reaction network. In the upper left, the set of reactions is given. **S** is the stoichiometric matrix of the system. Rows correspond to species and columns correspond to reactions in the same order as in the list of reactions. A schematic drawing of the reaction network is depicted in the lower half of the figure. The light gray reaction is added to the network for the computation of the second Hasse diagram of organizations in Fig. 10.2.

The set of organizations, together with the set inclusion $\subseteq$, forms a partially ordered set that can be visualized as a hierarchical structure in a Hasse diagram (see Figs. 10.2a and 10.2 for the Hasse diagram of the reaction network in Fig. 10.1). Linking to the dynamics, [8] has shown that all fixed-points of the system are instances of organizations, given that the dynamics of the system is described by a set of ordinary differential equations. This represents one of the most important aspects of chemical organization theory, since it allows one to infer statements about the possible qualitative states of a reaction network during simulation from its stoichiometry alone. Furthermore, a simulation of a network can be mapped to a movement in the Hasse diagram of organizations. For a simulation using ordinary differential equations, the floating-point concentrations need to be mapped to a set of present species. This is done by assuming only species present that have a concentrations above a predefined threshold Θ. This threshold could, for example, relate to a concentration equaling one molecule of a species. An example for the mapping of a simulation to the Hasse diagram of organizations can be found in Figs. 10.2b and 10.2c.

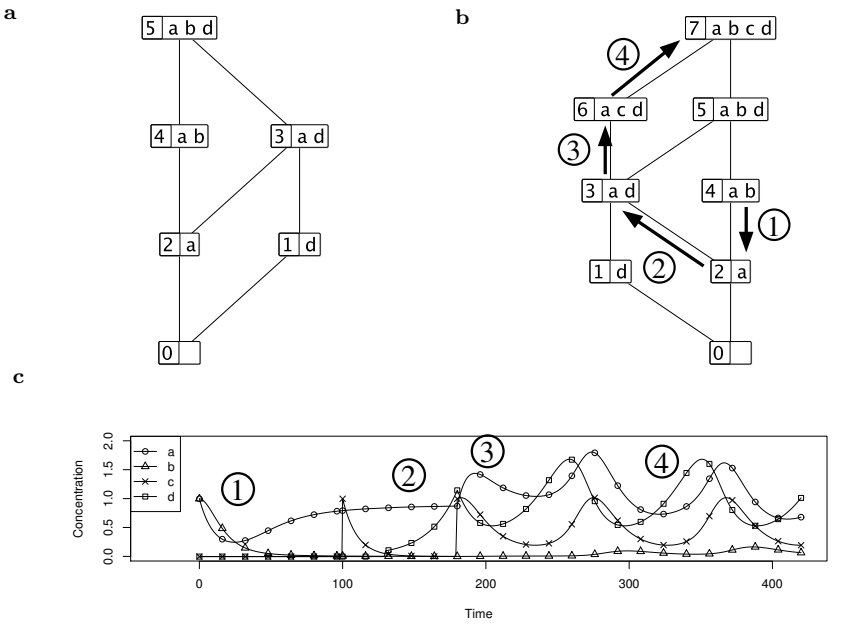

Fig. 10.2 Example for the movement in the Hasse diagram of organizations. Hasse diagrams of the chemical organizations of the reaction network in Fig. 10.1 **a** without and **b** with the self-replicating reaction for species d. **c** Example for the relation between the dynamics of a network and a movement in the Hasse diagram in part b. (1) At $t = 0$, the network is initialized with $[a] = 1$ and $[b] = 1$ corresponding to organization 4. Due to the kinetic laws, species b vanishes (i.e., its concentration falls below the predefined threshold Θ), corresponding to a downward movement into organization 2. There is no organization containing $\{a, c\}$; hence, $[c]$ raised to 1 at $t = 100$ rapidly vanishes (until $t = 130$). (2) At $t = 130$ $[d]$, is raised to 0.1 and rapidly accumulates, since the only reaction consuming d needs c as an additional educt, which is only present in very small amounts. Thus, the system moves to organization $\{a, d\}$, since still $[c] < \Theta$. (3) $[c]$ is raised to 1 to stop the accumulation of d. In consequence, the system moves upward into $\{a, c, d\}$. (4) Until $t = 180$, the decay of b was faster than its production through reaction 2. The peaks of the oscillations of $[a]$ starting in $t = 180$ allow a production of b greater than its decay. Thus, it accumulates, finally leading to an upward movement in the Hasse diagram when $[b] > \Theta$.

10.3.2 Algorithms for the Computation of Chemical Organizations

Two deterministic and one heuristic algorithm for the computation of chemical organizations have been developed [5]. While the deterministic algorithms can be applied to medium-scale networks of up to 100 reactions and molecules, the heuristic algorithm can find a subset of all organizations in larger networks. Thus, even networks with up to 1000 species and reactions have been analyzed [5].

Constructive Approach

The constructive approach for the computation of chemical organizations focuses on the Hasse diagram of organizations. The algorithm starts with the smallest organization of the network. This organization contains all species that can be generated by inflow reactions. If no inflow reaction is present, the smallest organization corresponds to the empty set. Starting from this initial organization, species sets containing this organization are being searched.

