



UNIVERSITY OF PISA

Department of Computer Science

Master Degree in Computer Science

Modeling Reaction Systems

Supervisor:

Prof. Roberto Bruni
Prof. Roberta Gori
Prof. Paolo Milazzo

Candidate:

Elvis Rossi

ACADEMIC YEAR 2025/2026



Contents

1	Introduction	1
1.1	Software Design and Key Features	1
2	Background	3
2.1	Reaction Systems	3
2.1.1	Example: A binary counter	5
2.1.2	Simple loops	5
2.2	SOS rules for reaction systems	6
2.3	Positive Reaction Systems	8
2.3.1	From RSs to Positive RSs	8
2.3.2	Minimization	9
2.3.3	Slicing	9
2.4	Bisimulation	10
2.4.1	Algorithms for evaluating bisimulation	11
3	Design	19
3.1	ReactionSystems	19
3.1.1	Entities and Translator	20
3.1.2	Set	20
3.1.3	Reaction	21
3.1.4	Process	22
3.1.5	Choices	23
3.1.6	Environment	23
3.1.7	System	24
3.1.8	Label	26
3.1.9	Graph	27
3.1.10	Slicing Trace	30
3.1.11	Bisimilarity and Bisimulation	30
3.1.12	Grammar and Separated Grammar	33
3.1.13	Experiments and Frequency	33
3.1.14	Instructions	34
3.2	ReactionSystemsGUI	36

4 Development	43
4.1 ReactionSystems	43
4.1.1 Entities and Translator	43
4.1.2 Set	44
4.1.3 Reaction	44
4.1.4 Process, Choices and Environment	44
4.1.5 System	45
4.1.6 Label	45
4.1.7 Graph	46
4.1.8 Slicing Trace	47
4.1.9 Bisimilarity and Bisimulation	47
4.1.10 Assert	47
4.1.11 Grammar	48
4.2 ReactionSystemsGUI	49
5 Testing and Validation	51
5.1 Tests	51
5.2 Validation	57
6 Conclusion	59
6.1 Summary of Contributions	59

Chapter 1

Introduction

Reaction Systems (RSs) are a successful computational framework inspired by biological system. The interaction between biochemical reactions and the functioning of single reactions are based on the mechanisms of facilitation and inhibition, which can be modeled and analyzed using RS.

In this work new software for modeling, analyzing and designing Reaction Systems is designed and developed, with focus on performance and user interface design.

A Reaction System consists of a set of entities and a set of reactions over them. Each reaction produces some set of entities P (called products) if enabled, meaning if a set R (called reactants) is wholly present and if a set I (called inhibitors) of entities is completely absent. The use of inhibitors induces non-monotonic behaviors that are difficult to analyze. Entities can also be provided by an external context sequence to simulate *in silico* biological experiment, expanded by structural operational semantics (SOS) rules to account for several biological experiments. In addition Positive RS, trace slicing, graph generation, bisimulation and more is available through an intuitive visual language with a graphical interface.

Despite the rich theoretical development of Reaction Systems, practical tools for working with RS models have lagged behind. Numerous versions of Prolog and Python programs have been developed[4], but have problems regarding performance and usability. Future developments may be hindered by this lack of software by this technological barrier.

This thesis aims to bridge the gab between the theoretical foundations of Reaction Systems and their practical application. To achieve this a new software platform for modeling, analyzing and designing Reaction Systems is proposed.

1.1 Software Design and Key Features

The core contribution here presented is a new software tool built from the ground up to support Reaction Systems modeling and analysis. Equal emphasis has been placed on performance and user experience. The software is implemented in Rust[27], a modern systems programming language chosen for its efficiency and reliability. Rust's strong per-

formance characteristics (memory safety, speed, and concurrency support) help ensure that even larger Reaction System models can be analyzed quickly, while its emphasis on code safety and clarity makes the tool more maintainable in the long term.

The platform provides two user interfaces to accommodate different user needs. A command-line interface (CLI) is available for quick integration in already existing pipelines. The CLI allows specification of Reaction Systems and instructions over them and is easily expandable to meet the need of the programmer.

A graphical user interface (GUI) is also available both as standalone native application and as a static web application running on WebAssembly[18]. The GUI lowers the learning curve for new users: instead of writing code or scripts, one can construct reactions, run simulations and view results through the same interface through interactive diagrams and controls. By providing both CLI and graphical native/web interfaces, the tool caters to a wide audience.

Key capabilities of the developed Reaction Systems software include:

- Graphical interface for modeling and simulation: An interactive GUI that allows users to graphically define RS components and simulate their behavior.
- Trace slicing: tools for examining execution traces in detail. Trace slicing allows a user to isolate and inspect specific segments of a reaction sequence, exploring causality between produced elements or inhibited reactions.
- Bisimulation analysis: support for formal bisimulation analysis, which enables comparing different Reaction Systems models for behavioral equivalence. The methods are not restricted to just analyzing RS, but are available for any graph-like structure.
- Conversion between Reaction Systems and Positive Reaction Systems: to better explore traces and causality, a more suitable model is provided. Conversion between systems is handled automatically and the defaults can be fine tuned or overridden.

Together, these features make the software a comprehensive environment for working with Reaction Systems. The user can construct, simulate and analyze the results through multiple lenses: from observing cyclic behaviors to checking formal equivalences and visualizing interaction networks.

Chapter 2

Background

2.1 Reaction Systems

Reaction Systems are a qualitative model inspired by biochemical processes. The behavior is described by reactions, each of them requiring some reactants R and requiring the absence of inhibitors I to produce some product elements P . “Elements” and “entities” will be used interchangeably to refer to elements of these sets.

Definition (Reaction). A reaction is a triplet $a = (R, I, P)$, where R, I, P are finite sets with $R \cap I = \emptyset$ and $R, I, P \neq \emptyset$. If S is a set such that $R, I, P \subseteq S$, then a is a reaction in S .

The reactions (R, I, P) operate over a finite set of entities S , called *background set*. The theory of RSs is based on three assumptions:

- no permanency, meaning entities vanish unless sustained by a reaction;
- no counting, meaning the exact quantity of each entity is irrelevant;
- threshold nature of resources, meaning if an entity is present, it is present for all possible reactions.

Definition. Let T be a finite set.

1. Let a be a reaction. Then a is enabled by T , denoted by $en_a(T)$ if $R_a \subseteq T$ and $I_a \cap T = \emptyset$. The result of a on T , denoted by $res_a(T)$, is defined by: $res_a(T) := P_a$ if $en_a(T)$ and $res_a(T) := \emptyset$ otherwise.
2. Let A be a finite set of reactions. The result of A on T , denoted by $res_A(T)$, is defined as $res_A(T) := \bigcup_{a \in A} res_a(T)$.

Note that by virtue of the second assumption (2) if two reactions $a, b \in A$, with both a and b enabled by T , then even if $R_a \cap R_b \neq \emptyset$, still both $P_a \subseteq res_A(T)$ and $P_b \subseteq res_A(T)$. Both reactions can use $R_a \cap R_b$ to produce their products. This would not be allowed in other models such as Petri nets [22], a common model of concurrent systems.

Let $rac(S)$ be the set of all the reactions in S .

Definition (Reaction System). A Reaction System (RS) is a pair $\mathcal{A} = (S, A)$ such that S is a finite set and $A \subseteq \text{rac}(S)$ is a finite set of reactions in S .

The set S is called the *background set* of \mathcal{A} ; its elements, called *entities*, represent molecular entities (e.g. atoms, ions, molecules) that may be present in the state of the system modeled by \mathcal{A} . The set A is called the *set of reactions* of \mathcal{A} . Since S is finite, so is A .

the behavior of a RS is formalized through the notion of an interactive process.

Definition (Interactive Process). Let $\mathcal{A} = (S, A)$ be a RS and let $n > 0$. An n -step interactive process in \mathcal{A} is a pair $\pi = (\gamma, \delta)$ such that $\gamma := \{C_i\}_{i \in [0, n]}$ is the context sequence and $\delta := \{D_i\}_{i \in [0, n]}$ is the result sequence, where $\forall i \in [0, n], C_i, D_i \subseteq S$, $D_0 = \emptyset$ and $\forall i \in [0, n-1], D_{i+1} := \text{res}_A(D_i \cup C_i)$. We call $\tau := W_0, \dots, W_n$ the state sequence, where $W_i := C_i \cup D_i$ for all $i \in [0, n]$.

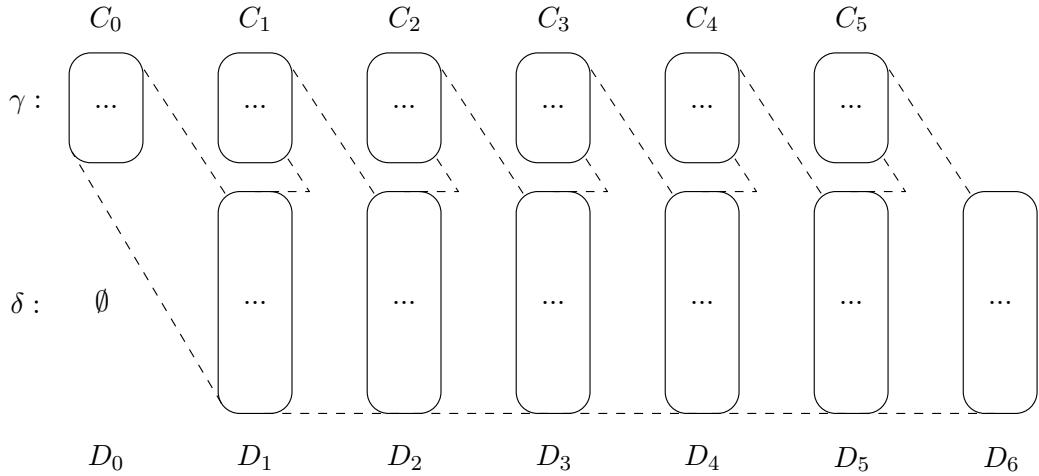


Figure 2.1: Interactive process of a reaction system.

