
Generic strategies for chemical space exploration

Jakob L. Andersen

Department of Mathematics and Computer Science,
University of Southern Denmark,
Campusvej 55, DK-5230 Odense M, Denmark

and

Max Planck Institute for Mathematics in the Sciences,
Inselstraße 22 D-04103 Leipzig, Germany
E-mail: jlandersen@imada.sdu.dk

Christoph Flamm*

Institute for Theoretical Chemistry,
University of Vienna,
Währingerstraße 17, A-1090 Wien, Austria
E-mail: xtof@tbi.univie.ac.at

Daniel Merkle*

Department of Mathematics and Computer Science,
University of Southern Denmark,
Campusvej 55, DK-5230 Odense M, Denmark
E-mail: daniel@imada.sdu.dk

*Corresponding authors

Peter F. Stadler

Institute for Theoretical Chemistry,
University of Vienna,
Währingerstraße 17, A-1090 Wien, Austria
Bioinformatics Group, Department of Computer Science,
and Interdisciplinary Center for Bioinformatics,
Härtelstraße 16-18, D-04107, Leipzig, Germany

and

Max Planck Institute for Mathematics in the Sciences,
Inselstraße 22 D-04103 Leipzig, Germany

and

Fraunhofer Institute for Cell Therapy and Immunology,
Perlickstraße 1, D-04103 Leipzig, Germany

and

Center for non-coding RNA in Technology and Health,
University of Copenhagen,
Grønnegårdsvej 3, DK-1870 Frederiksberg C, Denmark

and

Santa Fe Institute,
1399 Hyde Park Rd, Santa Fe, NM 87501, USA
E-mail: studla@bioinf.uni-leipzig.de

Abstract: The chemical universe of molecules reachable from a set of start compounds by iterative application of a finite number of reactions is usually so vast, that sophisticated and efficient exploration strategies are required to cope with the combinatorial complexity. A stringent analysis of (bio)chemical reaction networks, as approximations of these complex chemical spaces, forms the foundation for the understanding of functional relations in Chemistry and Biology. Graphs and graph rewriting are natural models for molecules and reactions. Borrowing the idea of partial evaluation from functional programming, we introduce partial applications of rewrite rules. A framework for the specification of exploration strategies in graph-rewriting systems is presented. Using key examples of complex reaction networks from carbohydrate chemistry we demonstrate the feasibility of this high-level strategy framework. While being designed for chemical applications, the framework can also be used to emulate higher-level transformation models such as illustrated in a small puzzle game.

Keywords: chemical spaces; generative chemistries; graph grammars; graph transformation; double pushout approach; graph binding; partial rule application; rule composition; borate stabilised Formose reaction.

Reference to this paper should be made as follows: Andersen, J.L., Flamm, C., Merkle, D. and Stadler, P.F. (2014) 'Generic strategies for chemical space exploration', *Int. J. Computational Biology and Drug Design*, Vol. 7, Nos. 2/3, pp.225–258.

Biographical notes: In 2011, Jakob L. Andersen received his Master's degree in Computer Science from the University of Southern Denmark. He is currently aiming at a double PhD in Computer Science under supervision of Daniel Merkle, University of Southern Denmark and Peter F. Stadler, Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany. His research interests include cheminformatics, graph rewriting and in general graph algorithms.

Christoph Flamm received a Mag. rer. nat. (Organic Chemistry, focus area: natural product synthesis) and a Dr. rer. nat. (Theoretical Chemistry, focus area: folding kinetics of RNA molecules) from the University of Vienna. In Fall 2006, he was conferred the 'venia docendi' and became Associated Professor in the Department of Theoretical Chemistry at the University of Vienna. CF's main expertise is the development of theoretical models and methods for complex biological systems and the implementation of software tools for their simulation and analysis. His current research interests are computational models of the structure, dynamics, and evolution of chemical reaction networks with special emphasis on metabolic networks.

Daniel Merkle is an Associate Professor at the Department of Mathematics and Computer Science at the University of Southern Denmark since 04/2008. He received his ‘Diplom in Computer Science’ in 1997 at the Department of Computer Science of the University of Karlsruhe (TH). In 2002 he received his PhD from the Institute of Applied Informatics and Formal Description Methods (AIFB), University of Karlsruhe (TH). From October 2002 until April 2008 he has been Assistant Professor at the Parallel Computing and Complex Systems Group, at the Department of Mathematics and Computer Science, University of Leipzig. His research interests are Cheminformatic and Parallel Computing.

Peter F. Stadler received his PhD in Chemistry from U. Vienna in 1990 and then worked as Assistant and Associate Professor for Theoretical Chemistry at the same School. In 2002 he moved to Leipzig as Full Professor for Bioinformatics. Since 1994 he is External Professor at the Santa Fe Institute. He is an External Scientific Member of the Max Planck Society since 2009 and an Corresponding Member Abroad of the Austrian Academy of Sciences since 2010.

This paper is a revised and expanded version of a paper entitled ‘Generic strategies for chemical space exploration’ presented at *International Conference on Intelligent Biology and Medicine (ICIBM)*, Nashville, TN, USA, August 11–13, 2013.

1 Introduction

The systematic computational exploration of chemical spaces have become topic of high practical relevance e.g., in drug design (Eberhardt et al., 2011; Dow et al., 2012; Reymond and Awale, 2012; Wong, 2012). Recent efforts to gain insights into the distribution of properties in chemical spaces include the construction of large databases of hypothetical compounds. The “chemical universe database GDB-17” (Ruddigkeit et al., 2012), for instance comprises 166.4 billion molecules of up to 17 atoms of C, N, O, S, and halogens covering the size and composition range of typical lead compounds. Beyond the diversity of molecules and their properties, however, potential synthesis pathways leading to them are a crucially important consideration in practice. This calls for methods to systematically explore chemical spaces in terms of restricted types of chemical reactions. Here we demonstrate how this can be achieved in a natural way by means of graph grammars in conjunction with efficient exploration strategies.

The *structural formula* of a chemical compound is a graph that represents the connectivity and mutual arrangements of its atoms. Atom types are given as vertex labels, while edges represent bond types. At this level of modelling, chemical reactions are naturally represented as graph transformations. Chemical reactions are explained and categorised in terms of *reaction mechanisms* that encapsulate the local changes of chemical bonds. In the formal framework of graph grammars, reaction mechanisms correspond to the productions (rules). Because of this conceptual alignment between chemistry and graph grammars, a variety of artificial chemistry models of different degree of chemical realism have been

devised on this basis (Dittrich et al., 2001). Of course, these purely combinatorial models of chemistry have their limitations. Deliberately disregarding the spatial embedding of molecules they cannot capture many aspects of stereochemistry and they are restricted to (over)simplified models of reactions energies and reaction kinetics. Graph grammar models are nevertheless of practical interest when the task is to explore large areas of chemical spaces and they provide a means of analysing regularities in very large reaction networks.

Several graph rewriting tools have become available in the recent past, see Fernández and Namet (2010) for an overview. Application areas beyond chemistry include model checking and verification, proof representation, and modelling control flow of programs among many others. A strategy language to control the application of graph rewriting rules has been presented in Fernández et al. (2012) for PORGY (Andrei et al., 2011; Pinaud et al., 2012). A strategy framework for exploring chemical spaces has very different design goals from the properties desirable for applications within different areas of computer science. For instance, chemical graph grammar rule frequently merge or split graphs, since connected components correspond to individual molecules. Hence a chemically motivated component handling is required. The theoretically reachable chemical spaces can be infinite, e.g., when the rules and the starting material allow polymers to form. Exploration thus may not halt except due an enforced resource (size) limitation. Decisions on how to expand the space are usually heavily influenced by chemical properties or additional data sources. Furthermore, the goal of an analysis might also be motivated by a chemical question such as the detection of chemical subspaces or the need to find specific chemical transformation patterns.

In this paper different types of ‘chemistries’ will be used to demonstrate the different aspects of the strategic construction framework for fractions of the chemical space. The Diels-Alder reaction (Diels and Alder, 1928), a sigmatrope cycloaddition reaction between a conjugated diene and an alkene will serve as an example for a combinatorial complex chemical space emerging from a single reaction rule. The formose reaction (Butlerov, 1861), which subsumes the formation of sugars from formaldehyde is used to explore the impact of changes in ‘chemistry’ (i.e. the set of reaction rules) on the structure and complexity of the chemical space. HCN chemistry is used to illustrate how construction of chemical space can be biased with experimental data.

The outline of the paper is as follows. In Section 2 we will present the underlying framework of graph transformation including the Double Pushout Approach. Secondly, in Section 3, we describe a method for efficient calculation of rule application. Thirdly, general strategies to explore chemical spaces will be introduced in Section 4. In Section 5 we will apply our framework to several before-mentioned complex chemical settings. Finally a puzzle game will be used to illustrate the generality of our approaches.

2 Formal framework

2.1 Chemical graph rewriting with the double pushout approach

Molecules are always represented by connected graphs. Chemical reactions, however, more often than not, involve two or more interacting molecules as their ‘input’ (educts) and there

is no guarantee that the ‘output’ (products) is connected. Thus we have to consider graph transformations that operate on not necessarily connected graphs. More precisely, we regard a graph G here as a multiset $\{g_1, g_2, \dots, g_{\#G}\}$ of its $\#G$ connected components. All graphs are simple, i.e., without loops and parallel edges. Double and triple bonds are viewed as edge labels rather than multiple edges.

Several abstract formalisms for graph transformation have been explored in the literature, see e.g., Rozenberg and Ehrig (1997) for a detailed introduction. We found that the so-called Double Pushout (DPO) approach provides the most intuitive direct encoding of chemical reactions and the closest connection to the language of chemistry. A DPO transformation rule $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ consists of three graphs L , R and K known as the left, right and context graph, respectively, and two graph morphisms l and r that determine how the context is embedded in the left and the right graph. The rule p can be applied to a graph G if the left graph L can be found in G and some additional consistency conditions are satisfied. This is modelled by the requirement that there is a *matching morphism* $m : L \rightarrow G$ that describe how L is contained in G . Intuitively, the copy of L is replaced within G by R in such a way that the context K is left intact, resulting in the transformed graph H . This operation, the derivation $G \xrightarrow{p,m} H$, is described in the framework of category theory by the requirement that the following commutative diagram exists:

$$\begin{array}{ccccc} L & \xleftarrow{l} & K & \xrightarrow{r} & R \\ \downarrow m & & \downarrow d & & \downarrow n \\ G & \xleftarrow{} & K & \xrightarrow{} & H. \end{array} \quad (1)$$


The derivation $G \xrightarrow{p,m} H$ implicitly define the intermediary graph D and the result graph H as well as morphisms $d : K \rightarrow D$ and $n : R \rightarrow H$ that fix how the context and the right graph of the rule are embedded in the intermediary and the result graph, respectively. In terms of molecules (connected components) we can write $\{g_1, g_2, \dots, g_{\#G}\} \Rightarrow \{h_1, h_2, \dots, h_{\#H}\}$.