In order to find such organizations, the self-maintenance condition is relaxed to the semi-self-maintenance condition. This condition requires that if a species of an organization is consumed by a reaction that it is also produced by another reaction of the organization. It can be shown that each species set that fulfills the self-maintenance condition also fulfills the semi-self-maintenance condition. Given an initial organization, the algorithm iteratively tries to add species to the current organization. Then the closure of the species set is computed. Thus, for all reactions whose educts are present in the species set, the products are added if they are not yet present. This step is iterated until no further product can be added. Subsequently, the algorithm tries to fulfill the semi-self-maintenance condition by adding species such that producing reactions are available to species that are only consumed. Any possible combination of such production reactions are tried. By iteratively using the resulting species sets as initial species sets, all species sets fulfilling the closure and the semi-self-maintenance condition are being found. Using linear programming, each species set is subsequently tested for the self-maintenance condition. If this condition is fulfilled, the species set represents an organization.

An approach to accelerate the search for organizations can be to consider only connected organizations. In this approach, every species that is added to the network has to take part in a reaction that uses at least one species from the current species set. By computing only the connected organizations in the original and a modified version of the reaction network, all organizations can be found.

Flux Based Approach

The second approach for the computation of chemical organizations focuses on the self-maintenance condition. All flux vectors satisfying this condition lie in a convex polyhedral cone $\mathcal{P}$ in the n-dimensional flux space. This cone is defined by its (unique) set of n-dimensional spanning vectors or extreme rays. Thus, each flux vector representing a solution of the self-maintenance condition can be written as a positive linear combination of the extreme rays of $\mathcal{P}$. The inequalities defining $\mathcal{P}$ are given by

$$\mathbf{Sv} \geq 0 \text{ and } \mathbf{v} \geq 0.$$

Various algorithms for the transformation of the representation of $\mathcal{P}$ by inequalities into the set of extreme rays have been proposed. The most prominent is represented by the the double description method [13].

The extreme rays are further simplified by a transformation into the corresponding set of reactions (i.e., the set of nonzero indices). Since all flux vectors satisfying the self-maintenance condition are positive linear combinations of extreme rays, the algorithm subsequently searches for combinations of the extreme rays that satisfy the closure condition. The search for such combinations is split into three steps. In the first step, all elementary organizations of the system are computed. These are organizations that fulfill two properties. First, they are reactive organizations; that is, each species they contain participates in at least one reaction of the organization. Second, there exists no set of other organizations such that their union is equal to any of the elementary organizations. Consequently, each reactive organization can be written as the union of elementary organizations. In the second step, the elementary organizations are combined, yielding all reactive organizations. In a last step, it is iteratively tested for each organization whether there exist species that can be added without changing the set of available reactions. Thus, all organizations are found in the final step.

Heuristic

Empirical results show that the most time-consuming step in the flux-based approach is represented by the computation of the extreme rays of the convex polyhedral cone $\mathcal{P}$ representing the solution space to the self-maintenance condition. The number of extreme rays is growing exponentially with system size and, thus, the time needed for their computation in larger networks can exceed practical limits.

Since the extreme rays are not necessary for the second and third steps of the flux-based approach, a heuristic approximating the set of elementary organizations can be used. A central observation is that since the reactions of each organization fulfill the self-maintenance condition, also the union of two such reaction sets fulfills the self-maintenance condition (i.e., there exists a positive linear combination of both flux vectors satisfying condition (1) of the self-maintenance condition). Consequently, given a set of organizations, we just need to search for combinations of these organizations that fulfill the closure condition to find additional organizations. If this initial set contains all elementary organizations, the full set of organizations of the complete system can be found.

To find such an initial set of organizations, a random walk strategy can be used. This strategy starts with a randomly chosen species that is added to the current species set. Then randomly selected species that are substrate or product to at least one reaction also consuming or producing the last added species are iteratively added. After computing the closure of the final species

set, it is tested whether the self-maintenance condition is fulfilled. In this case, an organization has been found.

The random walk strategy is iterated several times, yielding an initial set of organizations. Subsequently, we can search for combinations of those organizations that fulfill the closure condition to find all organizations that can be generated from unions of the initially found organization set. As demonstrated in [5], this approach can be used to find all organizations in reaction networks that contain even 1000 species if the number of organizations is small.

10.3.3 Application of Chemical Organization Theory

Chemical Organization Theory has seen a wide array of applications. In the field of Artificial Chemistry it has been applied to the study of chemical programming and chemical or prebiotic evolution. Although first introduced in the study of Artificial Chemistries, it has also seen many applications in the study of biological networks. One of its appealing properties is the possibility to predict the dynamic behavior of a model under simulation. This is of great importance if the kinetics governing a reaction network are partly or totally unknown, impossible to tackle in current solvers or too complex to analyze their potential behavior.

Application in Chemical Computing

Using chemical reactions as a base for computation has first been proposed by Banâtre and Métayer [1]. Appealing properties of such systems are the massive parallelism were the solution appears as an emergent global behavior [23]. These concepts have been used in *GAMMA* [2], *CHAM* [21], *P-systems* [28] and *MGS* [14]. Such systems are difficult to analyze and, thus, Chemical Organization Theory has been used to tackle some of the problems in their analysis.

One example is the analysis of chemical reaction networks derived from boolean networks. Boolean networks can be transformed into reaction networks using a method described in [18] and [23]. This can be done by transforming boolean variables into two species each: one corresponding to the `true` state of the variable and one corresponding to the `false` state of the variable. System states encompassing both representations of a variable at once can be resolved by the addition of a reaction mimicking mutual destruction (i.e., a reaction taking both reactions as educts and producing nothing).

An example of the transformation of the universal *Nand* operator, presented in [23], is outlined in Fig. 10.3. If the input of the boolean network is simulated by inflow reactions with each input variable being set (i.e., ei-

ther the species presenting the `true` or the `false` state is supplied), the only chemical organizations of the resulting network contains a species corresponding to the output of the *Nand* operator. Since every boolean function can be constructed from the *Nand* operator, arbitrary boolean networks can be transformed into reaction networks using this approach.