Note that C_i and D_i do not have to be disjointed.

$W_0 = C_0$ is called the initial state of π . If $C_i \subseteq D_i$ for all $i \in [1, n]$ then π is called context-independent. For context-independent interactive process, we can take $C_i = \emptyset$ for all $i = [1, n]$ without changing the state sequence.

In a context-independent state sequence $\tau = W_0, \dots, W_i, W_{i+1}, \dots, W_n$, during the transition from W_i to W_{i+1} all entities from $W_i - \text{res}_A(W_i)$ will not persist. This reflects the assumption of no permanency (1). Thus, if τ is not context-independent, an entity from a current state can also be sustained by the context C_{i+1} .

Definition (Sequence Shift). Let $\gamma = \{C_i\}_{i \in [0, n]}$ a context sequence. Given a positive integer $k \leq n$, let $\gamma^k := \{C_{i+k}\}_{i \in [0, n-k]}$.

2.1.1 Example: A binary counter

A reaction system can act as a cyclic n-bit counter in which external signals trigger increment or decrement operations. To build the counter, let $n > 0$ be an integer and define the background set as $\{p_0, p_1, \dots, p_{n-1}\} \cup \{dec, inc\}$.

Five sets of reactions describe a binary counter-like behavior:

$$\begin{aligned}
 a_j &= (\{p_j\}, \{dec, inc\}, \{p_j\}), \forall j \in [0, n] \\
 b_j &= (\{inc, p_0, p_1, \dots, p_{j-1}\}, \{dec, p_j\}, \{p_j\}), \forall j \in [0, n] \\
 c_{j,k} &= (\{inc, p_k\}, \{dec, p_j\}, \{p_k\}), \forall j, k \text{ s.t. } 0 \leq j < k < n \\
 d_j &= (\{dec\}, \{inc, p_0, p_1, \dots, p_j\}, \{p_j\}), \forall j \in [0, n] \\
 e_{j,k} &= (\{dec, p_j, p_k\}, \{inc\}, \{p_k\}), \forall j, k \text{ s.t. } 0 \leq j < k < n
 \end{aligned}$$

where reactions a cause the bits to be restrained in the next state if there is no operation, reactions b implement the increment operation by flipping the least significant zero to one, reactions c let the more significant bits remain, reactions d implements the decrement operation by flipping to one the bits when there is no one at a lower position, and reactions e let the more significant bits remain.

The complete RS \mathcal{B}_n is defined as follows: $\mathcal{B}_n = (S_n, B_n)$ where

$$S_n = \{p_0, p_1, \dots, p_{n-1}\} \cup \{dec, inc\}$$

and

$$B_n = \{a_j, b_j, d_j \mid 0 \leq j < n\} \cup \{c_{j,k}, e_{j,k} \mid 0 \leq j < k < n\}$$

To illustrate the system in action consider the sequence of contexts: $C_0 = \{p1, p3\}$, $C_1 = \emptyset$, $C_2 = \{inc\}$, $C_3 = \{inc\}$, $C_4 = \{dec\}$, $C_5 = \{dec, inc\}$. This gives the result sequence $\delta = \emptyset \cdot \{p1, p3\} \cdot \{p1, p3\} \cdot \{p0, p1, p3\} \cdot \{p2, p3\} \cdot \{p0, p1, p3\} \cdot \emptyset$ and state sequence $\tau = \{p1, p3\} \cdot \{p1, p3\} \cdot \{p1, p3, inc\} \cdot \{p0, p1, p3, inc\} \cdot \{p2, p3, dec\} \cdot \{p0, p1, p3, dec, inc\} \cdot \emptyset$ that in binary representation is $\{1010_2\} \cdot \{1010_2\} \cdot \{1010_2\} \cdot \{1011_2\} \cdot \{1100_2\} \cdot \{1011_2\} \cdot \{0000_2\}$ by ignoring inc and dec .

2.1.2 Simple loops

The result sequence of a RS, generated by applying the reactions, can be different from \emptyset for an infinite number of results. But if two results are the same, also the following result will be the same. This implies the existence of a loop of states. To identify a loop one can simply truncate the trace at position n and search for the result in position $n+1$ in the truncated trace.

If the entities provided by the context are not constant, the behavior might not be as easily described. By restricting only on a constant set provided by the context, we can be sure that the computation will find a loop.

These loops are called *lollipops*.

For example by providing the set $\{inc\}$ as the context at each step of the system from Example 2.1.1, we obtain the loop: $\{p1, p3\} \cdot \{p0, p1, p3\} \cdot \{p2, p3\} \cdot \{p0, p2, p3\} \cdot \{p1, p2, p3\} \cdot \{p0, p1, p2, p3\} \cdot \{\} \cdot \{p0\} \cdot \{p1\} \cdot \{p0, p1\} \cdot \{p2\} \cdot \{p0, p2\} \cdot \{p1, p2\} \cdot \{p0, p1, p2\} \cdot \{p3\} \cdot \{p0, p3\}$.

2.2 SOS rules for reaction systems

The behavior of a RS could be defined as a discrete time interactive process: a finite context sequence describes the entities provided by the environment at each step, the current state is determined by the union of the entities coming from the environment with those produced from the previous step and the state sequence is determined by applying all and only the enabled reactions to the set of entities available in the current state.

Given the context sequence, the semantics of RSs is uniquely determined and can be represented as a finite, deterministic and labeled or unlabeled transition system. RS have had defined a Labeled Transition System (LTS) semantics as seen in[7].

Definition (RS processes). Let S be a set of entities. An RS process P is any term defined by the following grammar:

P	$::=$	$[M]$	mixture process
M	$::=$	(R, I, P)	reaction
		D	set of entities
		K	context process
		$M M$	parallel composition
K	$::=$	0	nil context
		X	process variable
		$C.X$	set of entities
		$K + K$	followed by context
		$\text{rec}X.K$	non deterministic choice
			recursive operator

Where

$$R, I, P \subseteq S \text{ non empty sets of entities}$$

$$C, D \subseteq S \text{ possibly empty set of entities}$$

$$X \text{ a process variable}$$

An RS process P embeds a mixture process M obtained as the parallel composition of some reactions (R, I, P) , some set of currently present entities D (possibly the empty set \emptyset), and some context process K . For brevity sake $M_1|M_2|M_3 = \prod_{i \in \{1,2,3\}} M_i$. A process context K is a possibly nondeterministic and recursive system: the nil context 0 stops the computation; the prefixed context $C.K$ provides the entities in C and then uses K as the next context; the non deterministic choice $K_1 + K_2$ allows the context to behave either as K_1 or K_2 ; X is a process variable; and $\text{rec}X.K$ is the recursive operator of process algebras. For brevity sake $K_1 + K_2 + K_3 = \sum_{i \in \{1,2,3\}} K_i$.

Definition (RSs as RS processes). Let $\mathcal{A} = (S, A)$ be a RS, and $\pi = (\gamma, \delta)$ an n -step interactive process in \mathcal{A} with $\gamma = \{C_i\}_{i \in [0, n]}$ and $\delta = \{D_i\}_{i \in [0, n]}$. For any step $i \in [0, n]$, the corresponding RS process $[\mathcal{A}, \pi]_i$ is defined as follows:

$$[\mathcal{A}, \pi]_i := \left[\prod_{a \in A} a | D_i | K_{\gamma^i} \right]$$

where the context process $K_{\gamma^i} := C_i.C_{i+1} \dots C_n.\mathbf{0}$ is the sequentialization of the entities offered by γ^i .

Definition (Label). A label is a tuple $\langle W \triangleright R, I, P \rangle$ with $W, R, I, P \subseteq S$.

In a transition label $\langle W \triangleright R, I, P \rangle$, W records the set of entities currently in the system (produced in the previous step or provided by the context), R records the set of entities whose presence is assumed (either because they are needed as reactants on an applied reaction or because their presence prevents the application of some reaction), I records the set of entities whose absence is assumed, and P records the set of entities produced by the applied reactions.

Definition (Operational Semantics). The operational semantics of processes is defined by the set of SOS inference rules in figure 2.2.

$$\begin{array}{c}
\frac{}{D \xrightarrow{\langle D \triangleright \emptyset, \emptyset, \emptyset \rangle} \emptyset} \text{(Ent)} \quad \frac{}{C.K \xrightarrow{\langle C \triangleright \emptyset, \emptyset, \emptyset \rangle} K} \text{(Ctx)} \quad \frac{K \xrightarrow{\text{rec}X.K/X} K'}{K \xrightarrow{\langle W \triangleright R, I, P \rangle} K'} \text{(Rec)} \\
\\
\frac{K_1 \xrightarrow{\langle W \triangleright R, I, P \rangle} K'_1}{K_1 + K_2 \xrightarrow{\langle W \triangleright R, I, P \rangle} K'_1} \text{(Suml)} \quad \frac{K_2 \xrightarrow{\langle W \triangleright R, I, P \rangle} K'_2}{K_1 + K_2 \xrightarrow{\langle W \triangleright R, I, P \rangle} K'_2} \text{(Sumr)} \\
\\
\frac{}{(R, I, P) \xrightarrow{\langle \emptyset \triangleright R, I, P \rangle} (R, I, P)|P} \text{(Pro)} \quad \frac{J \subseteq I \quad Q \subseteq R \quad J \cup Q \neq \emptyset}{(R, I, P) \xrightarrow{\langle \emptyset \triangleright R, I, P \rangle} (R, I, P)|P} \text{(Inh)} \\
\\
\frac{M_1 \xrightarrow{\langle W_1 \triangleright R_1, I_1, P_1 \rangle} M'_1 \quad M_2 \xrightarrow{\langle W_2 \triangleright R_2, I_2, P_2 \rangle} M'_2 \quad (W_1 \cup W_2 \cup R_1 \cup R_2) \cap (I_1 \cup I_2)}{M_1|M_2 \xrightarrow{\langle W_1 \cup W_2 \triangleright R_1 \cup R_2, I_1 \cup I_2, P_1 \cup P_2 \rangle} M'_1|M'_2} \text{(Par)} \\
\\
\frac{M \xrightarrow{\langle W \triangleright R, I, P \rangle} M' \quad R \subseteq W}{[M] \xrightarrow{\langle W \triangleright R, I, P \rangle} [M']} \text{(Sys)}
\end{array}$$

Figure 2.2: SOS semantics of the reaction system process

$K \xrightarrow{\text{rec}X.K/X}$ denotes the process obtained by replacing in K every free occurrence of the variable X with its recursive definition $\text{rec}X.K$. The rule (Pro), executes the reaction (R, I, P) (its reactants, inhibitors, and products are recorded the label), which remains available at the next step together with P . The rule (Inh) applies when the reaction (R, I, P) should not be executed; it records in the label the possible causes for which the reaction is disabled: possibly some inhibiting entities ($J \subseteq I$) are present or some reactants ($Q \subseteq R$) are missing, with $J \cup Q \neq \emptyset$, as at least one cause is needed for

explaining why the reaction is not enabled.