In applications to modelling chemistry, several additional requirements must be satisfied. Conservation of mass and atom types dictates that the restrictions of r and l to the vertex sets (atoms) are bijective. Furthermore, m (and by extension d and n) are subgraph isomorphisms and hence injective. We note in passing that this guarantees the existence of a bijection $a : V(G) \rightarrow V(H)$ known as the *atom mapping*. In the DPO formalism, furthermore, the existence of an inverse production $p^{-1} = (L \xleftarrow{l} K \xrightarrow{r} R)$, corresponding to the reverse chemical reaction, is guaranteed. Some more basic properties of chemical graph grammars can be found in Andersen et al. (2013b). Figure 1 shows an example of a chemical derivation.

2.2 Proper derivations

Consider a valid derivation $\{g_1, g_2\} \xrightarrow{p,m} \{h_1, h_2\}$ and an arbitrary graph g' . Clearly, the derivation $\{g_1, g_2, g'\} \xrightarrow{p,m} \{h_1, h_2, g'\}$ is also valid because the images of m and n are contained in $\{g_1, g_2\}$ and $\{h_1, h_2\}$, respectively. The graph g' is irrelevant for the transformation. We call a derivation $G \xrightarrow{p,m} H$ *proper* if $\text{img } m \cap g_i \neq \emptyset$ for all $g_i \in G$. It is not hard to see that the inverse of a proper derivation is again proper. Throughout the following sections we will assume every derivation to be proper, unless otherwise stated.

Figure 1 Example of a chemical derivation from cyclohexadiene and isoprene using a Diels-Alder transformation. The edges changed by the transformation is shown in red and the vertices from K are shown in green. Note that edges shown in parallel are in the underlying graphs a single edge with a special label to encode a specific chemical bond (see online version for colours)




2.3 Derivation graphs

Chemical reaction networks can be represented as directed (multi)hypergraphs whose vertices are the molecules of the ‘chemical universe’ under consideration and whose hyperedges represent chemical reactions (Zeigarnik, 2000). Here, it is important to consider hyperedges as multisets to accommodate the stoichiometric coefficients, i.e., the multiplicities in which molecules enter a chemical reaction such as $2\text{H}_2 + \text{O}_2 \rightarrow 2\text{H}_2\text{O}$. Such networks can be constructed from experimentally observed data. An example is the network of organic chemistry (NOC) (Bishop et al., 2006; Fialkowski et al., 2005; Grzybowski et al., 2009), which shows a non-trivial organisation concentrated around a core region of about 300 synthetically important building blocks and industrial compounds. Metabolic networks consist of the enzymatically catalysed reactions constituting the chemical basis of modern life forms. They are available from dedicated databases, see e.g., Karp and Caspi (2011).

In the framework of graph grammar models, an analogous *derivation graph* can be defined. Its vertex set consists of the connected labeled graphs \mathcal{G} that represent the molecules. Directed hyperedges connect the multisets $G \subseteq \mathcal{G}$ and $H \subseteq \mathcal{G}$ only if there is a proper derivation $G \xrightarrow{p,m} H$. The conventions for visualising hyperedges adhere to the three examples in Figure 2.

Figure 2 A bipartite graph notation for directed (multi-)hypergraphs, in which the production itself is drawn as a special type of intermediate vertex, is used in most cases: (a) $\{g_1, g_2\} \xrightarrow{p} g_3$. We only make an exception for 1-to-1 transformations (isomerisation reactions); (b) $g_4 \xrightarrow{p} g_5$. Multiplicities are indicated by multiple arcs and (c) $\{g_6, g_6\} \xrightarrow{p} g_7$



3 Transformation by partial rule application

The core strategy to expand the underlying derivation graph is the discovery of new graphs by means of proper derivations implied by the direct application of rules. Given a rule $p = (L \leftarrow K \rightarrow R)$ and a set of graphs \mathcal{U} , the task is to find all proper derivations $G \xrightarrow{p} H, G \subseteq \mathcal{U}$ where G and H are multisets of graphs. This can be done by a testing of all k -multisubsets of \mathcal{U} for all $1 \leq k \leq \#L$. Since nearly all chemical reactions are mono-molecular or bi-molecular, we can restrict ourselves to $\#L \leq 2$, at least when elementary reactions are of primary interest. Still, the number of multisets is $O(|\mathcal{U}|^2)$. In the worst case, all unique multisets may give successful transformations, often leading to a combinatorial explosion that quickly becomes unmanageable. In the following section we show that a more detailed control of the multisets that are considered for transformation is desirable.


The key concept is partial rule composition (Andersen et al., 2013b), i.e., the binding of graphs to rules, resulting in partial rules that can be applied more efficiently in an exploration strategy. The idea is analogous to partial evaluation of functions by binding some of the variables. Full graph transformations are computed as repeated partial rule application in this framework. For the sake of brevity, we only sketch the idea here and omit a complete formal definition of partial rules.

A partial rule application of a rule $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ with $L = \{l_1, l_2, \dots, l_{\#L}\}$ to a graph G , is a generalisation of a full transformation of G in which only some but not all components of L do not match G . Thus L is partitioned into the matching part $\bar{L} \neq \emptyset$ and the non-matching remainder L' . The restriction \bar{l} of $l : K \rightarrow L$ to the pre-image \bar{K} of \bar{L} defines the partial transformation rule $\bar{p} = (\bar{L} \xleftarrow{\bar{l}} \bar{K} \xrightarrow{\bar{r}} \bar{R})$. Using the restricted matching morphism $\bar{m} : \bar{L} \rightarrow G$ it can be applied to G resulting in graph \bar{H} . The remainder L' of L gives rise to a new rule $p_G = (L' \xleftarrow{l'} K' \xrightarrow{r'} R_G)$ whose right graph consists of the transformed version of G as well the original right graph R , i.e., it contains both \bar{H} and R as subgraphs. A formal, diagrammatic representation is given in Figure 3(c). An abstract partial application is shown in Figure 3(a) and (b).

Given a not necessarily connected graph G and DPO transformation rule $p = (L \xleftarrow{l} K \xrightarrow{r} R)$, our task is to construct all partial rules obtainable by binding G to p . These partial rules can then be applied to further graphs, allowing for more efficient exploration strategies. The following algorithm enumerates these partial rules:

- 1 For all $l_i \in L$ find the set of all subgraph isomorphisms of l_i to G . That is, find $M_i = \{m \mid m : l_i \rightarrow G \text{ is a subgraph isomorphism}\}$ for $1 \leq i \leq \#L$.
- 2 For all nonempty subsets \bar{L} of L , construct all partial matching morphisms, \bar{m} , by merging morphisms from each $M_j, l_j \in \bar{L}$. Note, that each \bar{m} must be injective.
- 3 For each partial matching morphism, \bar{m} , apply p to G with \bar{m} to obtain a new rule $p_G = (L' \xleftarrow{l'} K' \xrightarrow{r'} R_G)$.

Figure 3 Partial application of some rule $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ to a graph G , with $L = \{l_1, l_2, l_3\}$ and $R = \{r_1, r_2\}$. The partial application is done through a partial matching morphism $\bar{m} : \bar{L} \rightarrow G$ with $\bar{L} = \{l_1, l_2\}$. The application results in a new rule, $p_G = (L' \xleftarrow{l'} K' \xrightarrow{r'} R_G)$ with $L' = \{l_3\}$, for which R is a subgraph of R_G . The transformed graph of G , called \bar{H} , is also a subgraph of R_G . Figure 3(c) is the diagram of subgraph relations for a general partial rule application



The partial matching morphisms constructed from considering $\bar{L} = L$ are actually full matching morphisms, and so the resulting rule has $L' = K' = \emptyset$. In this case p_G represents the creation of R_G from an empty graph, and $G \xrightarrow{p, \bar{m}} R_G$ is a valid derivation. If G is connected, the derivation will additionally be proper.

In the following section we will regard a rule p as a function on sets of graphs, defined provisionally as:

$$p(\mathfrak{U}) = \mathfrak{U} \cup \bigcup_{\substack{G \xrightarrow{p} H \\ G \subseteq \mathfrak{U}}} H.$$

That is, the result of applying p to a set of graphs, \mathfrak{U} , is \mathfrak{U} itself along with all graphs derivable from \mathfrak{U} using p .

3.1 Complex graph states

Consider the problem of applying a rule p twice to a set of graphs \mathfrak{U} . That is, finding $\mathfrak{U}_2 = p(\mathfrak{U}_1)$ for $\mathfrak{U}_1 = p(\mathfrak{U})$. By our definition of rule application we have $\mathfrak{U} \subseteq \mathfrak{U}_1$, so when the algorithm described above is used for evaluating $p(\mathfrak{U}_1)$ it will find not only new derivations but also all derivations found when evaluating $p(\mathfrak{U})$. We therefore use a more complex state than simply sets of graphs. A *graph state* F is defined as a pair of ordered sets of graphs $(\mathcal{U}, \mathcal{S})$ with $\mathcal{S} \subseteq \mathcal{U}$. The elements, \mathcal{U} and \mathcal{S} , will be referred to also as $U(F)$ and $S(F)$ respectively, where U and S are functions on the graph state. In the following we will denote

$U(F)$ as the *universe* of the graph state F and $S(F)$ as the *subset* of the state. The order of graphs in the subset and in the universe is independent and is arbitrary unless otherwise stated.

We define the application of a rule p to a graph state F in the following manner. Let H' be all connected graphs derivable from $U(F)$ with p such that at least one graph from $S(U)$ is being transformed in each derivation:

$$H' = \{h \in H \mid G \xrightarrow{p} H : G \subseteq U(F) \wedge G \cap S(F) \neq \emptyset\}. \quad (2)$$

The result $F' = p(F)$ is such that


$$U(F') = U(F) \cup H' \quad (3)$$

$$S(F') = H' \setminus U(F). \quad (4)$$

That is, the resulting universe contains the input universe and all derived graphs, and the resulting subset contains all new graphs which was not known before. The removal of known graphs from the output subset is motivated by the goal of exploring the underlying network of derivations.