The transformation of boolean networks into reaction networks can also be used to integrate a regulatory network modelled by a boolean formalism into a metabolic network. Doing this in a model of the central metabolism of *Escherichia coli* [18] showed that the chemical organizations of the resulting system corresponded to the growth on several available carbon sources. Additionally, the results of the knock-out of genes could be correctly predicted in most cases. The knock-outs can be simulated by removing the corresponding reactions from the network. The lethality of a knock-out can be assessed through the existence of no organization containing all species that are necessary for the organism to survive.

Design Principles in Chemical Computing

On a more abstract level, Chemical Organization Theory can be used as a guideline in the construction of systems devised for chemical computation [7]. The construction of such a system can be separated into two stages. In the first stage, only the set of species and reactions is constructed. The kinetics governing the reactions are devised in the second step. In the first step, Chemical Organization Theory is iteratively applied to analyze whether the constructed model exhibits the desired behavior based on the set of chemical organizations it contains. After this stage has been completed, kinetic laws can be assigned to the reactions. As argued in [7], such a design principle leads to more robust models. This is due to the fact that chemical organizations confer a structural stability to the system that is independent from the kinetic laws governing the reactions of the system.

A promising prospect of these design principles lies in the emerging field of Synthetic Biology [10]. The research in this area focuses on the combination of components like gene regulatory elements or proteins from different species to construct systems that exhibit a qualitatively new behavior. In [29], genes were used from several species to construct a recombinant strain of *Saccharomyces cerivisae* that could detect trace amounts of explosives. As outlined in [11], principles from engineering can be used to facilitate such developments organized in several layers. One of these layers includes the design of new systems. Since biological systems are often represented as reaction networks, a promising aspect of Chemical Organization Theory is the application in the construction of systems in Synthetic Biology.

a

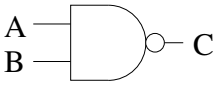

b

A	B	C
false	false	true
false	true	true
true	false	true
true	true	false

c

$$
\begin{aligned}
&\mathcal{R} = \{ \\
&1\,\overline{A} + 1\,\overline{B} \rightarrow 1\,C \quad (1)\\
&1\,\overline{A} + 1\,B \rightarrow 1\,C \quad (2)\\
&1\,A + 1\,\overline{B} \rightarrow 1\,C \quad (3)\\
&1\,A + 1\,B \rightarrow 1\,\overline{C} \quad (4)\\
&1\,A + 1\,\overline{A} \rightarrow \emptyset \quad (5)\\
&1\,B + 1\,\overline{B} \rightarrow \emptyset \quad (6)\\
&1\,C + 1\,\overline{C} \rightarrow \emptyset \quad (7)\\
&1\,C \rightarrow \emptyset \quad (8)\\
&1\,\overline{C} \rightarrow \emptyset \quad (9)\\
&\}
\end{aligned}
$$

d

$\mathcal{R} \cup \{\emptyset \rightarrow 1\,\overline{A} + 1\,\overline{B}\}$	$\{\overline{A},\ \overline{B},\ C\}$
$\mathcal{R} \cup \{\emptyset \rightarrow 1\,\overline{A} + 1\,B\}$	$\{\overline{A},\ B,\ C\}$
$\mathcal{R} \cup \{\emptyset \rightarrow 1\,A + 1\,\overline{B}\}$	$\{A,\ \overline{B},\ C\}$
$\mathcal{R} \cup \{\emptyset \rightarrow 1\,A + 1\,B\}$	$\{A,\ B,\ \overline{C}\}$

e

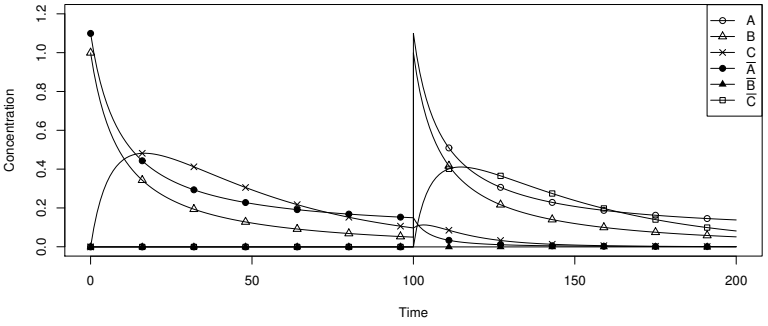

Fig. 10.3 Example for the conversion of a boolean *Nand* function into a reaction network. **a** *Nand* function with input variables A and B and output variable C. **b** Logic table of the *Nand* function. For every possible input of A and B, the output in C is given. **c** Reaction network modelling the *Nand* function. Plain symbol denotes `true` state of a variable and overlined symbol the `false` state. Reactions 5 to 7 are used to avoid inconsistent system states where the two species representing the state of a variable are present at the same time. Reactions 8 and 9 are used to reset the "output" of the network. **d** For every possible state of the input variables, the only organization of the system is given. It contains the input variables as well as a form of the output variable C. **e** Numerical simulation of the reaction network using mass-action kinetics. At $t = 0$, the network is initialized with $[\overline{A}] = 1.1$ and $[B] = 1.0$ corresponding to the input $A =$`false` and $B =$`true`. The output of the network is a rising concentration of the `true` form of variable C (line with crosses). At $t = 100$, another input is given by increasing the concentration of A to 1.1 and B to 1.0. Accordingly, this leads to an increase of the concentration of $\overline{C}$, corresponding to an output of $C =$`false` (line with squares).