2.3 Positive Reaction Systems

A particular kind of Reaction Systems, first seen in [Brodo_Bruni_Falaschi_Gori_Milazzo_Montag] are those without inhibitors. Such reactions are called positive and can be simply written as pairs (R, P) and are equivalent to (R, \emptyset, P) . One can always encode any standard RS $\mathcal{A} = (S, A)$ into an equivalent one without inhibitors. In order to track the absence of entities, a new “negative” entity is added for each original one. In any meaningful state $W = D \cup C$ there will always be either one between a and \bar{a} , but never both. As a consequence, for any entity $a \in S_C$, we must assume that the context will provide either a or \bar{a} . Define $\mathbf{S} := S \uplus \bar{S}$ and $\bar{S} := \{\bar{a} | a \in S\}$. A subscript D or C will be used to differentiate between entities related to reaction products and related to the context.

Definition (State consistency). A set $\mathbf{W} \subseteq \mathbf{S}$ is non-contradictory if for all entities $a \in S$ it holds that $\{a, \bar{a}\} \not\subseteq \mathbf{W}$. A non-contradictory state $\mathbf{W} \subseteq \mathbf{S}$ is consistent if, for any entity $a \in S$, either $a \in \mathbf{W}$ or $\bar{a} \in \mathbf{W}$ holds.

Definition (Positive RS[8]). A Positive RS is a Reaction System $\mathcal{A}^+ = (\mathbf{S}, A)$ that satisfies the following conditions:

1. Each reaction r in A is positive, i.e., $r = (\mathbf{R}, \emptyset, \mathbf{P})$ for some non-contradictory sets \mathbf{R} and \mathbf{P} .
2. Consistency preservation: for any consistent state \mathbf{W} , the result set $res_{\mathcal{A}^+}(\mathbf{W})$ must be consistent.

If one assumes that the initial state $\mathbf{D}_0 \subseteq \mathbf{S}_D$ is a non-contradictory set and that the sets $\mathbf{C}_0, \dots, \mathbf{C}_n \subseteq \mathbf{S}_C$ provided by the context are non-contradictory sets, the second condition of 2 guarantees that all result states traversed by the computation will be consistent as well.

2.3.1 From RSs to Positive RSs

For each standard RS $\mathcal{A} = (S, A)$ it is possible to construct a Positive RS $\mathcal{A}^+ = (\mathbf{S}, A^+)$ that exactly mimic the behavior of \mathcal{A} . The reactions in A^+ can be split in two categories: A_{pos}^+ that simply embeds the original reactions A and A_{neg}^+ whose reactions serve for negative entities bookkeeping.

For each reaction $(R, I, P) \in A$ there will be one positive reaction $(R \cup \bar{I}, P) \in A_{pos}^+$. Extra reactions are needed to track the absence of the products of the original reactions. They will be produced whenever no reaction in A that produces a is enabled. For this purpose, assume to collect all reactions in A that are capable of producing a : to ensure that none of them are enabled, we must make sure that, for each one, at least one reactant is absent or at least one inhibitor is present.

Definition (Prohibiting set). Let $\mathcal{A} = (S, A)$ be RS and $a \in S_D$ one of its entities. A non-contradictory set $\mathbf{T} \subseteq \mathbf{S}$ is a prohibiting set for a if for any reaction $(R, I, P \cup a) \in A$ we have that $\mathbf{T} \cap (I \cup \bar{R}) \neq \emptyset$. Denote the set of prohibiting sets for a with $Proh_{\mathcal{A}}(a)$.

Definition (Encoding RSs into PRSs). Let $\mathcal{A} = (S, A)$ be a RS, its encoding into a Positive RS is obtained by considering $\mathcal{A}^+ := (\mathbf{S}, A^+)$ whose set of positive reactions $A^+ := A_{pos}^+ \cup A_{neg}^+$ is defined as follows:

$$A_{pos}^+ := \{(R \cup \bar{I}, P) | (R, I, P) \in A\}$$

$$A_{neg}^+ := \bigcup_{a \in S} \{(\mathbf{T}, \bar{a}) | \mathbf{T} \in \text{Proh}_{\mathcal{A}}(a)\}$$

The resulting \mathcal{A}^+ satisfies the two conditions from Definition and thus is a Positive RS.

The states of the new Positive RS are in bijection with the states of the old system and can be proven that the two systems compute exactly the same states at each step.

2.3.2 Minimization

The procedure of converting a RS into a Positive RS can produce a system with many redundant reactions. The following rules are used to minimize the reactions after they are computed:

1. The reaction $r_1 = (R_1, I_1, P)$ can be omitted if a reaction $r_2 = (R_2, I_2, P)$ such that $R_2 \subseteq R_1$ and $I_2 \subseteq I_1$ is present.
2. If both reactions $r_1 = (R \cup \{a\}, I, P)$ and $r_2 = (R, I \cup \{a\}, P)$ are present, they can be replaced by $r = (R, I, P)$.

In general one can apply a minimization process to both standard and Positive RS and derive a simplified version of the original system with fewer reactions.

2.3.3 Slicing

In the context of programming, dynamic slicing is a technique that helps a user to debug a program by simplifying a partial execution trace, by pruning parts which are irrelevant and highlighting parts of the program which are responsible for the production of an error. In the case of RSs, the goal is to highlight how a subset of the elements in a state were originated. This includes the reactants and reactions that were responsible for producing them.

Starting from the pair $\frac{D_\sigma}{C_m}$ denoting the user's marking and proceeding backwards, apply iteratively a slicing step that deletes from the partial computation all information not related to D_σ . The sliced trace will contain only the subsets of entities and reactions which are necessary for deriving the marked entities.

Since the algorithm 1[8] can only capture dependencies related to reactants, but ignores the ones related to inhibitors, converting the RS into a Positive RS makes possible the tracking of the absence of entities via negative entities. Minimizing the Positive RS reduces the noise in the output and is thus desirable.

Algorithm 1 Trace Slicer

Input:- a reaction system \mathcal{A}

- a trace $T = \frac{D_0}{C_0} \xrightarrow{N_1} \dots \xrightarrow{N_m} \frac{D_m}{C_m}$
- a marking $D_\sigma \subseteq D_m$

Output: a sliced trace $\frac{D'_0}{C'_0} \xrightarrow{N'_1} \dots \xrightarrow{N'_m} \frac{D'_\sigma}{C'_m}$

```

1:  $D'_m := D_\sigma$ 
2: for  $i = \{m, m-1, \dots, 1\}$  do
3:    $D'_{i-1} := \emptyset$ 
4:    $C'_{i-1} := \emptyset$ 
5:    $N'_i := \emptyset$ 
6:   for  $r_j = (R_j, I_j, P_j) \in N_i$  such that  $(D'_i \cap P_j \neq \emptyset)$  do
7:      $N'_i := N'_i \cup \{j\}$ 
8:      $C'_{i-1} := C'_{i-1} \cup (R_j \cap S_C)$ 
9:      $D'_{i-1} := D'_{i-1} \cup (R_j \cap S_D)$ 
10:  end for
11: end for

```

2.4 Bisimulation

Given two distinct RS processes, the natural question to ask would be if their simulation is the same, or at least behaves the same. Bisimulation is one such relation, defined in terms of coinductive games, of fixed point theory and of logic. Bisimulation equivalence aims to identify transitions systems with the same branching structure, and which thus can simulate each other in a stepwise manner.

Definition (Transition System[2]). A transition system TS is a tuple $(S, Act, \rightarrow, I, AP, L)$ where:

- S is a set of states,
- Act is a set of actions,
- $\rightarrow \subseteq S \times Act \times S$ is a transition relation,
- $I \subseteq S$ is a set of initial states,
- AP is a set of atomic propositions,
- $L : S \rightarrow 2^{AP}$ is a labeling function.

TS is called finite if S , Act , and AP are finite.

The intuitive behavior of a transition system can be described as follows: the transition system starts in some initial state $s_0 \in I$ and evolves according to the transition relation \rightarrow . Given s as the current state, then a transition $s \xrightarrow{\alpha} s'$ is selected nondeterministically and taken, meaning the action α is performed and the transition system evolves from state s into the state s' . The labeling function L relates a set $L(s) \in 2^{AP}$ of atomic propositions to any state s . It intuitively stands for exactly those atomic propositions $\alpha \in AP$ which are satisfied by state s .