With the definition above we rewrite our initial example as; find $F_2 = p(F_1)$ for $F_1 = p(F)$ and $S(F) = U(F) = \mathcal{U}$. The application $p(F_1)$ can now only discover derivations with at least one graph from $S(F_1)$, which by definition contains only new graphs. Therefore, only new derivations are found. Figure 4 contains a visualisation of the example.

Figure 4 Illustration of the evaluation of $F_2 = p(F_1)$ for $F_1 = p(F)$ and some set of graphs, $S(F) = U(F) = \mathcal{U}$. Each derivation must use at least one graph from the input subset. Two abstract derivations are shown with the endpoints indicating in which sets the graphs are



The implementation utilises the algorithm for transformation by first partially applying the rule to the subset of the input state, and then afterwards the full universe.

4 Strategies

The previous section described how a rule p is applied to a state F to calculate a new state F' , and motivated this by the example of composition of rule application, $F' = p(p(F))$. Using the definition of a graph state, we generalise the interface for rule application into general strategies. A strategy is simply any function Q from and to the set of graph states.

In the following we introduce core strategies defined in the framework. Most of the strategies are parameterised, which we will note with brackets around these parameters. The application of a strategy Q with some fixed parameter, n , to a graph state F is thus denoted as $Q[n](F)$.

4.1 *Parallel*

A parallel strategy is defined in terms of a set of substrategies, $\{Q_1, Q_2, \dots, Q_n\}$. The result of applying a parallel strategy is the union of the results from applying the individual substrategies:


$$F' = \mathbf{parallel}[\{Q_1, Q_2, \dots, Q_n\}](F)$$

$$U(F') = \bigcup_{1 \leq i \leq n} U(Q_i(F))$$

$$S(F') = \bigcup_{1 \leq i \leq n} S(Q_i(F)).$$

A simple use of parallel strategies is to model the possibility of different reaction mechanisms happening simultaneously. As example, consider modelling the formose chemistry which consists of keto-enol tautomerism and aldol addition, both reversible reactions (see Appendix A for the grammar details). Let r_1 and r_2 denote the corresponding reactions from Appendix A, i.e., the enol-to-keto reaction pattern for the carbonyl group and the pattern for aldol addition. The parallel strategy $Q = \mathbf{parallel}[\{r_1, r_2\}]$ thus models that these two reactions can happen simultaneously as illustrated in Figure 5.

Figure 5 Application of a parallel strategy $Q = \mathbf{parallel}[\{r_1, r_2\}]$ to a graph state F , with r_1 being the transformation rule for the enol to keto conversion and r_2 being the transformation rule for aldol addition (see Appendix A). (a) The reaction network with the graph state F consisting of $U(F) = \{\text{formaldehyde, 1,2-ethenediol}\}$ and $S(F) = \{1,2-ethenediol\}$. (b) The reaction network after evaluation of $Q(F)$, with two new molecules; glycolaldehyde and glyceraldehyde. The resulting graph state F' has $U(F') = \{\text{formaldehyde, 1,2-ethenediol, glycolaldehyde, glyceraldehyde}\}$ and $S(F') = \{\text{glycolaldehyde, glyceraldehyde}\}$. In both networks the subset of the graph state is highlighted (see online version for colours)



4.2 Sequence

A sequence strategy, Q , is a composition of a list of substrategies, Q_1, Q_2, \dots, Q_n :


$$Q(F) = Q_n(\dots(Q_2(Q_1(F))\dots)).$$

To increase left-to-right readability of sequence strategies, we will use the notation $Q = Q_1 \rightarrow Q_2 \rightarrow \dots \rightarrow Q_n$. Additionally, if $Q_1 = Q_2 = \dots = Q_n = Q'$, we may use the normal notation for powers of functions, $Q = Q'^n$, for the sequence.

An example of the application of a sequence strategy can be seen in Figure 6, in which two sequential steps of the formose chemistry (parallel strategies) are derived starting from a graph state F with $U(F) = \{\text{formaldehyde, glycolaldehyde}\}$ and $S(F) = \{\text{glycolaldehyde}\}$.

Figure 6 Application of the sequence strategy

$Q = \text{parallel}[\{r_0, r_1, r_2, r_3\}] \rightarrow \text{parallel}[\{r_0, r_1, r_2, r_3\}]$ to the graph state F_0 , with r_i denoting the transformation rules for keto-enol tautomerism and reversible aldol addition. (a) The initial reaction network with F_0 in which $U(F_0) = \{\text{formaldehyde, glycolaldehyde}\}$ and $S(F_0) = \{\text{glycolaldehyde}\}$. (b) The intermediary reaction network after evaluation of the first step of the strategy. The difference in graph state is that 1,2-ethenediol is now added to the universe and subset, while glycolaldehyde no longer is in the subset. (c) The reaction network after complete evaluation of $Q(F_0)$. The final graph state F_2 has all four molecules in the universe and only glyceraldehyde in the subset. Note that in the last step of the strategy the reverse keto-enol reaction is discovered, but glycolaldehyde is already in the universe so it will not be added to the subset of F_2 . The subset of the graph state is highlighted in each network (see online version for colours)



4.3 Repetition

The sequencing strategy only allows composition of a fixed number of strategies, whereas the repetition strategy is used to compose a single strategy with itself many times.

A repetition strategy, Q , is parameterised by a non-negative integer, n , and an inner strategy Q' . The inner strategy is composed with itself until the graph state reaches a fixed point or its subset is empty, however at most n times:

$$Q = \mathbf{repeat}[Q', n] = Q'^k$$

$$k = \min\{0, 1, \dots, n\}, \text{ such that } Q'^k(F) = Q'^{k+1}(F) \vee S(Q'^{k+1}(F)) = \emptyset \vee k = n.$$

This means that if the graph state reaches a fixed point then that graph state is returned, and if the subset of the state becomes empty then the *previous* state is returned. We motivate this condition of a non-empty subset of a produced graph state by our definition of rule application, which requires at least one graph from the subset. By returning the last graph state with non-empty subset the repetition strategy can be used as a precomputation in a sequence to find a kind of closure under some inner strategy.

Note that if $k = 0$ the strategy becomes the identity strategy, i.e., the resulting graph state is the same as the input graph state. If n is set large enough to not limit the repetition, we call it unbounded repetition, and write it as $Q = \mathbf{repeat}[Q']$.

In Figure 6 the strategy for deriving two steps of the formose network is shown. As a generalisation the strategy $Q = \mathbf{repeat}[n, \mathbf{parallel}[\{r_0, r_1\}]]$ can be used to derive (at most) n steps of the network. Figure 7 shows another example using the repetition strategy, where all isomers of glyceraldehyde 3-phosphate (G3P) are generated.

4.4 Revive

Consider the following high-level description of a strategy: Given a single graph g , try to apply the rule p . If the application of p is successful, then let H denote all the produced graphs and return $H \setminus \{g\}$ (all graphs not already known). If the application of p is not successful, then intentionally $\{g\}$ should be returned. The simple strategy $Q = p$ applied to F with $S(F) = U(F) = \{g\}$ only partially achieves this, as illustrated in the following. Let $F' = Q(F)$ be the resulting graph state after evaluation of the strategy on F . Using the definition of the rule application strategy, equations (2)–(4), we get

- $S(F') = H \setminus \{g\}$ and $U(F') = H \cup \{g\}$ if p is successfully applied
- $S(F') = \emptyset$ and $U(F') = \{g\}$ if p can not be applied.

However, the desire was to have $S(F') = \{g\}$ in the unsuccessful case. The intention of the revive strategy is to provide a mechanism to model the desired behaviour. A rule application strategy discovers a (possibly empty) set of derivations. We say that a graph g is *consumed* in a rule application strategy if any of the discovered derivations $G \Rightarrow H$ have $g \in G$. In the natural way we extend this and say that a graph g is consumed by a strategy if it is consumed by any of its substrategies. A revive strategy, $\mathbf{revive}[Q']$, is parameterised by a single substrategy, Q' , and is defined as:

$$F' = \mathbf{revive}[Q'](F)$$

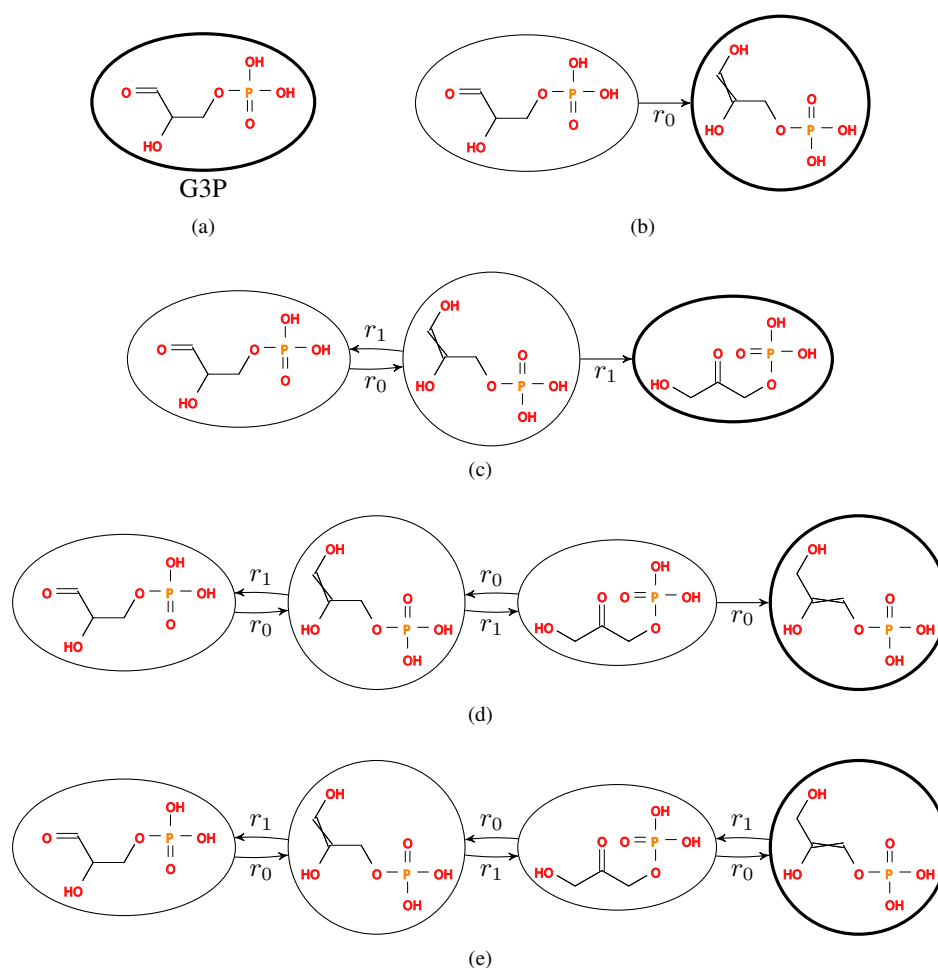
$$U(F') = U(Q'(F))$$

$$S(F') = S(Q'(F)) \cup \{g \in S(F) \mid g \in U(F') \wedge g \text{ is not consumed in } Q'.\}$$

That is, any graph from the input subset which is still in the output universe and was not consumed, will be added to the output subset. The high-level described example to illustrate

the problem with a simple rule p can now be solved with the strategy $Q = \mathbf{revive}[p]$. If the application of p is unsuccessful, then g is not consumed and will be added to the resulting subset.