Using Chemical Organization Theory for Model-checking

The correspondence of each fixed-point of a reaction network to an organization of the same system can be used to check the consistency of reaction networks. This property can even be extended to states of the system where the concentration of some species grows [19]. Thus, Chemical Organization Theory can be used to check which species of a reaction network can be present in a long-term simulation of the system. Since the analysis does not require an explicit knowledge of the kinetics of the reactions, this approach also extends to networks with partly or completely unknown kinetics.

Analyzing, for example, the reaction network depicted in Fig. 10.1 without reaction (4) and the associated Hasse diagram of organizations in Fig. 10.2a, we find that the largest organization does not contain species c. Thus, neither this species nor reactions using it as a substrate can persist in a long-term simulation of the reaction network.

Study of Chemical Evolution

Chemical Organization Theory can also be used to study chemical evolution, a process thought to have preceded the appearance of the first organism on Earth. Matsumaru et al. [24] analyzed the evolution of a chemistry of binary strings interacting with each other. The interaction between two strings produced new strings and, thus, could be transformed into a set of reaction rules. In a first setting, a system without mutation events was examined. It was found that the diversity of the system (i.e., the number of different strings diminished over time) as did the number of chemical organizations. Finally, only two organizations remained. In a second step, mutation events randomly flipping a position of a small number of strings at given time-points were introduced. In this setting, interesting effects could be observed. These events increased the diversity both in the set of species and organizations even though there is no trivial relation between the numbers of organizations and species.

Events in the evolution of a reaction network can be interpreted as either an upward, downward or sideward movement in the set of organizations of the system. A downward movement corresponds to the disappearance of certain organizations of a set of organizations and an upward movement corresponds to organizations appearing anew. A sideward movement corresponds to a mixture of both types of events. The sideward movement can be caused by species appearing anew, resulting in a upward movement instantaneously followed by a downward movement. While downward movements usually correspond to a decrease in diversity, upward movements correspond to an increase.

10.4 Software for the Computation of Chemical Organizations

The three central aspects of the software for the analysis using Chemical Organization Theory are the computation of chemical organizations, the visualization of the Hasse diagram of organizations and the integration into existing software frameworks for the analysis of complex reaction networks. These aspects are offered by the applications of *OrgTools*. The applications can be downloaded from [17]. A schematic overview on the applications of *OrgTools* and their integration into SBW is given in Fig. 10.4.

Three different tools allow the computation of the chemical organizations of reaction networks. Each of them is dedicated to some commonly appearing setting in which chemical organizations are being computed.

OrgAnalysis aims to offer some basic network editing capabilities and an integration into SBW [31]. Thus, it integrates into a rich framework of commonly used tools for network editing, simulation and analysis.

OrgFinder, in contrast, represents a basic command line tool which offers almost the same options for organization computation as *OrgAnalysis*. Thus, is can be easily integrated into scripts and allows a simple "conversion" of reaction networks into the corresponding set of chemical organizations.

OrgBatch is focused on the batch computation of chemical organizations. This is especially of interest if, for example, the evolution of a reaction network is simulated and the chemical organizations of different time steps of the network should be analyzed.

The central concept in Chemical Organization Theory is represented by the Hasse diagram of organizations. *OrgView* visualizes the Hasse diagram of organizations and thus allows a structured analysis of the chemical organizations of a reaction network. Additionally, it is possible to compare different Hasse diagrams at once and, for example, analyze the evolution of chemical organizations in the course of the evolution of a reaction network.

All of these tools use SBML [16] as a standard for the description of reaction networks. Together with an integration into SBW [31], many commonly used tools for the design and analysis of reaction networks as described in Section 10.2 become easily accessible.

10.4.1 OrgAnalysis

OrgAnalysis offers simple network editing capabilities integrated into an application for organization computation. It can either be started directly or by any other SBW-compliant application through its SBW interface. As it integrates into the Analysis section of SBW, it can be called with a string representing a reaction network in the SBML format. When called, for example,