Definition (Bisimulation Equivalence[2]). Let $TS_i = (S_i, Act_i, \rightarrow_i, I_i, AP, L_i)$, $i \in \{1, 2\}$, be transition systems over AP. A bisimulation for (TS_1, TS_2) is a binary relation $\mathcal{R} \subseteq S_1 \times S_2$ such that:

- $\forall s_1 \in I_1 (\exists s_2 \in I_2. (s_1, s_2) \in \mathcal{R})$ and $\forall s_2 \in I_2 (\exists s_1 \in I_1. (s_1, s_2) \in \mathcal{R})$
- for all $(s_1, s_2) \in \mathcal{R}$ it holds:
 1. $L_1(s_1) = L_2(s_2)$
 2. if $s'_1 \in Post(s_1)$ then there exists $s'_2 \in Post(s_2)$ with $(s'_1, s'_2) \in \mathcal{R}$
 3. if $s'_2 \in Post(s_2)$ then there exists $s'_1 \in Post(s_1)$ with $(s'_1, s'_2) \in \mathcal{R}$.

TS_1 and TS_2 are bisimulation-equivalent (bisimilar), denoted $TS_1 \sim TS_2$, if there exists a bisimulation \mathcal{R} for (TS_1, TS_2) .

Where $Post(s)$ is the set of successors of s defined as

$$Post(s) := \bigcup_{\alpha \in Act} \left\{ s' \in S \mid s \xrightarrow{\alpha} s' \right\}$$

An intuitive way to see bisimulation is by framing it as a game between an attacker and a defender: the attacker wants to disprove the equivalence between two processes s and t , the latter tries to show that s and t are equivalent. Each turn the attacker picks one process and one transition $s \xrightarrow{\alpha} s'$, the defender must reply by picking one transition $t \xrightarrow{\alpha} t'$ of the other process with exactly the same label α . The game ends either with the attacker winning by finding a transition with no equivalent one in the other process or with the attacker losing by having no transitions available.

2.4.1 Algorithms for evaluating bisimulation

Follows a definition of a partition, used extensively in the following algorithms:

Definition (Partition). A partition of S is a set $\{B_0, \dots, B_k\}$, $k \geq 0$ of non-empty subsets of S such that:

- $B_i \cap B_j = \emptyset$, for all $0 \leq i < j \leq k$,
- $S = B_0 \cup B_1 \cup \dots \cup B_k$.

An equivalence relation over the set of states S can be represented as a partition of the states. The sets B_i are called blocks.

Let π and π' be two partitions of S . π' is a refinement of π if for each block $B' \in \pi'$ there exists some block $B \in \pi$ such that $B' \subseteq B$.

The algorithm of Kanellakis and Smolka[1] Given a transition system $T = (S, Act, \rightarrow, I, AP, L)$, let $\pi = \{B_0, \dots, B_k\}$, $k \geq 0$ be a partition of the set of states S . The algorithm due to Kanellakis and Smolka is based on the notion of splitter.

Definition (Splitter). A splitter for a block $B_i \in \pi$ is a block $B_j \in \pi$ such that, for some action $\alpha \in Act$, some states in B_i have α -labeled transitions whose target is a state in B_j and others do not.

Intuitively, thinking of blocks as representing approximations of equivalence classes of processes with respect to strong bisimilarity, the existence of a splitter B_j for a block B_i in the current partition indicates that we have a reason for distinguishing two groups of sets of states in B_i , namely those that afford an α -labeled transition leading to a state in B_j and those that do not. Therefore B_i can be split by B_j with respect to action α into the two new blocks:

$$B_i^1 = \left\{ s \mid s \in B_i \text{ and } s \xrightarrow{\alpha} s' \text{ for some } s' \in B_j \right\} \text{ and} \\ B_i^2 = B_i \setminus B_i^1.$$

This splitting results in the new partition:

$$\pi' = \{B_o, \dots, B_{i-1}, B_i^1, B_i^2, B_{i+1}, \dots, B_k\}$$

which is a refinement of π .

The algorithm of Kanellakis and Smolka iterate the splitting of some blocks B_i by some blocks B_j with respect to some action α until no further refinement of the current partition is possible. The resulting partition coincides with bisimilarity over the input labeled transition systems when the initial partition π_{initial} is chosen to be equal to S and is called the coarsest stable partition.

Definition (Stable Partition and Coarsest Stable Partition). A set $B \subseteq S$ is stable with respect to a set $I \subseteq S$ if either $B \subseteq \text{pre}(I)$ or $B \cap \text{pre}(I) = \emptyset$.

A partition π is stable with respect to a set I if each block $B \in \pi$ is stable with respect to I .

A partition π is stable with respect to a partition π' if π is stable with respect to each block $B' \in \pi'$. A partition π is stable if it is stable with respect to itself.

The coarsest stable refinement of a partition π_{initial} is a stable partition that is refined by any other stable partition that refines π_{initial} .

Note that $B \subseteq S$ is stable with respect to a block $C \subseteq S$ if and only if C is not a splitter for B .

Follows the pseudocode for the algorithm of Kanellakis and Smolka[1].

The algorithm uses the function `split(B, α, π)` which given a partition π , a block $B \in \pi$ and an action α , splits B with respect to each block in π and action α .

Theorem (Kanellakis and Smolka). When applied to a finite labeled transition system with n states and m transitions, the algorithm of Kanellakis and Smolka computes the partition corresponding to bisimilarity in time $O(n \cdot m)$. \square

Proof of correctness relies on the fact that when *changed* is false, there is no splitter for any of the blocks in π . Moreover, if we denote by π_i the partition after the i -iteration of the main loop, we have $\sim \subseteq \sim_i \subseteq \pi_i$. Thus the algorithm terminates with $\pi = \sim$.

```

1:  $\pi := S$ 
2:  $changed := \text{true}$ 
3: while  $changed$  do
4:    $changed := \text{false}$ 
5:   for each block  $B \in \pi$  do
6:     for each action  $\alpha$  do
7:       sort the  $\alpha$ -labeled transitions from states in  $B$ 
8:       if  $\text{split}(B, \alpha, \pi) = \{B_1, B_2\} \neq \{B\}$  then
9:         refine  $\pi$  by replacing  $B$  with  $B_1$  and  $B_2$ 
10:         $changed := \text{true}$ 
11:      end if
12:    end for
13:  end for
14: end while

```

```

1: procedure SPLIT( $B, \alpha, \pi$ )
2:   choose some state  $s \in B$ 
3:    $B_1, B_2 := \emptyset$ 
4:   for each state  $t \in B$  do
5:     if  $s$  and  $t$  can reach the same set of blocks in  $\pi$  via  $\alpha$ -labeled transitions then
6:        $B_1 := B_1 \cup \{t\}$ 
7:     else
8:        $B_2 := B_2 \cup \{t\}$ 
9:     end if
10:   end for
11:   if  $B_2$  is empty then
12:     return  $\{B_1\}$ 
13:   else
14:     return  $\{B_1, B_2\}$ 
15:   end if
16: end procedure

```

The algorithm of Paige and Tarjan[10] Performance of the algorithm by Kanellakis and Smolka can be significantly improved through the use of more complex data structures. Paige and Tarjan proposed an algorithm that utilizes information about previous splits to make future splits more efficient. A simple algorithm over a one symbol alphabet is presented, followed by an algorithm that converts any LTS into a one symbol LTS.

The Paige-Tarjan algorithm is based on the following observation. Let B be stable with respect to S , and let S be partitioned into S_1 and S_2 . Then, if $B \cap S = \emptyset$, B is stable with respect to both S_1 and S_2 . Otherwise B can be split into three blocks:

$$\begin{aligned} B_1 &= B \setminus \text{pre}(S_2), \\ B_{12} &= B \cap \text{pre}(S_1) \cap \text{pre}(S_2), \\ B_2 &= B \setminus \text{pre}(S_1). \end{aligned}$$

The improvement in complexity that the Paige-Tarjan algorithm provides over the algorithm by Kanellakis and Smolka stems from the fact that three-way splitting can be performed in time proportional to the size of the smaller of the two blocks S_1, S_2 .

Two types of splitter can be identified: simple and compound splitters.

simple splitters are used to split blocks in π into two disjoint subsets as done in the algorithm of Kanellakis and Smolka.

Definition (Simple splitting). Let π be a partition and let B be a set of states in S . Define $\text{split}(B, \pi)$ as the following procedure:

For each block $B' \in \pi$ such that B' is not stable with respect to B , replace B' by the blocks

$$\begin{aligned} B'_1 &= B' \cap \text{pre}(B) \quad \text{and} \\ B'_2 &= B' \setminus \text{pre}(B). \end{aligned}$$

B is a splitter for π when $\text{split}(B, \pi) \neq \pi$, in which case π is refined with respect to B and $\text{split}(B, \pi)$ is the partition that results from that refinement.

Some useful properties follow:

Lemma. [1]

1. Stability is preserved by refinement: if π refines π' and π' is stable with respect to a set of states I , then so is π .
2. Stability is preserved by union: if π is stable with respect to sets I_1 and I_2 , then π is also stable with respect to $I_1 \cup I_2$.
3. Assume that $B \subseteq S$. Let π_1 and π_2 be two partitions of S such that π_1 refines π_2 . Then $\text{split}(B, \pi_1)$ refines $\text{split}(B, \pi_2)$.
4. Assume that $B, B' \subseteq S$. Let π be a partition of S . Then

$$\text{split}(B, \text{split}(B', \pi)) = \text{split}(B', \text{split}(B, \pi))$$

□

In order to implement the algorithm efficiently, it is useful to reduce the problem to that of considering a labeled transition system without deadlocked states, meaning without states with no outgoing edge. This can be done easily by preprocessing the initial partition $\pi_{initial}$ by splitting each block $B \in \pi_{initial}$ into:

$$\begin{aligned} B_1 &= B \cap pre(S) \quad \text{and} \\ B_2 &= B \setminus pre(S). \end{aligned}$$

B_2 will never be split again by the refinement algorithm. Therefore run the refinement algorithm starting from the partition $\pi'_{initial} = \{B_1 \mid B \in \pi_{initial}\}$.

In order to find splitters efficiently, some additional information is kept. The algorithm maintains another partition X such that

- π is a refinement of X and
- π is stable with respect to X .