Figure 7 The strategy $Q = \mathbf{repeat}[\mathbf{parallel}[\{r_0, r_1\}]]$ applied to the initial graph state F_0 with $U(F_0) = S(F_0) = \{\text{G3P}\}$ (shown in (a)). (b)–(d) The intermediary reaction networks from evaluation of $Q(F_0)$. Each step discovers a new isomer which constitutes the new subset. Additionally, the reaction to the previous isomer is discovered. However, this molecule is already in the universe of the current graph state and is therefore not added to the subset. (e) The final step in the repetition results in an empty subset as only known molecules (those in the universe) are rediscovered. The graph state from (d) is therefore the result. In all networks the subset of the current graph state is highlighted (see online version for colours)



As another example, consider the following problem. Two graphs, g_1 and g_2 and the transformation rule p , as illustrated in Figure 8 are given. We wish to develop a strategy to transform all edge labels using rule p , with the intend to use this strategy as a precomputation for a subsequent strategy. That is, the subset of the graph state after evaluation of the strategy

must contain the completely transformed graphs in the subset. The strategy $Q = \mathbf{repeat}[p]$ may seem like the most intuitive approach to model this process. However, the evaluation of $Q(F)$ with $S(F) = U(F) = \{g_1, g_2\}$ does not give the intended result, which is illustrated in Figure 9.

Figure 8 Graphs and transformation rule for the example of the semantics of revive strategies





Figure 9 Illustration of the application of $\mathbf{repeat}[p]$ to F with $S(F) = U(F) = \{g_1, g_2\}$. Only the subset of the graph states are shown. The first application of p results in two new graphs, g_3 and g_4 , but as p can only be applied to g_4 the final subset is only a single graph, g_5 , instead of both g_3 and g_5



The problem is that the repetition strategy will continue as long as *any* new graph can be discovered, and does not preserve the most derived graphs in the subset. Using the strategy $\mathbf{repeat}[\mathbf{revive}[p]]$ correctly solves the problem. A chemical example for the revive strategy will be given in the results section.

4.5 Derivation predicates

For the purpose of precise modelling and the problems with combinatorial explosion it is convenient to limit the possibilities of expansion. We define two variations of the concept of derivation predicates, which both introduce extra constraints in equation (2) to prune unwanted derivations. The strategy $\mathbf{leftPredicate}[P, Q']$ is defined by the predicate P on a multiset of graphs and a transformation rule, and by the substrategy Q' . A candidate derivation from the graphs G with the rule p found by Q' , is only fully calculated and accepted if $P(G, p)$ is true. A right predicate strategy, $\mathbf{rightPredicate}[P, Q']$ is also defined by a predicate and a substrategy, though with the predicate P evaluating a complete derivation. Thus, a derivation $G \xrightarrow{p} H$ is only accepted if $P(G \xrightarrow{p} H)$ is true.

As example, given a strategy Q' we wish to produce only graphs with at most 42 vertices (atoms, in a chemical context). This requires a right predicate strategy as information about the right side of the derivation (the products) are needed. This can be specified with the following strategy:

$$Q = \mathbf{rightPredicate}[P, Q']$$

$$P(G \xrightarrow{r} H) \equiv \forall h \in H : |V(h)| \leq 42.$$

Instead, we might want to restrict that some molecule g should not be an educt in any reaction with the transformation rule being r . This constraint does not require the information of a complete derivation, and may as such be formulated as a left predicate strategy:

$$Q = \mathbf{leftPredicate}[P, Q']$$

$$P(G, p) \equiv \neg(r = p \wedge g \in G)$$

with Q' being an arbitrary strategy.

4.6 Filter, sort, take and add

To facilitate more elaborate use of strategies in a functional style we define several strategies which correspond to functions on lists in other languages. As a graph state is composed of both a universe and a subset, all of these strategies are defined in two variations.

A filter strategy is parameterised by a predicate on a graph and a graph state:

$$F' = \mathbf{filterSubset}[P](F) \qquad F' = \mathbf{filterUniverse}[P](F)$$

$$U(F') = U(F) \qquad U(F') = \{g \in U(F) \mid P(g, F)\}$$

$$S(F') = \{g \in S(F) \mid P(g, F)\} \qquad S(F') = \{g \in S(F) \mid P(g, F)\}.$$

A sorting strategy is parameterised with a predicate on two graphs and a graph state, used as a less-than predicate in a stable sort of a list of graphs:

$$F' = \mathbf{sortSubset}[P](F) \qquad F' = \mathbf{sortUniverse}[P](F)$$

$$U(F') = U(F) \qquad U(F') = \mathbf{stableSort}[P](U(F))$$

$$S(F') = \mathbf{stableSort}[P](S(F)) \qquad S(F') = S(F).$$

The choice that the sorting algorithm must be stable is motivated by the desire to allow lexicographical sorting by sequencing several sorting strategies.

A take strategy is parameterised with a natural number:

$$F' = \mathbf{takeSubset}[n](F) \qquad F' = \mathbf{takeUniverse}[n](F)$$

$$k = \min\{n, |S(F)|\} \qquad k = \min\{n, |U(F)|\}$$

$$U(F') = U(F) \qquad U(F') = \{U(F)_1, U(F)_2, \dots, U(F)_k\}$$

$$S(F') = \{S(F)_1, S(F)_2, \dots, S(F)_k\} \qquad S(F') = S(F) \cap U(F').$$

An addition strategy appends a given set of graphs to either the universe and optionally also to the subset:

$$\begin{array}{ll}
 F' = \mathbf{addSubset}[\{g_1, g_2, \dots, g_n\}](F) & F' = \mathbf{addUniverse}[\{g_1, g_2, \dots, g_n\}](F) \\
 U(F') = U(F) \cup \{g_1, g_2, \dots, g_n\} & U(F') = U(F) \cup \{g_1, g_2, \dots, g_n\} \\
 S(F') = S(F) \cup \{g_1, g_2, \dots, g_n\} & S(F') = S(F).
 \end{array}$$

An example usage of these strategies is the procedure of ranking graphs according to some property, take the best n graphs for subsequence expansion, i.e.:

$$Q' = \mathbf{sortSubset}[P] \rightarrow \mathbf{takeSubset}[n]$$

Note that the sorting predicate P can be based on any external data such as results from wet lab experiments. As example we have used mass spectrometry data to bias the expansion towards high intensity molecules (see the Results and Discussion section).

The addition strategies can be used both for injecting new graphs in the middle of a strategy, but we also find them convenient simply for uniform left-to-right writing of a strategy application. E.g., given a (large) strategy Q we wish to apply to the graph state F , we can write:

$$F' := \mathbf{addUniverse}[U(F)] \rightarrow \mathbf{addSubset}[S(F)] \rightarrow Q$$

with the interpretation $F' = Q(F)$.

4.7 Implementation remarks

The strategies are implemented in C++ as part of a library, to allow easy extension at the user level. Extensions can vary from simple graph state manipulating strategies to complete replacement of the underlying transformation formalism. The library is aimed at chemical graph transformation, with special optimisation for molecules (e.g., use of canonical SMILES strings for graph isomorphism (Weininger, 1988; Weininger et al., 1989)), but is not restricted to the domain of chemistry. The current implementation uses VF2 (Cordella et al., 2004) to find subgraph isomorphisms, and as a fall-back algorithm for isomorphism check for general graphs. Furthermore, the library utilises data structures and procedures for molecule handling from the graph grammar library (GGL) (Benkő et al., 2003). A Python module with bindings to the C++ library is also implemented to allow easy development of expansion strategies.

5 Results

In this section we will apply our strategy framework to three different chemical systems and present results on how to systematically explore complex chemical universes:

- for the Diels-Alder reaction system we will repeatedly merge molecules with isoprene
- we will compare chemical universes of basic formose chemistry with and without using borate as inhibitor motivated by a recent experiment by Ricardo et al. (2004)

- we will present a strategy to explore the complex chemical spaces of hydrogen cyanide polymerisation and hydrolysis product in order to show how to integrate mass spectrometry results in our framework.

In order to easily illustrate subspaces that are also expected to exist in a chemical setting, we will apply the strategy framework to a small puzzle game (Appendix C).


5.1 The Diels-Alder reaction

The Diels-Alder reaction is one of the most useful reactions in organic chemistry and has heavily influenced total synthesis in the last decades (Nicolaou et al., 2002). The explosion of the chemical space by applying this reaction several times will be biased by the strategy framework. The reaction is shown in an example derivation in Figure 1, while the starting molecules, isoprene and cyclohexadiene, are shown in Figure 10. Let $p = (L \leftarrow K \rightarrow R)$ be the transformation rule modelling the Diels-Alder reaction. The intention of the rule is that it is applied to two molecules, but this constraint is not encoded in the rule. We therefore first wrap p with a derivation predicate:

$$Q_p = \mathbf{leftPredicate}[P, p] \quad P(G, p') \equiv \#G = 2.$$

This means that all derivations $G \xrightarrow{p} H$ must have $|G| = 2$.

Figure 10 The starting molecules, (a) isoprene and (b) cyclohexadiene, for application of the Diels-Alder reaction. The molecules are shown in two versions; one with all vertices explicit and chemical interpretation of edge labels (left), and one version in standard chemical visualisation



A generic breadth-first exploration of the chemical space can be done with the following strategy:

$$Q_{\text{BFS}} = \mathbf{addSubset}[\{\text{isoprene, cyclohexadiene}\}] \rightarrow \mathbf{repeat}[Q_p, n].$$

However, for $n = 4$ the strategy already discovers 825 new graphs through 1278 derivations.¹ The number of subgraph isomorphism queries throughout the evaluation is 74591. In Appendix B, Figure 14 the resulting derivation graph for just $n = 2$ is shown.