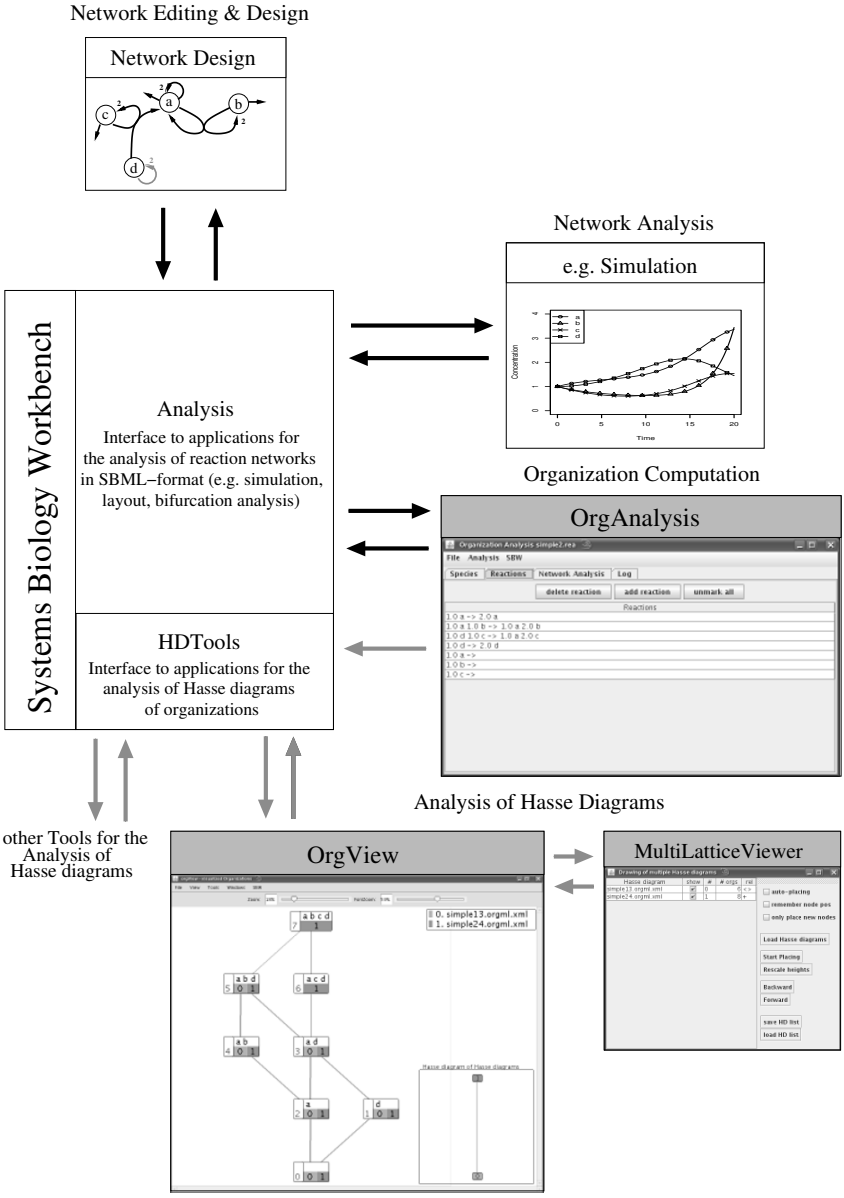

Fig. 10.4 Schematic drawing of the integration of OrgTools into SBW. Applications with gray headings denote *OrgTools*. Black arrows denote reactions networks being passed and gray arrows correspond to Hasse diagrams. First, the network is designed using, for example, *CellDesigner*. Then the network can be passed through the Analysis section of SBW (e.g., for simulation or computation of chemical organizations (*OrgAnalysis*)). The Hasse diagram of organizations can then be passed to *OrgView* for visualization. *OrgView* interfaces to *MultiLatticeViewer* to control the display of several Hasse diagrams at once. If a second Hasse diagram is loaded into *OrgView*, it is automatically forwarded to *MultiLatticeViewer*.

from a network design tool interfacing with SBW, the corresponding reaction network is directly loaded without the prior need of saving the network to a file.

Network Editing

Since there is a whole range of dedicated network editing software available, *OrgAnalysis* only concentrates on editing capabilities that help in the analysis of chemical organizations. Thus, it offers an analysis mechanism that identifies species that are only produced or consumed. Since the self-maintenance condition requires each species of an organization to be producible if it is consumed (cf. Section 10.3.1), the addition of an inflow for a species that is only consumed can lead to new organizations and, thus, to new qualitative behavior of the model during simulation. The metabolite section of *OrgAnalysis* allows one to add outflow or inflow reactions for species. Alternatively, it is possible to set a species to external; that is, it is considered present at any time and buffered by reactions outside of the system. These capabilities are complemented by a simple text-based mechanism for editing reactions.

Computation of Chemical Organizations

All of the presented algorithms for the computation of chemical organizations (Section 10.3.2) are integrated into *OrgAnalysis*. While the flux-based approach has been completely implemented in Java, the heuristic and the constructive approaches use the linear programming library *LpSolve* [3] shipped along with *OrgTools* as a subroutine.

All algorithms offer a range of options to focus the computation of chemical organizations onto certain aspects. Thus, the constructive algorithm by default only computes the set of connected organizations. This option can be of importance if the network contains many organizations that are a superset to a certain reactive organization, without containing any additional reactions [5]. Thus, they contain molecules that do not participate in any reaction within the organization. Additionally, it is possible to calculate regulatory organizations. Following an approach outlined in [18] and [23], this allows the integration of boolean rules governing the availability of reactions in the network. Some reactions can, for example, be constrained to zero flux if certain conditions formulated in a boolean formula are met (see Section 10.3.3 for details on the integration of boolean formalisms into reaction networks).

SBW Interface

The SBW interface allows the integration of *OrgAnalysis* into some commonly used network analysis tools. Additionally, it serves as interface between *OrgAnalysis* and *OrgView* to visualize the computed Hasse diagrams. Several Hasse diagrams can be visualized in the same instance of *OrgView* to allow an easier comparison of the results. Additionally, *OrgAnalysis* can call any application that possesses an interface to the Analysis section of SBW.

Model-checking

As outlined in Section 10.3.3, Chemical Organization Theory can be used to analyze the long-term behavior of reaction networks during simulation. Thus, *OrgAnalysis* offers the capability to directly examine the chemical organizations of a network and identify species and reactions that cannot be present in a long-term simulation. In some cases, only the short-time behavior is of interest. Consequently, if some species or reactions were found absent from any organization, the analysis is repeated by taking into account species whose concentration is either initially nonzero or is set to a positive value at a given time-point. If this second step detects species or reactions that do not appear in any organization, they cannot have a positive concentration, respectively, flux during simulation. This hints to possible modelling inconsistencies [19].