Initially $X = \{S\}$. Follows a general outline of the algorithm.

```

1: while  $\pi \neq X$  do
2:   Find a block  $I \in X \setminus \pi$ 
3:   Find a block  $B \in \pi$  such that  $B \subseteq S$  and  $|B| \leq \frac{|I|}{2}$ 
4:   Replace  $I$  withing  $X$  with the two sets  $B$  and  $I \setminus B$ 
5:   Replace  $\pi$  with split( $I \setminus B$ ,  $\text{split}(B, \pi)$ ).
6: end while

```

The efficiency of the above algorithm relies on the heuristic for the choice of the block B at line 3 and on the use of three-way splitting to implement line 5 efficiently.

Suppose that we have a partition π that is stable with respect to a set of states I that is a union of some of the blocks in π . Assume also that π is refined first with respect to a non-empty set $B \subset I$ and then with respect to $I \setminus B$. Two properties can be observed:

- Refining π with respect to B splits a block $D \in \pi$ into two blocks $D_1 = D \cap pre(B)$ and $D_2 = D \setminus pre(B)$ if, and only if, D is not stable with respect to B .
- Refining further `split(B, π)` with respect to $I \setminus B$ splits the block D_1 into two blocks $D_{11} = D_1 \cap pre(S \setminus B)$ and $D_{12} = D_1 \setminus D_{11}$ if, and only if, D_1 is not stable with respect to $S \setminus B$.

A block I is simple if it is also a block of π and is compound otherwise. Note that a compound block I contains at least two blocks of π . A compound block can be partitioned into B and $I \setminus B$ in such a way that both of the properties 2.4.1 hold. If π is stable with respect to I , either $D \subseteq pre(I)$ or $D \cap pre(I) = \emptyset$. If D is not stable with respect to B , it holds that $D \not\subseteq pre(B)$ and $D \cap pre(B) \neq \emptyset$. Therefore, $D \subseteq pre(I)$. One can thus infer that $D_{12} = D_1 \setminus D_{11} = D_1 \cap (pre(B) \setminus pre(I \setminus B))$, which is the crucial observation underlying the implementation of the algorithm.

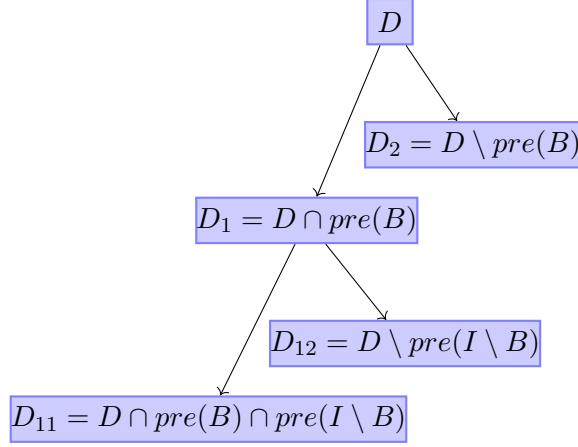


Figure 2.3: Three-way splitting of a block.

Depicted in 2.3 the result of a three-way split of a block D .

The time performance of the algorithm by Paige and Tarjan relies on the following observations:

- Each state in the input labeled transition system is in at most $\log(n+1)$ blocks B used as refining sets, since each is at most half the size of the previous one.
- A refinement step with respect to a block B , as shown by Paige and Tarjan[16], can be implemented in time $O(|B| + \sum_{b \in B} |\text{pre}(b)|)$ by means of use of appropriate data structures.

The algorithm has thus an $O(m \log n)$ time bound.

log-space reduction of bisimilarity checking over a one-letter action set

In order to apply the previous algorithm to a LTS with arbitrary number of labels, there needs to be first a reduction to an equivalent LTS with only one symbol as label. Assume two given processes P and Q over an LTS T with the set of actions $\{\alpha_1, \alpha_2, \dots, \alpha_l\}$. Let T' be the modified LTS which contains all the process of T together with some additional ones defined in the following way: for every transition $P_1 \xrightarrow{\alpha_i} P_2$ in T we add into T'

- two transitions $P_1 \rightarrow P_{(P_1, \alpha_i, P_2)}$ and $P_{(P_1, \alpha_i, P_2)} \rightarrow P_2$ where $P_{(P_1, \alpha_i, P_2)}$ is a newly added state, and
- a newly added path of length i from $P_{(P_1, \alpha_i, P_2)}$.

Finally for every process P in T we create in T' a newly added path of length $l+1$ starting from P . A small optimization can be added by sorting the frequency of labels and thus creating the lowest possible number of auxiliary nodes for each label.

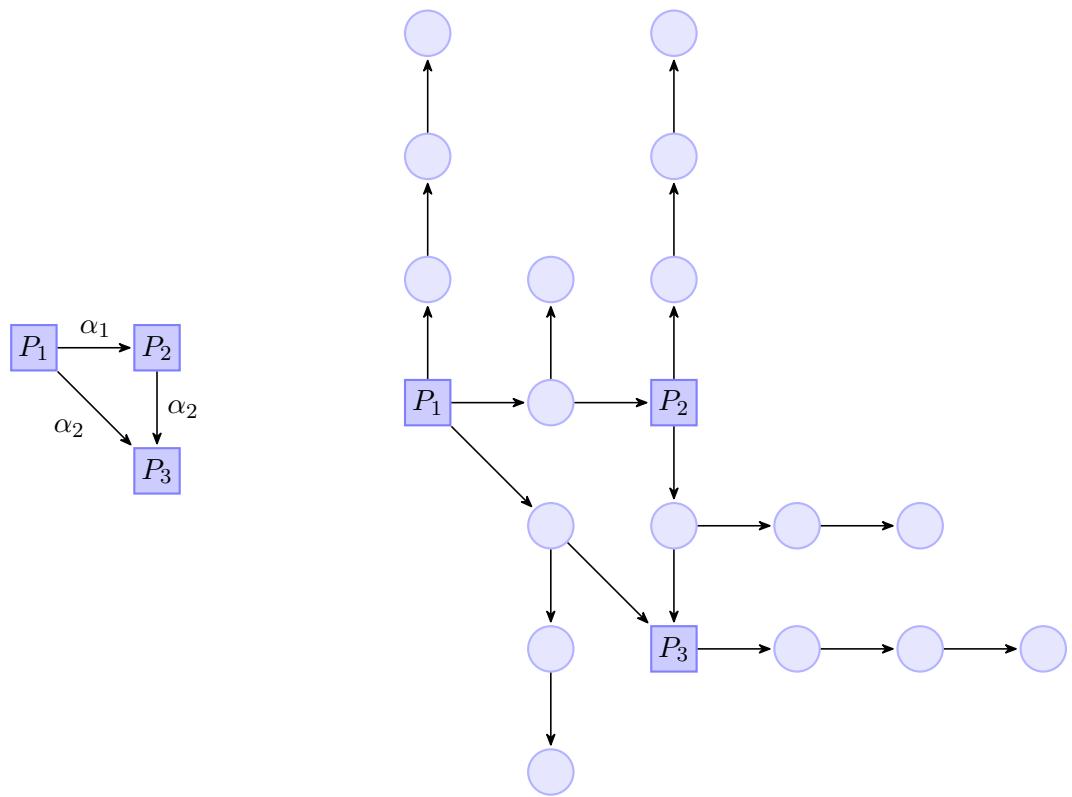


Figure 2.4: Example of reduction

Chapter 3

Design

Two main problems were targeted during the design: simulating the behavior of Reaction Systems, RS processes and other operations on LTS, and interacting with the user in a intuitive manner. The programming language chosen was Rust[27], since it offered good performance and ease of development. Two Git repositories are provided: ReactionSystems[19] and ReactionSystemsGUI[20].

The ReactionSystems project follows a modular architecture and clear design principles to mirror the theoretical model; it implements procedures over RS as pure rust functions and is structured as a library. It also provides a crude Command Line Interface (CLI) for some of the functions provided. The CLI expects a path to a file with a RS definition and some instructions to execute.

The code is organized in workspaces in order to reduce compilation time and aid code reuse. In the second Git repository a native and web application is implemented in Rust and in WebAssembly[18] generated from Rust code. The web application consists of only static files and as such may be served by a simple HTTP server.

In the signature of the functions, types will be displayed in teletype font. The decorator `?` will be used for both option types and for result types without distinction.

3.1 ReactionSystems

The design is structured to faithfully implement the reaction system formalism while remaining flexible. It provides a foundation that matches theoretical definitions (ensuring correctness) and supports further expansion (such as adding optimization, visualization, or integration with other tools) by maintaining a clean separation between the model representation and the execution logic. Since the language Rust supports object-oriented programming via traits, but lacks generic inheritance, the design of the basic building blocks of RSs are designed around this limitation. Usually a basic trait is provided for each of them and an extension of the trait is implemented for all structures that implement the basic trait.

Since it is not practical for a user to specify the structures in Rust, a syntax for the basic structures has been specified. This syntax tries to remain compatible, as much as

possible, with ones from previous software. To develop the parser, LALRPOP[9] was chosen as the parser generator framework. LALRPOP code is transpiled to Rust code via macros and then compiled to machine code.

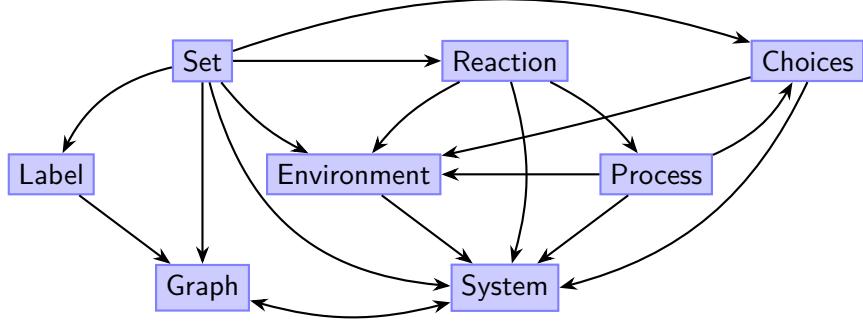


Figure 3.1: Basic structures and relationships between them

3.1.1 Entities and Translator

Entities are the most basic data structure that a RS need to keep track of. They don't have a specified interface and are instead treated only in sets.