We now decide to only look at the subspace of molecules which are derived by repeatedly merging molecules with isoprene, starting with cyclohexadine. The following strategy implements this specification:

$$\begin{aligned}
 Q_{\text{subspace}} = & \mathbf{addUniverse}[\{\text{isoprene}\}] \rightarrow \mathbf{addSubset}[\{\text{cyclohexadine}\}] \\
 & \rightarrow \mathbf{leftPredicate}[P_{\text{init}}, Q_p] \rightarrow \mathbf{filterUniverse}[P_{\text{filter}}] \\
 & \rightarrow \mathbf{repeat}[Q_p, n]
 \end{aligned} \tag{5}$$

with

$$\begin{aligned}
 P_{\text{init}}(G, p') & \equiv G = \{\text{isoprene, cyclohexadine}\} \\
 P_{\text{filter}}(g, F) & \equiv g \neq \text{cyclohexadine}.
 \end{aligned}$$

This first computes all possible proper derivations $\{\text{isoprene, cyclohexadine}\} \xrightarrow{p} H$, then removes cyclohexadine from the graph state to prevent further derivations. In the end it uses breadth-first expansion for at most n steps. This strategy, with $n = 3$ (i.e., 4 expansion steps including the very specific first step) discovers only 165 new graphs through 236 derivations,² and uses 5524 subgraph isomorphism queries. The derivation graph with $n = 2$ is visualised in Figure 11.

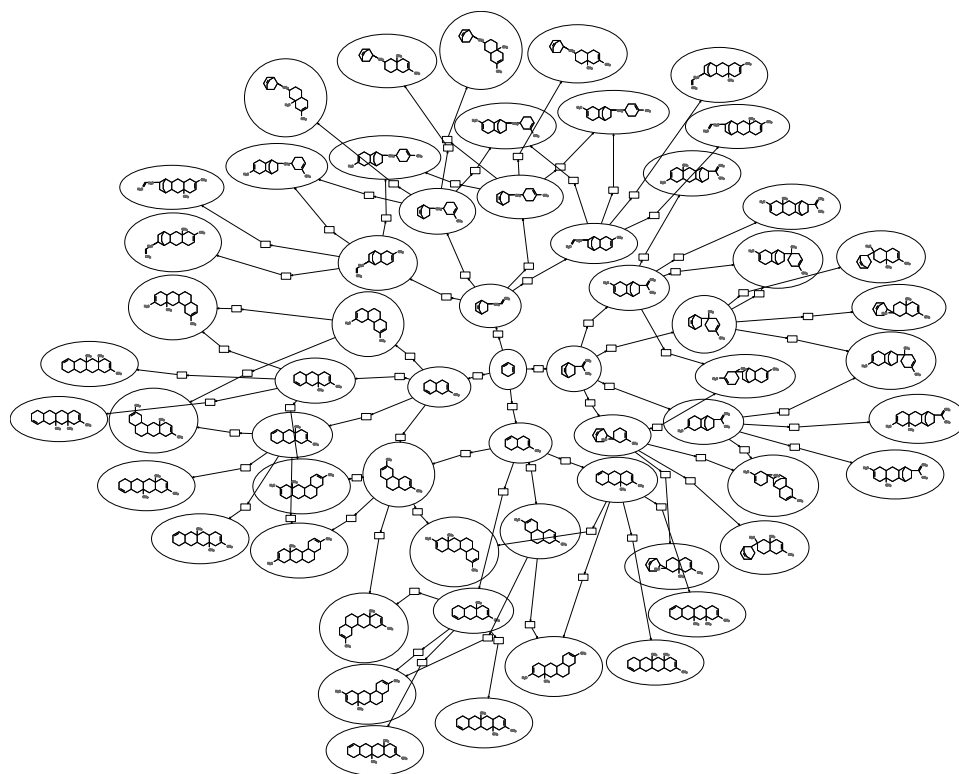
5.2 Borate stabilised Formose reaction

Sugars, or more general carbohydrates, a broad class of organic compounds, can be viewed as polymers of formaldehyde units. The reactivity of carbohydrates is dominated by their carbonyl and their vicinal alcohol functional groups. In particular the enolised form of a carbonyl group may attack another one (in keto form), resulting in the formation of a new carbon-carbon bond. This reaction is known as *aldol addition* (see Figure 6(c)). If the carbon atom adjacent to a carbonyl group carries an alcohol functionality, then the *enolisation reaction* of the carbonyl group erases the ‘information’ at which carbon atom the carbonyl functionality was located before the enolisation. This effect allows the carbonyl group to ‘travel’ along the carbohydrate backbone (see Figure 7(e)). Both reactions are responsible for the meta-stability of carbohydrates and result in complex carbohydrate mixtures when repeated again and again as for instance under the conditions of the formose reaction (Decker et al., 1982). The formose reaction has been extensively discussed as a possible prebiotic route to higher carbohydrates in particular five-carbon sugars, such as ribose, needed for the formation of nucleotides (the building blocks of RNA) (Benner et al., 2012). Unfortunately, if the formose reaction is not stopped in time the reaction mixture turns into black ‘tar’. Therefore, some stabilising mechanism compatible with prebiotic environments, that prevent the destruction of interesting sugars, is indispensable to keep the formose reaction as a plausible prebiotic scenario for higher carbohydrate formation. The addition of borate, capable of binding vicinal diols, to the reaction mixture has been identified as such a stabilising mechanism, that biases the outcome of the formose reaction towards high yields of five-carbon sugars (Ricardo et al., 2004). In the following we illustrate how expansion strategies can be exploited to carve out the differences between the formose reaction networks with and without borate.

The basic formose reaction consists of two types of reversible reaction patterns, keto-enol tautomerism and aldol reaction. As they are reversible they are modelled by two transformation rules each. These are shown in Appendix A as transformation rule r_0, \dots, r_3 ,

while the two initial molecules, formaldehyde and glycolaldehyde, are shown in Figure 12(a) and (b) respectively. To keep the model simple we use a borate-like molecule, Figure 12(c), with just two hydroxyl groups instead of a complete molecule. To enable the formation of borate complexes we use the transformation rule shown in Figure 12(d). This reaction pattern is described in Benner et al. (2010) as inhibiting keto-enol tautomerism by making the hydrogen atoms attached to the carbon atoms non-acidic. To approximate this behaviour we relabel these vertices from H to D, thereby preventing the reaction pattern of enolisation (r_0 in Appendix A) from matching at these locations. The relabeling is done with the reaction 'hToD', Figure 12(e).

Figure 11 The derivation graph resulting from evaluating the expansion strategy Q_{subspace} , equation (5). To minimise clutter, the vertex with isoprene and the corresponding edges are not shown, although isoprene is involved in any reaction (the resulting chemical reaction network is a hypergraph)



The formose chemistry contains an infinite number of molecules, so to limit the scope of the exploration we prune any reaction which creates molecules with more than five carbon atoms. This is formulated with a right predicate strategy around the application of the basic formose reaction patterns:

$$\text{rightPredicate}[P_{\#C}, \text{parallel}[\{r_0, r_1, r_2, r_3\}]]$$

$$P_{\#C}(G \xrightarrow{R} H) \equiv \forall h \in H : h \text{ has at most 5 carbon atoms.}$$

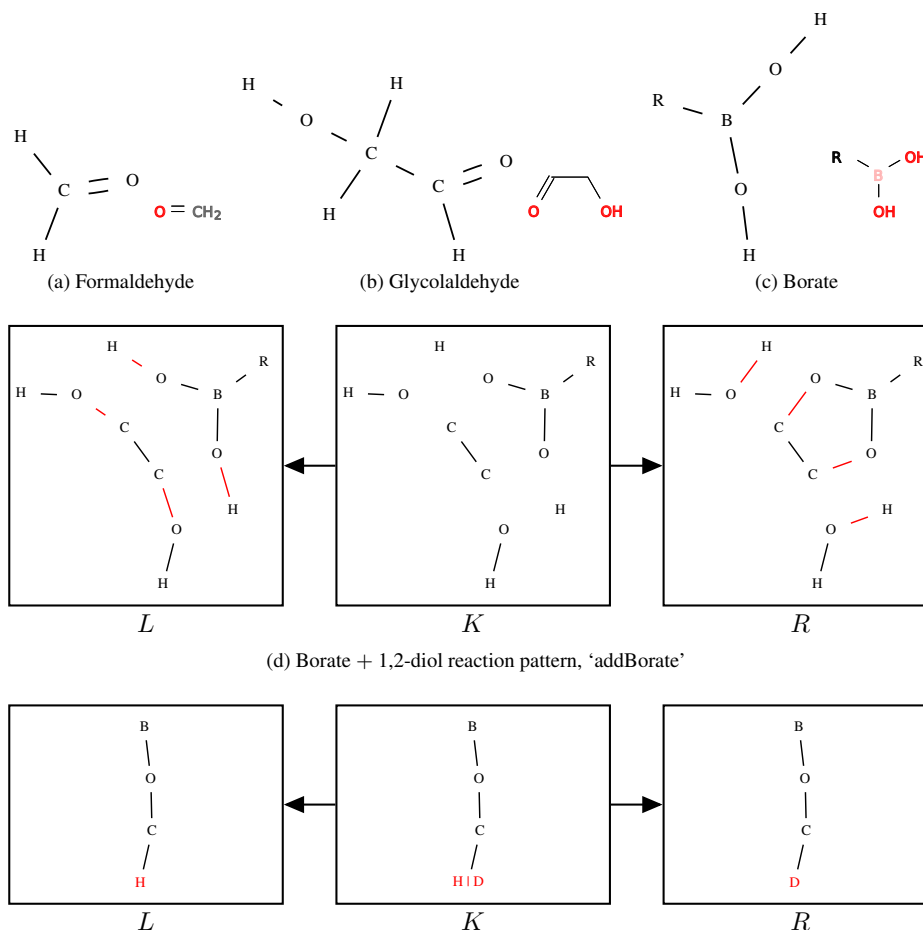
As a reference, we generate the non-inhibited reaction network with the strategy Q_{BFS} :

```

 $Q_{\text{BFS}} = \mathbf{addUniverse}\{\{\text{formaldehyde}\}\}$ 
 $\rightarrow \mathbf{addSubset}\{\{\text{glycolaldehyde}\}\}$ 
 $\rightarrow \mathbf{repeat}\left[ \right.$ 
 $\quad \mathbf{rightPredicate}\left[ P_{\#C}, \mathbf{parallel}\{\{r_0, r_1, r_2, r_3\}\}\right]$ 
 $\left. \right]$ .