10.4.2 Other Applications for Organization Computation

OrgBatch

In some applications (e.g., in the simulated evolution of a reaction network), it becomes desirable to compute the organizations of a whole list of reaction networks at once. In this case *OrgBatch* can be used. It implements the constructive and the flux-based approach for organization computation, including various options that can be used with both algorithms. Additionally, it is possible to process files by size in order to obtain results for those networks first where computation time is expected to be shortest. For some networks, the computation of organization does not finish in reasonable time. As discussed in [5], this might be due to network size or network topology. In such a case, computation can be aborted either directly or after a user-specified time-span.

OrgFinder

In order to offer an the simplest possible mechanism for the computation of chemical organizations, the commandline tool *OrgFinder* can be used. By just providing a network file in SBML format, the organizations are computed directly and the Hasse diagram is written into a file in the XML-based *OrgML* format. Both deterministic algorithms for organization computation with the same options as in *OrgAnalysis* can be used.

10.4.3 OrgView

The purpose of *OrgView* is the visualization of Hasse diagrams of organizations. The Hasse diagrams can be loaded either through the SBW interface or from a file. In order to analyze a Hasse diagram, an appropriate placing of its nodes is necessary. Thus, different placing algorithms have been implemented. Multiple lattices can be loaded at once to allow an easier comparison of the chemical organizations of different networks.

Loading Hasse Diagrams

Hasse diagrams can be loaded into *OrgView* in two ways. First, *OrgView* can be called with a computed Hasse diagram from *OrgAnalysis* via the *HDTools* section of the SBW menu. Second, a Hasse diagram can be loaded from a file. This file can be either in the XML-based *OrgML* or in the text-based *.ltc* format.

Since information on the network is necessary for the visualization of the Hasse diagram, the corresponding network file in SBML format is additionally required. To facilitate the access to the network file, *OrgView* automatically searches for a network file obeying the following naming convention. For an *OrgML* file with the name *network.orgml.xml* or *network.ltc*, the corresponding network file is expected to be *network.xml* in the SBML format.

Placing of Hasse Diagrams

The nodes of the Hasse diagram of organizations are normally arranged with growing size from bottom to top. To allow a more user-friendly display of lattices, two placing algorithms have been implemented. Both offer fine-tuning of options for user-oriented demands. The simulated annealing-based placer yields the best results in positioning of nodes but has a high requirement in computation time for large Hasse diagrams. The parallel placer offers good results with a reduced requirement in computation time. Additionally, avail-

able options include the display of species only in nodes, where they appear anew; that is, each node only displays species that do not appear in any organization that are subset to this organization. Furthermore, the display of species can be completely switched off in order to concentrate on the topology of the Hasse diagram.

Displaying Several Hasse Diagrams at Once

A useful tool in organization analysis is the comparison of the chemical organizations of several reaction networks. Thus, *OrgView* allows the display of several Hasse diagrams at once. The Hasse diagrams of all networks are united and organizations that appear in several networks are indicated. A so-called "Hasse diagram of Hasse diagrams" allows the analysis of the relations of the organization of the networks between each other. Thus, if two networks contain the same set of organizations, they are shown as one node in this meta-diagram. If one network contains a subset of the organizations of another network, the node corresponding to the first Hasse diagram lies below the node corresponding to the second one (see Fig. 10.5 for an example).

The display of several Hasse diagrams at once is facilitated by *MultiLatticeViewer* integrated into *OrgView*. It allows one to load a list of lattices and the specification of the lattices to display. Additionally, it interfaces with SBW. Thus, if the same instance of *OrgView* is called several times with a Hasse diagram from *OrgAnalysis*, the Hasse diagrams are loaded into *MultiLatticeViewer*. Such a functionality is useful when designing a reaction

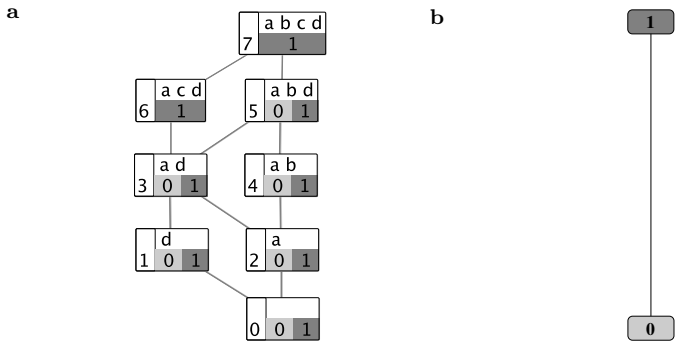

Fig. 10.5 **a** Integrated view of the two Hasse diagrams from Fig. 10.2. **b** Hasse diagram of Hasse diagrams of Fig. 10.2. The Hasse diagrams to which each organization belongs is indicated by a number below the set of species. Hasse diagram 0 contains the chemical organizations of the initial network from Fig. 10.1 and Hasse diagram 1 contains the organizations after the addition of one reaction producing species *d*. All organizations in Hasse diagram 0 appear in Hasse diagram 1. Hence, the Hasse diagram of Hasse diagram consists of a node labeled "0" below the node labeled "1."

network and examining the changes in the set of chemical organization that appear through addition or removal of certain reactions.

10.4.4 Organization Computation on the Web

Since one of the algorithms for organization computation and all of the presented tools have been implemented in Java, an applet that offers a similar functionality is available from [17]. The advantage of an applet is that it can be run without any prior installation on all systems that have a Java-capable web browser. The applet integrates the basic functionalities of *OrgAnalysis* and *OrgView*. Additionally, it contains an interface to a repository of curated biological networks in SBML format, the BioModels database [26] and a list of sample networks that have been studied using chemical organization theory. A tutorial on the use of the applet is linked on [17].