Positive elements are also defined and have a state, either **Positive** or **Negative**.

Since internally entities are represented as integers, a structure that keeps track of assignment between strings and integer is provided (**Translator**). This poses a problem with the default methods for formatting available in Rust, since for the trait **Display** and **Debug** only the structure itself can be used to generate the string. The trait **PrintableWithTranslator** and the structure **Formatter** solve this issue by incorporating the **Translator** into the struct. **Display** is then implemented on the generic structure **Translator**.

3.1.2 Set

The common procedures required for all sets are:

- `is_subset(a, b) → bool`, which should return true if $a \subseteq b$;
- `is_disjoint(a, b) → bool`, which should return true if $a \cap b = \emptyset$;
- `union(a, b) → set`, which should return $a \cup b$;
- `push(a, b)`, which should replace a with $a \cup b$ in place;
- `intersection(a, b) → set`, which should return $a \cap b$;
- `subtraction(a, b) → set`, which should return $a \setminus b$;
- `len(a) → int`, which should return the number of elements in a ;
- `is_empty(a) → bool`, which should return true if a has no elements, false otherwise;

- `contains(a, e : entity) → bool`, which should return $e \in a$;
- `add(a, e : element)`, which should add the element e to a in place.

Some other procedures are required for ease of use:

- `extend(a, b?)`, which should extend a with $a \cup b$ if b is a non-null value, and leave a unchanged otherwise, similar to `push` (3.1.2).

Two other procedures are implemented for all structs that implement the `BasicSet` trait:

- `iter(a) → iterable`, which returns an iterator over the elements of the set a ;
- `split(a, trace : [set]) → ([set], [set])?`, which returns the prefix and the loop part of a trace.

Both normal sets and positive sets satisfy this interface, but have additional specific functions for converting between the two.

The syntax for sets is:

<i>Set</i>	\coloneqq	$\{S\}$	<i>PositiveSet</i>	\coloneqq	$\{S\}$
<i>S</i>	\coloneqq	s, S	<i>S</i>	\coloneqq	<i>state s, S</i>
	$ $	s		$ $	<i>state s</i>
	$ $	ϵ		$ $	ϵ
Where			Where		
s		is a string	s		is a string
ϵ		is the empty string	ϵ		is the empty string

Syntax 3.1: Syntax for Set

Syntax 3.2: Syntax for Positive Set

3.1.3 Reaction

The methods required for all reactions are:

- `enabled(r, state : set) → bool`, which returns true if the reaction is enabled given the entities supplied by `state`;
- `compute_step(r, state : set) → set?`, which returns the products of the reaction if it is enabled by `state`.

All reactions that satisfy the basic trait automatically implement the following methods:

- `find_loop(rs : [reaction], entities : set, q : set) → ([set], [set])`, which finds a loop and returns the sets that make up the prefix and the loop separately;
- `find_only_loop(rs : [reaction], entities : set, q : set) → [set]`, which finds a loop and returns the sets that form it;

- `find_prefix_len_loop(rs : [reaction], entities : set, q : set) → (integer, [set])`, which finds a loop and returns the length of the prefix and the sequence of sets that compose the loop;

$$\begin{aligned} \text{Reaction} &::= [s, s, s] \\ &\quad | \quad [r: s, i: s, p: s] \end{aligned}$$

Where

s is a set, see 3.1

Syntax 3.3: Syntax for Reaction

$$\begin{aligned} \text{PositiveReaction} &::= [s, s, s] \\ &\quad | \quad [r: s, i: s, p: s] \end{aligned}$$

Where

s is a positive set, see 3.2

Syntax 3.4: Syntax for Positive Reaction

3.1.4 Process

Process structures mirror the structure of RS processes as described in Section 2.2. Since there is not much behavior that is shared between implementations and since usually they are used with pattern matching, the trait that describe a basic process is very simple.

- `concat(a, b) → process`, which returns a new process $a|b$ flattened with regards to parallel composition;
- `all_elements(a) → set`, which returns all the entities used in the process;
- `filter_delta(a, id : entity) → set?`, which returns the first rule $X = Q.X$ for any symbol X and any set Q .

$ \begin{array}{lcl} Process & ::= & [P] \\ P & ::= & C, P \\ & & C \\ C & ::= & (C) \\ & & \text{nill} \\ & & s.C \\ & & C+C \\ & & ?r?.C \\ & & \langle i, C \rangle .C \\ & & x \end{array} $ <p>Where</p> $ \begin{array}{ll} s & \text{is a set,} \\ & \text{see 3.1} \\ r & \text{is a reaction,} \\ & \text{see 3.3} \\ i & \in \mathbb{N} \\ x & \text{is a variable name} \end{array} $	$ \begin{array}{lcl} PositiveP & ::= & [P] \\ P & ::= & C, P \\ & & C \\ C & ::= & (C) \\ & & \text{nill} \\ & & ps.C \\ & & C+C \\ & & ?pr?.C \\ & & \langle i, C \rangle .C \\ & & x \end{array} $ <p>Where</p> $ \begin{array}{ll} ps & \text{is a positive set,} \\ & \text{see 3.2} \\ pr & \text{is a positive} \\ & \text{reaction, see 3.4} \\ i & \in \mathbb{N} \\ x & \text{is a variable name} \end{array} $
--	--

Syntax 3.5: Syntax for Process

Syntax 3.6: Syntax for Positive Process

3.1.5 Choices

Since one RS process may have more than one possible next system when evaluating, there is a need to express all possible choices for next states. The structure `choices` represents all those possible continuations, associating a set with a process. The set signifies all the entities that are provided by the context by choosing that context. One particular operation called `shuffle` is needed: given two choices structures c_1 and c_2 where $c_i : \text{set} \rightarrow \text{process}, i \in \{1, 2\}$, it generates a new choices structure c' such that $\forall s_1 \in \text{domain}(c_1).(\forall s_2 \in \text{domain}(c_2).\text{domain}(c') \ni (s_1 \cup s_2) \wedge c'(s_1 \cup s_2) = \text{concat}(c_1(s_1), c_2(s_2)))$, where \rightarrow symbolizes a partial function. Intuitively it is all the possible combinations of two parallel processes.

3.1.6 Environment

An environment can be thought as an association between variable names and processes. The basic interface requires the following methods:

- `get(a, k : entity) → process`, which returns the process associated with the variable k ;
- `all_elements(a) → set`, which returns all the entities used in any of the processes;
- `unfold(a, context : process, s : set) → choices?`, which returns the list of choices for the context, given the process definitions environment and is used to generate the next systems with the SOS rules.

These methods are automatically implemented for all `BasicEnvironment`:

- `lollipops_decomposed`,
- `lollipops_decomposed_named`,
- `lollipops_prefix_len_loop_decomposed`,
- `lollipops_prefix_len_loop_decomposed_named`,
- `lollipops_only_loop_decomposed`,
- `lollipops_only_loop_decomposed_named`.

They all try to find a loop and return some information about the found loop. The `_named` variants require a variable symbol for which in the environment there is an association to a process with the form $X = Q.\text{rec}(X) = Q.X$, where Q is a set and X is a variable name. The others instead finds all the symbols that satisfy the constraint and uses them all. Function `lollipops_decomposed` returns the trace of sets for the prefix and the trace of sets for the loop for each recursive variable.

`lollipops_prefix_len_loop_decomposed` returns the length of the prefix and the trace of the loop for each recursive variable. `lollipops_only_loop_decomposed` returns the trace of the loop for each recursive variable.

$$\begin{array}{lcl} \text{Environment} & ::= & [E] \\ E & ::= & x = c, \quad E \\ & | & x = c \\ & | & \epsilon \end{array}$$

Where

$$\begin{array}{ll} c & \text{is a process, see 3.5} \\ x & \text{is a variable name} \\ \epsilon & \text{is the empty string} \end{array}$$

Syntax 3.7: Syntax for Environment

$$\begin{array}{lcl} \text{PositiveEnvironment} & ::= & [E] \\ E & ::= & x = pc, \quad E \\ & | & x = pc \\ & | & \epsilon \end{array}$$

Where

$$\begin{array}{ll} pc & \text{is a positive process, see 3.6} \\ x & \text{is a variable name} \\ \epsilon & \text{is the empty string} \end{array}$$

Syntax 3.8: Syntax for Positive Environment

3.1.7 System

The basic interface for systems is only the following methods:

- `to_transitions_iterator(sys)` → iterator that yields (label, system)
- `to_slicing_iterator(sys)` → iterator that yields
(set, set, [int], system)
- `context_elements(sys)` → set
- `products_elements(sys)` → set

The method `to_transitions_iterator` should return an iterator over all the possible evaluations of the system. Likewise `to_slicing_iterator` should return an iterator over the same outgoing edges, but with information that support the creation of a trace to be used for slicing.

The two methods `context_elements` and `products_elements` should return the set of entities that are related to the context and the one related to the reactions. Since it may be a computationally expensive calculation, the result is cached in the structures. The heuristic that decides which elements belong to the context and which belong to the reactions might give incorrect results, so methods are provided that override the calculated values.