```

Figure 12 (a)–(c) The starting molecules of the borate inhibited formose reaction. The three molecules are shown both as explicit graphs with all vertices and in standard chemical visualisation. The borate molecule (c) is modelled with only two hydroxyl groups to simplify the model. (d) The reaction pattern for forming borate complexes with 1,2-diols. This rule additionally has a matching constraint: none of the carbon atoms may be an endpoint of a double bond. To approximate the subsequent non-reactivity of the hydrogens on the carbon atoms we relabel them to D using the reaction ‘hToD’ (e). This relabeling is in the context graph, K , represented with the annotation $\text{H} \mid \text{D}$ (see online version for colours)



Not all molecules can actually bind with borate and must therefore be preserved while the other molecules form complexes. This is modelled with a revive strategy around the actual complex forming reaction pattern, ‘addBorate’. After the potential forming of a borate complex, the relevant hydrogen atoms must be made inactive using the rule ‘hToD’. The number of relevant hydrogens may not be the same for all molecule and therefore the relabeling strategy is embedded in both a repeat and revive strategy. This models the notion of ‘as many times as possible’ on a collection of molecules. The reaction network with borate inhibition can thus be calculated by the following strategy:

$$\begin{aligned}
 Q_{\text{borate}} = & \mathbf{addUniverse}\{\{\text{formaldehyde, borate}\}\} \\
 & \rightarrow \mathbf{addSubset}\{\{\text{glycolaldehyde}\}\} \\
 & \rightarrow \mathbf{repeat}[\\
 & \quad \mathbf{revive}[\text{addBorate}] \\
 & \quad \rightarrow \mathbf{repeat}[\mathbf{revive}[\text{hToD}]] \\
 & \quad \rightarrow \mathbf{rightPredicate}[P_{\#C}, \mathbf{parallel}\{\{r_0, r_1, r_2, r_3\}\}] \\
 & \quad] .
 \end{aligned}$$

Let \mathcal{G} denote the set of molecules used and generated by the evaluation of Q_{borate} on the empty graph state. This set of molecules contain both borate complexes and simple carbohydrates without boron. To canonicalise the molecules we can use the strategy

$$\begin{aligned}
 Q_{\text{canon}} = & \mathbf{addSubset}[\mathcal{G}] \rightarrow \mathbf{repeat}[\mathbf{revive}[\text{dToH}]] \\
 & \rightarrow \mathbf{repeat}[\mathbf{revive}[\text{removeBorate}]]
 \end{aligned}$$

with ‘removeBorate’ being the inverse transformation rule of ‘addBorate’, and ‘dToH’ being the inverse of ‘hToD’. Note that ‘removeBorate’ requires water molecules as educts, but if ‘addBorate’ was ever used in Q_{borate} these molecules must be in \mathcal{G} .

As a variant of the network, we also calculate the network with a an extra molecule, dihydroxyacetone, in the subset:

$$\begin{aligned}
 Q_{\text{borate}}^+ = & \mathbf{addUniverse}\{\{\text{formaldehyde, borate}\}\} \\
 & \rightarrow \mathbf{addSubset}\{\{\text{glycolaldehyde, dihydroxyacetone}\}\} \\
 & \rightarrow \mathbf{repeat}[\\
 & \quad \mathbf{revive}[\text{addBorate}] \\
 & \quad \rightarrow \mathbf{repeat}[\mathbf{revive}[\text{hToD}]] \\
 & \quad \rightarrow \mathbf{rightPredicate}[P_{\#C}, \mathbf{parallel}\{\{r_0, r_1, r_2, r_3\}\}] \\
 & \quad] .
 \end{aligned}$$

In Figure 13 the reference reaction network created with Q_{BFS} is shown. Reactions in black are active only in the basic formose reaction case with formaldehyde and glycolaldehyde as set of input molecules. If borate is added to the input set of molecules, the reactions highlighted in blue are active, while the rest of the network is inactive. Finally if dihydroxyacetone is added to the input set of molecules the reactions highlighted in green are activated in addition to the blue part of the network. The evaluation of Q_{borate} leaves only the blue reactions, which are selective pathways from glycolaldehyde (C2a) to five-carbon sugars (C5b, C5₁, C5₂) active, while the rest of the network is shut down via

borate inhibition. These pathways rely on a constant replenishment of glycolaldehyde. Here dihydroxyacetone (C3k) comes into play. C3k can only be formed from within the formose network via retro-aldol reaction from higher carbohydrates. If added to the reaction network an catalytic loop is activated (sub-network in green: C3k, C3e, C4k, C4e, C5b, retro-aldol red dashed arrow to C3e and C2a) supporting the blue sub-network since C2a ends up as some five-carbon sugars in the blue sub-network. C3e enters another round in the cycle to construct another C2a. These computational results are in very good agreement with the experimental results presented in Kim et al. (2011).

5.3 HCN polymerisation and hydrolysis biased by mass spectrometry results

Hydrogen cyanide (HCN) is a known prebiotic precursor of amino acids as well as many other molecules relevant to present-day biology. It has been used to synthesise adenine already in 1961 (Oró and P., 1961) amino acids (Ferris et al., 1974), as well as many other molecules relevant to present-day biology (Ferris et al., 1974, 1978; Voet and Schwartz, 1983; Miyakawa et al., 2002; Saladino et al., 2004; Borquez et al., 2005; Matthews and Minard, 2006), and it is also known to play a key role also in sugar synthesis (Ritson and Sutherland, 2012). In Andersen et al. (2013a) graph grammar approaches and mass spectrometry results were integrated in order to generate a chemical network with highly likely polymerisation/hydrolysis products. In the first step of the wetlab experiments acid-catalysed HCN polymers were created, in the second step the polymers were hydrolysed under different conditions. The mass spectrometry results of the wetlab experiments were used in order to bias the chemical space exploration performed with the strategy framework. A detailed discussion of the results including a large variety of adenine pathways and autocatalytic processes within the inferred chemical space can be found in Andersen et al. (2013a). Here we focus on the description of the used strategies.

The model of the HCN chemistry is based on many transformation rules which are shown in detail in the web supplement of Andersen et al. (2013a). For the purpose of a concise strategy description we let \mathfrak{R} denote the set of needed transformation rules. The expansion strategy is aimed at modelling the wetlab experiments and thus consist of the sequencing of a strategy for polymerisation with a strategy for hydrolysis. As these two strategies are quite similar we only state the hydrolysis strategy, $Q_{\text{hydrolysis}}$.

Ideally, a simple breadth-first expansion strategy, **repeat**[\mathfrak{R}], can be used to expand the network but due to the sheer combinatorial explosion only very few steps can be calculated. Instead the following strategy can be used to prune the expansion:

$$Q_{\text{hydrolysis}} = \mathbf{addSubset}[\{\text{HCN}, \text{NH}_4, \text{H}_2\text{O}, \text{OH}^-\}]$$

$$\rightarrow \mathbf{repeat} [$$

$$\quad \mathbf{leftPredicate}[P, \mathbf{parallel}[\mathfrak{R}]]$$

$$\quad \rightarrow \mathbf{filterUniverse}[P_{\text{isomer}}]$$

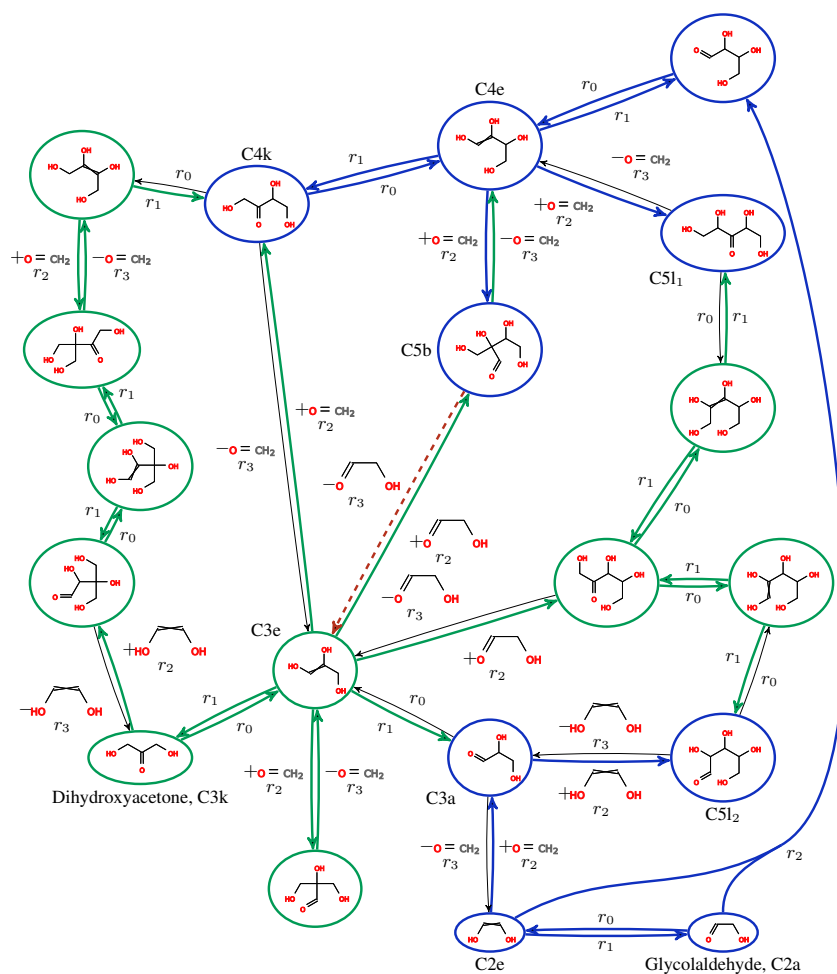
$$\quad \rightarrow \mathbf{sortUniverse}[P_{\text{intensity}}]$$

$$\quad \rightarrow \mathbf{takeUniverse}[20]$$

$$\quad \rightarrow \mathbf{addUniverse}[G_{\text{small}}]$$

$$]$$

Figure 13 The reaction network of the formose chemistry as calculated with the strategy Q_{BFS} . The blue subnetwork correspond to the borate inhibited network calculated with Q_{borate} . The green and blue networks together with the red reaction (C5b to C3e) correspond the network calculated with Q_{borate}^+ , i.e., with dihydroxyacetone as an input compound. Note that this particular model does not include stereochemical properties, and that the molecule depictions are made using Open Babel (O’Boyle et al., 2011), which for instance means crossing double are used to indicate the unspecified stereo. Each reaction is annotated with the reaction pattern, r_i , used to realise the concrete reaction. For the aldol reactions, r_2 and r_3 , the secondary educt (+) or product (-) is additionally shown. The addition of borate in Q_{borate}^+ is done with the strategy **revive**[*addBorate*], meaning that at most 1 borate is added in each iteration. The red reaction is no longer available if the addition is done with the strategy **repeat**[**revive**[*addBorate*]], meaning “add as many as possible” (see online version for colours)



where the predicates are defined as

$$P(G, p) \equiv \text{at most 1 molecule of } G \text{ has molar mass greater than } 50$$

$$P_{\text{isomer}}(g, F) \equiv \text{true iff the normalised Boltzmann factor of } g \text{ is above}$$

a certain threshold.