Analysis View

The analysis view allows the loading of network files by pasting a file in SBML format or a URL linking to such a file into the editor. After loading a network, it can be edited either directly in a text editor displaying the SBML or a more human-readable proprietary format, or via the same tools offered in *OrgAnalysis*. Thus, species and reactions can be added, deleted or modified. Additionally, species can be specified as external (i.e., adding an inflow and an outflow reaction). Subsequently, the chemical organizations of the reaction network can be computed and visualized. As in *OrgAnalysis*, it is additionally possible to analyze the long-term behavior of the model during simulation (see Sections 10.3.3 and 10.4.1 for more details). Thus, species and reactions that cannot be present in a long-term simulation of the reaction network or even cannot have any positive concentration or flux during the simulation can be identified.

Interface to the BioModels Database

To enlarge the list of networks available to the applet, an interface to the BioModels database [26] has been implemented. This database contains a whole list of curated biological models. When selecting a model, information available from the corresponding SBML file is displayed. Double-clicking on a model in the list loads it into the applet.

Visualization View

Some basic functionalities of *OrgView* are offered by the applet: Two placing algorithms for Hasse diagrams are available; some of the most commonly used options for the display of Hasse diagrams can also be specified; additionally, the visualized Hasse diagram can be exported into the .eps format.

Web-Integration

Another aspect of the applet is the possibility to directly access some of its capabilities from any webpage. It can be called with several parameters automatizing some of the most commonly used functions. It is, for example, possible to directly load a reaction network from a URL and display the Hasse diagram of organizations in one step. Details on this integration are given in the tutorial in [17].

10.5 Conclusion

In this chapter, *OrgTools*, a set of applications dedicated at the analysis of chemical reaction networks from Artificial Chemistry and Systems Biology, has been presented. These applications integrate into SBW and, thus, come along with a framework that covers almost any aspect in network analysis and design. Chemical Organization Theory is a good example of how work in Artificial Chemistry can benefit the work in Systems Biology. The applications of this concept reach from the study of boolean networks, the evolution of the first life to the study of "real-world" networks from Systems Biology. One of the most appealing properties of Chemical Organization Theory that help the study of such systems is the correspondence of each fixed-point of a reaction network to an organization. Since only information on the algebraic structure of the network (i.e., the set of species and reactions) is needed, even networks where information about them is unavailable can be analyzed. Thus, the species sets in the fixed points of a reaction network can be determined without needing to solve the corresponding kinetic laws analytically. In consequence, even the possible phenotypes of a reaction network (i.e., the set of species that can coexist over a long period of simulation-time) can be determined without an explicit knowledge of the kinetics of the network.

In a simplified setting, the installation of a whole system for the analysis of chemical organizations might be too much of an afford. Tools for the computation of organizations without the necessity of installation have been developed. In consequence, it is possible to compute chemical organizations and perform most of the analysis that has been described in this work only

using a Java-capable web browser. Thus, many tools tailored to the specific needs of the user are available.

Acknowledgements I thank Pietro Speroni di Fenizio and Maciej Komosinski for helpful comments on the manuscript.

References

1. Banâtre, J.P., Métayer, D.L.: A new computational model and its discipline of programming. Tech. Rep. RR-0566, INRIA (1986)
2. Banâtre, J.P., Métayer, D.L.: The gamma model and its discipline of programming. Sci. Comput. Program. **15**(1), 55–77 (1990)
3. Berkelaar, M., Eikland, K., Notebaert, P.: lp solve: Open source (mixed-integer) linear programming system, version 5.5 (2005). URL `http://tech.groups.yahoo.com/group/lp_solve/`
4. Centler, F., Fenizio, P., Matsumaru, N., Dittrich, P.: Chemical organizations in the central sugar metabolism of *Escherichia coli*. Math. Model. Biol. Syst. **1**, 105–119 (2007)
5. Centler, F., Kaleta, C., di Fenizio, P.S., Dittrich, P.: Computing chemical organizations in biological networks. Bioinformatics **24**(14), 1611–1618 (2008)
6. Chickarmane, V., Paladugu, S.R., Bergmann, F., Sauro, H.M.: Bifurcation discovery tool. Bioinformatics **21**(18), 3688–3690 (2005)
7. Dittrich, P., Matsumaru, N.: Organization-oriented chemical programming. In: Hybrid Intelligent Systems, 2007. HIS 2007. 7th International Conference on, pp. 18–23 (2007)
8. Dittrich, P., Speroni di Fenizio, P.: Chemical organization theory. Bull. Math. Biol. **69**(4), 1199–1231 (2007)
9. Dittrich, P., Ziegler, J., Banzhaf, W.: Artificial chemistries – a review. Artificial Life **7**, 225–275 (2001)
10. Drubin, D.A., Way, J.C., Silver, P.A.: Designing biological systems. Genes. Dev. **21**(3), 242–254 (2007)
11. Endy, D.: Foundations for engineering biology. Nature **438**(7067), 449–453 (2005)
12. Fontana, W., Buss, L.W.: 'The arrival of the fittest': Towards a theory of biological organization. Bull. Math. Biol. **56**, 1–64 (1994)
13. Fukuda, K., Prodon, A.: Double description method revisited. In: Selected papers from the 8th Franco-Japanese and 4th Franco-Chinese Conference on Combinatorics and Computer Science, pp. 91–111. Springer-Verlag, London, UK (1996)
14. Giavitto, J.L., Michel, O.: MGS: a rule-based programming language for complex objects and collections. In: M. van den Brand, R. Verma (eds.) Electr. Notes Theor. Comput. Sci., vol. 59, pp. 286–304. Elsevier Science Publishers (2001)
15. Hucka, M., Finney, A., Bornstein, B.J., Keating, S.M., Shapiro, B.E., Matthews, J., Kovitz, B.L., Schilstra, M.J., Funahashi, A., Doyle, J.C., Kitano, H.: Evolving a lingua franca and associated software infrastructure for computational systems biology: the Systems Biology Markup Language (SBML) project. Syst. Biol. (Stevenage) **1**(1), 41–53 (2004)
16. Hucka, M., Finney, A., Sauro, H.M., Bolouri, H., Doyle, J.C., Kitano, H., Arkin, A.P., Bornstein, B.J., Bray, D., Cornish-Bowden, A., Cuellar, A.A., Dronov, S., Gilles, E.D., Ginkel, M., Gor, V., Goryanin, I.I., Hedley, W.J., Hodgman, T.C., Hofmeyr, J.H., Hunter, P.J., Juty, N.S., Kasberger, J.L., Kremling, A., Kummer, U., Novère, N.L., Loew, L.M., Lucio, D., Mendes, P., Minch, E., Mjolsness, E.D., Nakayama, Y., Nelson, M.R., Nielsen, P.F., Sakurada, T., Schaff, J.C., Shapiro, B.E., Shimizu, T.S., Spence,