Other methods are implemented for all structures that satisfy the previous interface:

- `unfold(sys)` → `choices`[?], which, by calling the same method of the environment, returns the list of choices for the context;
- `run(sys)` → [system][?], which computes the sequence of systems for the leftmost execution;
- `digraph(sys)` → `graph`[?], which computes the graph generated by the execution of the system;
- `target(sys)` → (integer, set)[?], which returns the state in one of the terminal states and the number of steps to arrive at the last state;
- `slice_trace(sys)` → `trace`[?], which generates, similarly to `run`, a trace appropriate to run slicing calculations over;
- `lollipops(sys)` → [[set], [set]], similar to the method `lollipops_decomposed` provided by `environment`.
- `lollipops_only_loop_named(sys, el : element)` → [[set]][?], similar to the method provided by `environment`, returns the sequence of entities in the loop individuated by the variable name `el` in the environment.

$$\begin{aligned}
 System & ::= \text{Environment:}e \\
 & \quad \text{Initial Entities:}s \\
 & \quad \text{Context:}c \\
 & \quad \text{Reactions: } (R) \\
 R & ::= r;R \\
 & \quad | \\
 & \quad \epsilon
 \end{aligned}$$

Where

e	is an environment, see 3.7
s	is a set, see 3.1
c	is a Process, see 3.5
r	is a reaction, see 3.3
ϵ	is the empty string

Syntax 3.9: Syntax for System

While writing systems, occurs often to forget an element or swap elements names. This type of user error is particularly difficult to spot since most element names are not easily recognizable and the system grammar is particularly dense in information. A useful method to mitigate these problems is provided: **statistics**. For the structure **System** a static analysis of the entities of the system is run and various parameters are checked and reported.

3.1.8 Label

The label structure holds the information about how entities are used in the production of a system and are the labels on the edges of the graphs. Since the only use is to hold data, no meaningful method is required.

$$\begin{aligned}
 Label & ::= [\text{Entities:}s, \\
 & \quad \text{Context:}s, \\
 & \quad \text{Reactants:}s, \\
 & \quad \text{ReactantsAbsent:}s, \\
 & \quad \text{Inhibitors:}s, \\
 & \quad \text{InhibitorsPresent:}s, \\
 & \quad \text{Products:}s]
 \end{aligned}$$

Where

s	is a set, see 3.1
-----	-------------------

Syntax 3.10: Syntax for Label

```

PositiveLabel ::= [ Entities:ps,
                    Context:ps,
                    Reactants:ps,
                    ReactantsAbsent:ps,
                    Inhibitors:ps,
                    InhibitorsPresent:ps,
                    Products:ps ]

```

Where
 ps is a positive set, see 3.2

Syntax 3.11: Syntax for Positive Label

3.1.9 Graph

The project uses petgraph[5] as graph data structure library. *petgraph* provides several graph types, but the only one used is **Graph**, since it provided the best performance during testing. The library has methods for converting the graph structures into Dot Language[12] and GraphML File Format[24]. The Dot methods where found to be not powerful enough and where partially rewritten in the file `dot.rs`.

Custom formatting of the graphs was a key requirement, so domain specific languages are provided to customize the appearance of the generated formats. Four structures are provided:

- `NodeDisplay`, to specify the text displayed on each node;
- `EdgeDisplay`, to specify the text displayed on each edge;
- `NodeColor`, to specify the color of each node;
- `EdgeColor`, to specify the color of each edge.

Follows the BNF for each of the languages:

```

NodeDisplay ::= E
                  | E "s" NodeDisplay
E          ::= Hide
                  | Entities
                  | MaskEntities S
                  | ExcludeEntities S
                  | Context
                  | UncommonEntities
                  | MaskUncommonEntities S

```

Where
 S is a sets of entities
 s is a possibly empty string

Syntax 3.12: Syntax for NodeDisplay

`Hide` ignores the content of the node and prints the empty string, `Entities` prints the list of entities currently available in the system, `MaskEntities` S prints the list of entities masked by a specified set S , `ExcludeEntities` S prints the list of entities except for the entities specified by the set S , `Context` prints the context of the system, `UncommonEntities` prints only the entities that are not shared between all the nodes in the graph, `MaskUncommonEntities` S prints the entities not shared between all the nodes in the graph and masked by a specified set S .

```

EdgeDisplay ::= E
    | E "s" EdgeDisplay
E ::= Hide
    | Products
    | MaskProducts S
    | UncommonProducts
    | UncommonMaskProducts S
    | Entities
    | MaskEntities S
    | UncommonEntities
    | UncommonMaskEntities S
    | Context
    | MaskContext S
    | UncommonContext
    | UncommonMaskContext S
    | Union
    | MaskUnion S
    | UncommonUnion
    | UncommonMaskUnion S
    | Difference
    | MaskDifference S
    | UncommonDifference
    | UncommonMaskDifference S
    | EntitiesDeleted
    | MaskEntitiesDeleted S
    | UncommonEntitiesDeleted
    | UncommonMaskEntitiesDeleted S
    | EntitiesAdded
    | MaskEntitiesAdded S
    | UncommonEntitiesAdded
    | UncommonMaskEntitiesAdded S

```

Where

S is a sets of entities

s is a possibly empty string

Syntax 3.13: Syntax for `EdgeDisplay`

Four version of each base option is available: `normal`, `Mask` which masks the normal set of entities with a specified set, `Uncommon` which considers only the entities that are not shared between all edges of the graph, and `UncommonMask` which combines the two functionalities. The base options return the corresponding entities available in the label.

```

NodeColor ::= !" C "
             | E || NodeColor
E      ::= Entities op S ?"C"
             | Context.Nill ?"C"
             | Context.RecursiveIdentifier ( x )?"C"
             | Context.EntitySet op S ?"C"
             | Context.NonDeterministicChoice ?"C"
             | Context.Summation ?"C"
             | Context.WaitEntity ?"C"
op     ::= ==                                or =
             | <                                or ⊂
             | <=                               or ⊆
             | >                                or ⊃
             | >=                               or ⊇

```

Where

<code>C</code>	is a string that specifies the color of the node
<code>S</code>	is a sets of entities
<code>x</code>	is a variable

Syntax 3.14: Syntax for `NodeColor`

The `NodeColor` structure assigns the first correct color to the node. The structure can be thought of as a list of pairs; each pair has an entry that evaluated returns true or false, and an entry that holds the desired color of the node. To find the correct color, the list is scanned until the first pair that returns true and the color is assigned. If no pair returns true, a default value is assigned, specified after `!`. The possible functions expressible by the grammar are the ones expressed by `E` and query either the entities available or the current context.

```

EdgeColor ::= !"C"
| E || EdgeColor
E ::= Entities op S ?"C"
| Context op S ?"C"
| T op S ?"C"
| Reactants op S ?"C"
| AbsentReactants op S ?"C"
| Inhibitors op S ?"C"
| PresentInhibitors op S ?"C"
| Products op S ?"C"
op ::= == or =
| < or ⊂
| <= or ⊆
| > or ⊃
| >= or ⊇

```

Where

<i>C</i>	is a string that specifies the color of the node
<i>S</i>	is a sets of entities
<i>x</i>	is a variable

Syntax 3.15: Syntax for EdgeColor

EdgeColor behaves in a similar manner as NodeColor, except the base structure is a `Label`, so every field is a `Set`.

3.1.10 Slicing Trace

Only one structure for slicing trace is provided, but is made to work with both RS and Positive RS with generics. The only method they have is `slice(trace, marking : set) → trace?` which returns, if successful, a new sliced trace.

3.1.11 Bisimilarity and Bisimulation

In the workspace `bisimilarity` the algorithms by Kanellakis and Smolka, and Paige and Tarjan are implemented over generic graphs. Instead of an implementation over graphs with generic parameters, the input have to implement generic traits from the `petgraph` library, making it possible to use with different types of graph, for example spare graphs or matrix graphs.

One key feature was the ability to control via a domain specific language the labels on the edges of the graphs. The developed language is able to also specify values over nodes such that nodes with equal value may be collapsed into one node with outgoing and incoming edges inherited from the original nodes. The code for the typechecking and execution is available in the library `assert`.

The language has way to define subroutines or functions, has no while loop and limited for loop construction, so that the execution always terminates.

<i>Assert</i>	$::= \text{label}\{Tree\}$	<i>label</i> is replaced by other strings to differentiate languages
<i>Tree</i>	$::= Tree;Tree$ $\text{if } E \text{ then } \{Tree\};$ $\text{if } E \text{ then } \{Tree\}$ $\text{else } \{Tree\};$ $\text{let } x = E;$ $\text{let } x.Qualifier = E;$ $\text{return } E;$ $\text{for } x \text{ in } Range \{Tree\};$	
<i>E</i>	$::= \text{unaryP}(E)$ $E.\text{unaryS}$ $(E \text{ binary } E)$ $\text{binaryP}(E, E)$ <i>Term</i>	
<i>Term</i>	$::= \text{true}$ false x i l <i>set</i> $'s'$ element $"s"$ string (E)	
<i>Range</i>	$::= \{E\}$ $\{E..E\}$	iterate over set iterate over integer range

Where

<i>S</i>	is a sets of entities, see 3.1
<i>i</i>	$\in \mathbb{Z}$
<i>x</i>	is a variable name
<i>l</i>	is a label, see 3.1
<i>set</i>	is a set, see 3.1
<i>s</i>	is a string

Syntax 3.16: Syntax for **Assert**

Continues on the next page.

```

unaryP ::= empty
| length
| tostr
| toel string to element
| Entities
| Context
| Reactants
| ReactantsAbsent
| Inhibitors
| InhibitorsPresent
| Products
| AvailableEntities
| AllReactants
| AllInhibitors
| SystemEntities
| SystemContext
| source source of edge
| target target of edge
| label label of an edge
| neighbours node's neighbours
| system node's system
binary ::= && logical and, set intersection
| ||
| ~~ logical xor, set xor
| < less, set inclusion
| <= less equal, set inclusion or equal
| > greater, reverse set inclusion
| >= greater equal,
| reverse set inclusion or equal
| ==
| !=
| +
| *
| ^
| /
| % quotient
| :
| :: remainder
| concatenation
binaryP ::= substr
| min
| max
| commonsubstr

```

Syntax 3.16: Syntax for **Assert** (Continued)

```

unaryS  ::=  Entities
          |  length
          |  tostr
          |  toel
```

Syntax 3.16: Syntax for `Assert` (Continued)

Four different languages are provided as modules in the workspace `Assert`:

- `relabel`
- `grouping`
- `positive_relabel`
- `positive_grouping`

`relabel` has, in its syntax, the first token equal to `label`; it also has as special variables `label` and `edge` that function as input, with `edge.label == label`. It's used for grouping labels before calculating bisimilarity. `grouping` has, in its syntax, the first token equal to `node` and has as special variables `entities` and `node` that function as input, with `node.system.SystemEntities == entities`. It's used for grouping nodes. `positive_relabel` and `positive_grouping` have as first token `label` and `node` respectively and behave in a similar manner compared to their respective functions over RS.