The factor is calculated based on the isomers of g in $U(F)$

$$P_{\text{intensity}}(g_1, g_2, F) \equiv \text{intensity}(g_1) > \text{intensity}(g_2)$$

The intensities are found in the mass spectrometry data using the molar masses.

That is, the input graph state is augmented with basic food molecules. Then the main hydrolysis step is repeated until no new molecules are found. The main step first expands the network under the constraint that at least one small molecule is an educt in each reaction, which limits the growth of the molecules to be linear as opposed to exponential. The subsequent three steps prune the graph state of unlikely molecules, first by calculating normalised Boltzmann factors within each class of isomers. Then the mass spectrometry data from the wetlab experiments are used to select the 20 molecules with highest intensity for the next expansion step. These pruning steps might have removed the basic food molecules, and they are therefore reintroduced. Additionally the molecules immediately derivable from the food molecules are added. The evaluation of the overall HCN strategy take considerably longer (hours) to calculate than the previous examples. The bulk of the time is however spent on calculating energy value used the Boltzmann factors. For further details see Andersen et al. (2013a).

6 Conclusions

We have introduced here a generic framework to specify and execute strategies for the systematic exploration of spaces of graphs. Our generative approaches use the Double Pushout formalism to derive new graphs. Since this task is of immediate practical relevance in chemistry, we designed our framework and implementation with the aim of high efficiency in this particular domain of application. As performance was a particular focus of our work, we use state-of-the-art subgraph isomorphism check methods and we heavily employ hashing techniques in the checks for graph isomorphism; in order to infer proper derivations of new molecules with full or partial rule application we do *not* use a straightforward method to enumerate all possible left-hand-sides of derivations. Instead we employ partial rule applications, a method that shows theoretically as well as empirically a much better performance. The latter aspect will be discussed in more detail elsewhere.

As showcase examples we have considered complex systems of chemical reactions. For Diels-Alder reactions, which is plagued by a very rapid combinatorial explosion, we used the strategy framework to guide the exploration to emphasise products of repeated isoprene addition instead of unconstrained combinations of reactions products. This is of relevance e.g., in terpene chemistry and biosynthesis. In the case of the formose reaction we show how the strategies framework can be applied to explaining the effects of additional reactants on a given reaction network. In particular, we can in rule based manner also determine which reactions are effectively superseded by new ones, so that additional reactants can lead to a reduction of chemical network. The strategies framework thus serves not only as a convenient tool for exploration but allows also a detailed modelling of constraints in chemical networks.

Although the design was clearly chosen with systems chemistry and systems biology applications in mind, the strategy framework introduced here is however by no means limited

to chemical applications. Another promising application is the emulation of higher-level rules. In the DPO graph grammar formalism, the size of a subgraph that is affected by a transformation is by construction bounded by the left graph of the production that is to be applied. Apparently simple operations on a graph, such as “contract a clique in G to a single vertex”, however, do not have such a bound since the clique sizes depend only on the input graph. Hence, such rules cannot be specified directly as productions in a DPO graph grammar. In the Appendix 6 we use the well-known Catalan game (increate games, 2013) to show how our strategy framework can be applied to emulate this type of higher-level rules.

In order to analyse chemical reaction networks as created by our strategy framework, there exist several mathematical techniques that we plan to apply to our generated networks. Two of the most prominent ones are Flux Balance Analysis (Kauffman et al., 2003) and Elementary Mode Analysis (Klamt and Stelling, 2003). Note, that these methods are usually not applied to dynamically created reaction networks as produced by our framework. We aim at detecting new well-defined chemical reaction pattern. Furthermore, we expect to identify highly connected subgraphs in chemical spaces, that are connected via a small number of bridging reaction, similar to our observation for the Catalan game.

Acknowledgements

This work was supported in part by the Volkswagen Stiftung Proj. No. I/82719, the COST-Action CM0703 ‘Systems Chemistry’, and the Danish Council for Independent Research, Natural Sciences.

References

- Andersen, J., Andersen, T., Flamm, C., Hanczyc, M., Merkle, D. and Stadler, P. (2013a) ‘Navigating the chemical space of hcn polymerization and hydrolysis: guiding graph grammars by mass spectrometry data’, *Entropy*, Vol. 15, No. 10, pp.4066–4083.
- Andersen, J., Flamm, C., Merkle, D. and Stadler, P. (2013b) ‘Inferring chemical reaction patterns using graph grammar rule composition’, *J. Sys. Chem.*, Vol. 4, No. 1, p.4.
- Andrei, O., Fernández, M., Kirchner, H., Melançon, G., Namet, O. and Pinaud, B. (2011) ‘PORGY: strategy driven interactive transformation of graphs’, *Proceedings of the 6th International Workshop on Computing with Terms and Graphs (TERMGRAPH 2011)*, Vol. 48 of *Electronic Proceedings in Theoretical Computer Science*, pp.54–68.
- Benkő, G., Flamm, C. and Stadler, P.F. (2003) ‘A graph-based toy model of chemistry’, *J. Chem. Inf. Comput. Sci.*, Vol. 43, No. 4, pp.1085–1093.
- Benner, S.A., Kim, H-J. and Carrigan, M.A. (2012) ‘Asphalt, water and the prebiotic synthesis of ribose, ribonucleosides, and RNA’, *Acc. Chem. Res.*, Vol. 45, No. 12, pp.2025–2034.
- Benner, S.A., Kim, H-J., Kim, M-J. and Ricardo, A. (2010) ‘Planetary organic chemistry and the origins of biomolecules’, *Cold Spring Harb. Perspect. Biol.*, Vol. 2, pp.a003467.
- Bishop, K., Klajn, R. and Grzybowski, B. (2006) ‘The core and most useful molecules in organic chemistry’, *Angew. Chem. Int. Ed.*, Vol. 45, pp.5348–5354.
- Borquez, E., Cleaves, H.J., Lazcano, A. and Miller, S.L. (2005) ‘An investigation of prebiotic purine synthesis from the hydrolysis of HCN polymers’, *Orig. Life Evol. Biosph.*, Vol. 35, pp.79–90.
- Butlerov, A.M. (1861) ‘Einiges über die chemische structure der körper’, *Zeitschrift für Chemie*, Vol. 4, pp.549–560.

- Cordella, L., Foggia, P., Sansone, C. and Vento, M. (2004) 'A (sub) graph isomorphism algorithm for matching large graphs', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 26, No. 10, p.1367.
- Decker, P., Schweer, H. and Pohlmann, R. (1982) 'Bioids : X. identification of formose sugars, presumable prebiotic metabolites, using capillary gas chromatography/gas chromatography-mass spectrometry of n-butoxime trifluoroacetates on ov-225', *J Chromatogr A*, Vol. 244, pp.281–291.
- Diels, O.P.H. and Alder, K. (1928) 'Synthesen in der hydroaromatischen reihe', *Justus Liebig's Annalen der Chemie*, Vol. 460, pp.98–122.
- Dittrich, P., Ziegler, J. and Banzhaf, W. (2001) 'Artificial chemistries – a review', *Artificial life*, Vol. 7, No. 3, pp.225–275.
- Dow, M., Fisher, M., James, T., Marchetti, F. and Nelson, A. (2012) 'Towards the systematic exploration of chemical space', *Org. Biomol. Chem.*, Vol. 10, pp.17–28.
- Eberhardt, L., Kumar, K. and Waldmann, H. (2011) 'Exploring and exploiting biologically relevant chemical space', *Curr. Drug Targets*, Vol. 12, pp.1531–1546.
- Fernández, M. and Namet, O. (2010) 'Strategic programming on graph rewriting systems', *Proceedings of the 1st International Workshop on Strategies in Rewriting, Proving, and Programming (IWS 2010)*, Vol. 44 of *Electronic Proceedings in Theoretical Computer Science*, pp.1–20.
- Fernández, M., Kirchner, H. and Namet, O. (2012) 'A strategy language for graph rewriting', *Proceedings of the 21st International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2011)*, Vol. 7225 of *Lecture Notes in Computer Science*, pp.173–188.
- Ferris, J.P., Joshi, P.C., Edelson, E.H. and Lawless, J.G.J. (1978) 'HCN: a plausible source of purines, pyrimidines and amino acids on the primitive earth', *J. Mol. Evol.*, Vol. 11, pp.293–311.
- Ferris, J.P., Wos, J.D., Nooner, D.W. and Oró, J. (1974) 'Chemical evolution. XXI. The amino acids released on hydrolysis of HCN oligomers', *J. Mol. Evol.*, Vol. 3, pp.225–231.
- Fialkowski, M., Bishop, K., Chubukov, V., Campbell, C. and Grzybowski, B. (2005) 'Architecture and evolution of organic chemistry', *Angew. Chem. Int. Ed.*, Vol. 44, pp.7263–7269.
- Grzybowski, B., Bishop, K., Kowalczyk, B. and Wilmer, C. (2009) 'The "wired" universe of organic chemistry', *Nature Chemistry*, Vol. 1, pp.31–36.
- increpare games (2013) *Catalan*, Accessed 4 February, 2013.
- Karp, P. and Caspi, R. (2011) 'A survey of metabolic databases emphasizing the MetaCyc family', *Arch. Toxicol.*, Vol. 85, pp.1015–1033.
- Kauffman, K.J., Prakash, P. and Edwards, J.S. (2003) 'Advances in flux balance analysis', *Curr. Opin. Biotechnol.*, Vol. 14, No. 5, pp.491–496.
- Kim, H-J., Ricardo, A., Illangkoon, H.I., Kim, J.K., Carrigan, M.A., Frye, F. and Benner, S.A. (2011) 'Synthesis of carbohydrates in mineral-guided prebiotic cycles', *J. Am. Chem. Soc.*, Vol. 133, No. 24, pp.9457–9468.
- Klamt, S. and Stelling, J. (2003) 'Two approaches for metabolic pathway analysis?', *Trends Biotechnol.*, Vol. 21, No. 2, pp.64–69.
- Matthews, C.N. and Minard, R.D. (2006) 'Hydrogen cyanide polymers, comets and the origin of life', *Faraday Discuss.*, Vol. 133, pp.393–401 and 427–452.
- Miyakawa, S., Cleaves, H.J. and Miller, S.L. (2002) 'The cold origin of life: B. Implications based on pyrimidines and purines produced from frozen ammonium cyanide solutions', *Origins Life Evol. Biosphere*, Vol. 32, pp.209–218.
- Nicolaou, K., Snyder, S. and Montagnon, T. and Vassilikogiannakis, G. (2002) 'The Diels-Alder Reaction in total synthesis', *Angew. Chem. Int. Ed.*, Vol. 41, pp.1668–1698.
- O'Boyle, N., Banck, M., James, C., Morley, C., Vandermeersch, T. and Hutchison, G. (2011) 'Open babel: An open chemical toolbox', *Journal of Cheminformatics*, Vol. 3, No. 1, p.33.