H.D., Stelling, J., Takahashi, K., Tomita, M., Wagner, J., Wang, J., Forum, S.B.M.L.: The Systems Biology Markup Language (SBML): a medium for representation and exchange of biochemical network models. Bioinformatics **19**(4), 524–531 (2003)
17. Kaleta, C.: Organization computation on the web. URL `http://orgtheo.artificialchemistries.org/`
18. Kaleta, C., Centler, F., di Fenizio, P.S., Dittrich, P.: Phenotype prediction in regulated metabolic networks. BMC Syst. Biol. **2**(1), 37 (2008)
19. Kaleta, C., Richter, S., Dittrich, P.: Using chemical organization theory for model-checking (2009). Submitted
20. Kitano, H., Funahashi, A., Matsuoka, Y., Oda, K.: Using process diagrams for the graphical representation of biological networks. Nat. Biotechnol. **23**(8), 961–966 (2005)
21. Ma, W., Johnson, C., Brent, R.: Implementation of the Chemical Abstract Machine on a MIMD computer. Tech. rep., CS Lab, ANU, Australia (1996)
22. Matsumaru, N., Centler, F., di Fenizio, P.S., Dittrich, P.: Chemical organization theory applied to virus dynamics (Theorie chemischer organisationen angewendet auf infektionsmodelle). Inform. Technol. **48**(3), 154–160 (2006)
23. Matsumaru, N., Centler, F., Speroni di Fenizio, P., Dittrich, P.: Chemical organization theory as a theoretical base for chemical computing. Int. J. Unconv. Comp. **3**(4), 285–309 (2007)
24. Matsumaru, N., di Fenizio, P.S., Centler, F., Dittrich, P.: On the evolution of chemical organizations. In: S. Artmann, P. Dittrich (eds.) Proc. 7th German Workshop on Artificial Life, pp. 135–146. IOS Press, Amsterdam, NL (2006)
25. Matsumaru, N., Lenser, T., Hinze, T., Dittrich, P.: Toward organization-oriented chemical programming: A case study with the maximal independent set problem. In: F. Dressler, I. Carreras (eds.) Advances in Biologically Inspired Information Systems, *Studies in Computational Intelligence (SCI)*, vol. 69, pp. 147–163. Springer-Verlag, Berlin, Heidelberg, New York (2007)
26. Novére, N.L., Bornstein, B., Broicher, A., Courtot, M., Donizelli, M., Dharuri, H., Li, L., Sauro, H., Schilstra, M., Shapiro, B., Snoep, J.L., Hucka, M.: BioModels database: A free, centralized database of curated, published, quantitative kinetic models of biochemical and cellular systems. Nucleic Acids Res. **34**(Database issue), D689–D691 (2006)
27. Pfeiffer, T., Sánchez-Valdenebro, I., no, J.C.N., Montero, F., Schuster, S.: METATOOL: For studying metabolic networks. Bioinformatics **15**(3), 251–257 (1999)
28. Păun, G.: From cells to computers: computing with membranes (P systems). Biosystems **59**(3), 139–158 (2001)
29. Radhika, V., Proikas-Cezanne, T., Jayaraman, M., Onesime, D., Ha, J.H., Dhanasekaran, D.N.: Chemical sensing of DNT by engineered olfactory yeast strain. Nat. Chem. Biol. **3**(6), 325–330 (2007)
30. Ramsey, S., Orrell, D., Bolouri, H.: Dizzy: Stochastic simulation of large-scale genetic regulatory networks. J. Bioinform. Comput. Biol. **3**(2), 415–436 (2005)
31. Sauro, H.M., Hucka, M., Finney, A., Wellock, C., Bolouri, H., Doyle, J., Kitano, H.: Next generation simulation tools: the systems biology workbench and biospice integration. OMICS **7**(4), 355–372 (2003)
32. SBML Community: Systems Biology Markup Language (SBML) website. URL `http://www.sbml.org/`
33. Schuster, S., Fell, D.A., Dandekar, T.: A general definition of metabolic pathways useful for systematic organization and analysis of complex metabolic networks. Nat. Biotechnol. **18**(3), 326–332 (2000)
34. Varma, A., Palsson, B.: Metabolic flux balancing: Basic concepts, scientific and practical use. Bio/Technology **12**, 994–998 (1994)