- Oró, J. and Kimball, A.P. (1961) 'Synthesis of purines under possible primitive earth conditions. I. Adenine from hydrogen cyanide', *Arch. Biochem. Biophys.*, Vol. 94, pp.217–227.
- Pinaud, B., Melançon, G. and Dubois, J. (2012) 'PORGY: A visual graph rewriting environment for complex systems', *Comput. Graph. Forum*, Vol. 31, No. 3pt4, pp.1265–1274.
- Reymond, J.-L. and Awale, M. (2012) 'Exploring chemical space for drug discovery using the chemical universe database', *ACS Chem. Neurosci.*, Vol. 3, pp.649–657.
- Ricardo, A., Carrigan, M.A., Olcott, A.N. and Benner, S.A. (2004), 'Borate minerals stabilize ribose', *Science*, Vol. 303, p.196.
- Ritson, D. and Sutherland, J.D. (2012) 'Prebiotic synthesis of simple sugars by photoredox systems chemistry', *Nat. Chem.*, Vol. 4, pp.895–899.
- Rozenberg, G. and Ehrig, H. (1997) *Handbook of Graph Grammars and Computing by Graph Transformation*, Vol. 1, World Scientific, Singapore.
- Ruddigkeit, L., van Deursen, R., Blum, L.C. and Reymond, J.-L. (2012) 'Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17', *Journal of Chemical Information and Modeling*, Vol. 52, No. 11, pp.2864–2875.
- Saladino, R., Crestini, C., Costanzo, G. and DiMauro, E. (2004), 'Advance in the prebiotic synthesis of nucleic acids bases: implications for the origin of life', *Curr. Org. Chem.*, Vol. 8, pp.1425–1443.
- Voet, A. and Schwartz, A. (1983) 'Prebiotic adenine synthesis from HCN-evidence for a newly discovered major pathway', *Bioorg. Chem.* Vol. 12, pp.8–17.
- Weininger, D. (1988) 'SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules', *J. Chem. Inf. Comput. Sci.*, Vol. 28, No. 1, pp.31–36.
- Weininger, D., Weininger, A. and Weininger, J.L. (1989) 'SMILES 2. Algorithm for generation of unique SMILES notation', *J. Chem. Inf. Comput. Sci.*, Vol. 29, No. 2, pp.97–101.
- Wong, Y.-S. (2012) 'Exploring chemical space: recent advances in chemistry', *Chemical Genomics and Proteomics*, Springer, Vol. 800, pp.11–23.
- Zeigarnik, A.V. (2000) 'On hypercycles and hypercircuits in hypergraphs', in Hansen, P., Fowler, P.W. and Zheng, M. (Eds.): *Discrete Mathematical Chemistry*, Vol. 51 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, American Mathematical Society, Providence, RI, pp.377–383.

Notes

¹In this scenario we regard derivations which only differ in the matching morphism as duplicates. The evaluation of the strategy takes in the order of 10 seconds with a Intel® Core™ i5-2500K CPU (3.30 GHz).


²In this scenario we regard derivations which only differ in the matching morphism as duplicates. The evaluation of the strategy takes in the order of 8 seconds with a Intel® Core™ i5-2500K CPU (3.30 GHz).

Appendix A: Transformation rules for the formose chemistry


The main formose chemistry consists of two reversible reactions, keto-enol tautomerism and aldol addition. These reaction patterns are listed below as four transformation rules, r_0 to r_3 , one for each direction. Additionally, for modelling borate inhibition we use a borate addition rule, r_4 . The inverse of this rule, r_5 , is used for generating the underlying molecule without borate. The rules r_6 and r_7 are used for converting between acidic and non-acidic

hydrogens in borate complexes. Note that the context graph, K , of r_6 and r_7 also uses the labeling scheme “L label | R label”, with the meaning that the vertex changes label from ‘L label’ to ‘R label’.


A.1 r_0 , keto-enol tautomerism, keto-to-enol (see online version for colours)




A.2 r_1 , keto-enol tautomerism, enol-to-keto (see online version for colours)




A.3 r_2 , aldol reaction, addition (see online version for colours)





A.4 r_3 , aldol reaction, splitting (see online version for colours)



A.5 r_4 , borate reaction, addition (see online version for colours)

The rule has the following matching condition: none of the adjacent edges of the carbon vertices may represent a double bond.

A.6 r_5 , borate reaction, splitting (see online version for colours)A.7 r_6 , acidic to non-acidic hydrogen (see online version for colours)


A.8 r_7 , non-acidic to acidic hydrogen (see online version for colours)**Appendix B: Additional Diels-Alder chemistry figure**

Figure 14 shows the derivation graph obtained from the breadth-first expansion of the Diels-Alder chemistry. The number of expansion steps is only 2.

Appendix C: solving the catalan game


The Catalan game (increpare games, 2013) is a puzzle game in which the player in each level is presented with a simple undirected graph without labels. The goal is to transform the graph into a single vertex using the following rewriting rule; given a vertex v with degree exactly 3, identify v with its neighbours and preserve simpleness of the graph by identifying parallel edges and deleting loops. Figure 15 shows level 1 with the intermediary graphs towards the goal graph with a single vertex.

The transformation in the game can not be formulated as a single rule in the DPO formalism, because such rules must explicitly match the vertices and edges which are changed, while the Catalan transformation needs to change arbitrarily many edges. In the following we show how the strategies can be used to implement a move in the game, using only DPO rules.

Let g be the graph from some Catalan level, with all edge labels set to the empty string and all vertex labels set to the arbitrarily chosen label '0'. A high-level description of a move is:

- 1 find a vertex v with at least 3 neighbours and mark it by changing the label to 'A'. Mark the 3 matched neighbours with the label 'R'
- 2 if possible, find another fourth neighbour of v and mark v with 'FAIL'
- 3 discard all graphs with a vertex with the label 'FAIL'
- 4 for all edges e with both end-vertices having label 'R', remove e
- 5 for all edges ur with u having label '0' and r having label 'R', add uv if it does not exist already and then remove ur
- 6 for all edges ur with u having label '0' and r having label 'R', remove ur
- 7 remove all neighbours of v having label 'R'
- 8 unmark v by changing the label to '0'.

Figure 14 The derivation graph resulting from evaluating the breadth-first expansion strategy $Q_{\text{BFS}} = \text{addSubset}[\{\text{isoprene, cyclohexadine}\}] \rightarrow \text{repeat}[Q_p, 2]$ (on an empty graph state). To minimise clutter, the vertex with isoprene and the corresponding edges are not shown



Step 3 can be implemented with a filtering strategy while the other steps each require a transformation rule. The following strategy can be used to solve a level, in the sense that if a graph with a single vertex with label '0' is found, then a path to that graph is equivalent to a solution. The details of the transformation rules (mark, markForFail, removeInterR, reattachExternal, removeAttached, removeR and unmark) are shown in Appendix D.

$$Q_{\text{catalan}} = \text{addSubset}[\{g\}] \rightarrow \text{altRuleApp}[\text{repeat}[\text{mark} \rightarrow \text{revive}[\text{markForFail}] \rightarrow \text{filterUniverse}[P_{\text{fail}}] \rightarrow \text{repeat}[\text{revive}[\text{removeInterR}]] \rightarrow \text{repeat}[\text{revive}[\text{reattachExternal}]] \rightarrow \text{repeat}[\text{revive}[\text{removeAttached}]] \rightarrow \text{removeR} \rightarrow \text{unmark}]]]$$

$P(g', F) \equiv$ no vertex of g' has the label 'FAIL'.

With strategy Q_{catalan} all 56 levels of Catalan could be solved, all but one level took less than 10 minutes of computation time. Figure 16(b) exemplarily shows the derivation graph created when executing the strategy with g encoding level 25 of the game, and Figure 16(a) show the initial level graph. The resulting derivation graph is, in contrast to chemical reaction networks, not a hypergraph. However, the graph clearly illustrates subspaces that are connected via a small number of bridging edges. Such subspaces are also expected in chemical reaction networks.

Figure 15 Level 1 of the Catalan game and the intermediary graphs during transformation to a graph with a single vertex





Figure 16 The derivation graph created during expansion of level 25 of the Catalan game. A path equivalent to a solution is highlighted (see online version for colours)




Appendix D: Transformation rules for the Catalan game


The following sections contain visualisation of the rules used in the strategy to solve a level in the Catalan game. Vertices and edges shown in red are those being changed during

transformation. For some vertices the change is only a change of label. The label in the context graph, K , is for those in the format 'L | R' with L and R being the label in the left and right side of the rule.


D.1 mark (see online version for colours)




D.2 markForFail (see online version for colours)




D.3 removeInterR (see online version for colours)




D.4 reattachExternal (see online version for colours)



D.5 removeAttached (see online version for colours)



D.6 removeR (see online version for colours)



D.7 unmark (see online version for colours)