The template language requires two structures to function relating to the input of the language: a type structure and a value structure. The trait `SpecialVariables` holds all the necessary functions that need to be implemented for the special variables to function. Finally the generic language can have the two functions `typecheck` and `execute` implemented.

The language is very limited and is only designed for simple algorithms since there is no scoping. Typechecking consists in only asserting acceptable types for unary and binary functions, correct type in range declarations and for all return statements to have the same type.

A version for Positive RS is also provided and reflects the previous grammar with basic types replaced with their positive versions.

3.1.12 Grammar and Separated Grammar

Two workspaces are provided for parsing the structures above. `Grammar` creates only one endpoint that parses a system and a list of instructions. Those instructions are then executed via the library `execution`. A simple CLI has been implemented in the workspace `analysis`, with proper error formatting for LALRPOP errors.

3.1.13 Experiments and Frequency

An experiment is a list of weights and a list of sets of same length. The sets are used as entities given in addition to the context entities when computing the RS. The resulting

trace is then synthesized into relative frequencies. The methods offered by `Frequency` and `PositiveFrequency` are:

- `naive_frequency(sys : system) → frequency?`, which computes the relative frequency of each entity in all traversed states, assuming the computation is finite;
- `loop_frequency(sys : system, symbol : IdType) → frequency`, which computes the relative frequency of each entity in each state of the encountered loop, assuming the system stabilizes in a loop;
- `limit_frequency(experiment : [set], reactions : [reaction], entities : set) → frequency?`, which computes the relative frequency of each entity in the states of the last loop by providing repeatedly the sets in the experiment until the system stabilizes in a loop;
- `fast_frequency(experiment : [set], reactions : [reaction], entities : set, weights : [int]) → frequency?`, which computes the weighted relative frequency of each entity in any of the loops.

$$\begin{array}{lcl}
 \text{Experiment} & ::= & \text{Weights} : W \text{ Sets} : S \\
 W & ::= & i, \ W \\
 & | & i \\
 & | & \epsilon \\
 S & ::= & s, \ S \\
 & | & s \\
 & | & \epsilon
 \end{array}$$

Where

$$\begin{array}{ll}
 s & \text{is a sets of entities, see 3.1} \\
 i & \in \mathbb{Z} \\
 \epsilon & \text{is the empty string}
 \end{array}$$

Syntax 3.17: Syntax for Experiment

3.1.14 Instructions

The command line interface provided by the workspace `analysis` expects as input a path to a file with a RS and some instructions, reads the file, executes the instructions and returns the result in the forms specified. The syntax for specifying instructions is as follows:

$$\begin{array}{lcl}
 \text{Instructions} & ::= & \text{System } Is \\
 & | & \text{Deserialize } (path) \ Is \\
 Is & ::= & I, Is \\
 & | & I \\
 & | & \epsilon
 \end{array}$$

Syntax 3.18: Syntax for Instructions

The instruction `Stats` returns static information and statistics about the system. `Target` and `Run` return the results of the function with the same name defined in section 3.1.7. `Loop` returns the result of the function `lollipops_only_loop_named` described in the same section. `Frequency` returns the results of `naive_frequency` as described in section 3.1.13. `LimitFrequency` and `FastFrequency` require a path to an experiment, with syntax described in 3.1.13, and they return the result of `limit_frequency` and `fast_frequency` respectively. `Digraph` may take a function *group* that groups together nodes that have the same output, described in section 3.1.11. The graphs can be either saved as a Dot or GraphML file or serialized directly. `Bisimilarity` takes a path to another instruction file, whose instructions are ignored, and returns if the two systems are bisimilar. The results can be saved to a file or print to screen or both.

3.2 ReactionSystemsGUI

During development of ReactionSystems, a need for a more intuitive interaction with the structures presented itself. Since all the operations on the types were already limited and structured, a visual programming language was chosen as the best fit.

The library `egui_node_graph2`[28] was chosen since it offered customizability, performance and ease of programming. The library unfortunately lacked compatibility with the most recent version of `egui`[13], so it is included as a workspace and modified to fit better the need of the project. This way a couple of visual bugs present in the original code have also been fixed.

`egui_node_graph2` is based on the library `egui`, which is an immediate mode GUI. It differentiates itself from retained mode GUIs by having all the elements specified at every frame; this eases programming at the expense of performance. The trade-off is favorable since most of the computation will be on the algorithms over RS and the number of elements of the UI will remain small in most cases.

All the functions previously described are available as “nodes” in the GUI program. Each takes one or more inputs, colored by type, and prevents wrong types from connecting, reducing user error when connecting similarly colored types.

Since at every step all of the GUI is recalculated, a robust cache structure is needed. The cache developed keeps track of the modified nodes and only recomputes if necessary, exploiting the structure of the graph.

The library `egui_node_graph2` was also chosen for its ability to create a web application directly from Rust code. The web application is limited; there is no interaction with the file system and no true multi-threading. These limitations are imposed by WebAssembly itself, not by the transpilation from Rust.

The native application executes the expressed instructions in a separate thread and returns the result to the GUI thread to be displayed. Thus the web application may “freeze” and become unresponsive with long calculations.

Both native and web applications have the ability to save the current state and resume. The saved state is stored in the browser cache in the web application and in special directories in the native one:

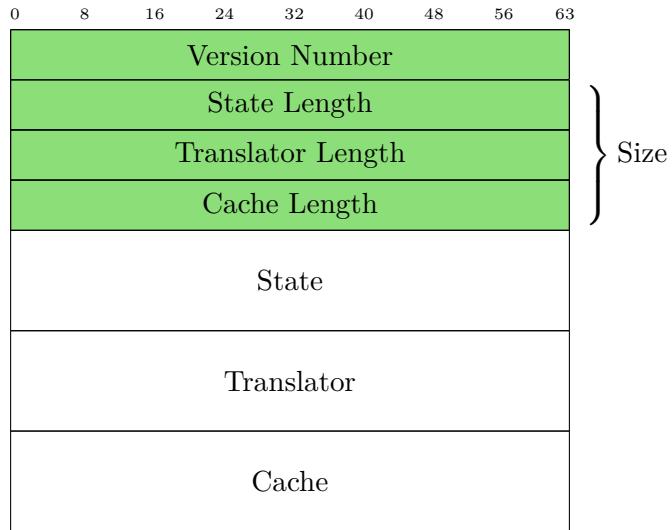


Figure 3.2: Save file structure

- Linux: `/home/UserName/.local/share/Reaction-Systems`
- macOS: `/Users/UserName/Library/Application Support/Reaction-Systems`
- Windows: `C:\Users\UserName\AppData\Roaming\Reaction-Systems\data`

The native application also has the ability to save and load the state from a file. The files have by default the extension “.ron”. The web version has no ability to interact with the file system due to a limitation of WebAssembly.

The file structure can be seen in figure 3.2, where “state” refers to the state of the GUI, “translator” refers to the `Translator` structure used to encode entities names into fixed sized integers, and “cache” refers to the cache structure for the GUI. Version number is a little-endian `u64` that encodes the version number of the application; if different from the version of the application, a warning will be issued, but the application will try and load the state anyway. Each “length” field is a little-endian `u64` and indicates the length in bytes of the corresponding field.

The user can request the result of a computation by interacting with the button “*Set active*” under most of the windows. A panel on the right of the screen appears with the computed result. The nodes “Save string to file”, “Save SVG” and “Save Rasterized SVG” instead have a button “*Write*” that writes to file the result. The node “Read a file” has an extra button “*Update file*” that reads again the file from disk since a filewatcher has not been implemented.

Since the generated graphs were often times immediately converted to DOT files and rendered to SVG, a native renderer is included that can create PNG images of the supplied graph. They are then rendered to screen. This reduces greatly the time switching between software to achieve the same result.

 Error	Structure that holds error messages
 String	A string
 Path	A path to a file
 Svg	A structure for creating and rendering SVG
 PositiveInt	Integer in \mathbb{N}
 Symbol	A single symbol, see section 3.1.1
 System	see section 3.1.7
 Environment	see section 3.1.6
 Set	see section 3.1.2
 Context	see section 3.1.4
 Reactions	see section 3.1.3
 Experiment	see section 3.1.13
 PositiveSystem	see section 3.1.7
 PositiveEnvironment	see section 3.1.6
 PositiveSet	see section 3.1.2
 PositiveContext	see section 3.1.4
 PositiveReactions	see section 3.1.3
 Trace	see section 3.1.10
 PositiveTrace	see section 3.1.10
 Graph	see section 3.1.9
 PositiveGraph	see section 3.1.9
 DisplayNode	see section 3.1.9
 DisplayEdge	see section 3.1.9
 ColorNode	see section 3.1.9
 ColorEdge	see section 3.1.9
 AssertFunction	see section 3.1.11
 GroupFunction	see section 3.1.11
 PositiveAssertFunction	see section 3.1.11
 PositiveGroupFunction	see section 3.1.11

Table 3.19: Types in ReactionSystemsGUI with associated color.

String			Creates a string
Path			Creates a path from a string
Read file			Reads a file into a string
Save string to file	 		Save a string to a file
Symbol			Creates a symbol from a string
Sleep			Waits for selected number of seconds
Dot file to SVG			Parses a Dot file string into an SVG
Save SVG	 		Saves an SVG
Save Rasterized SVG	 		Saves an SVG as a picture
Create System			Creates system from string
Create Positive System			Creates positive system from system
Compose System	  		Composes system from individual structures
Compose Positive System	  		
Decompose System		  	Decomposes system into individual structures
Decompose Positive System		  	
Environment			
Positive Environment			
Set			
Positive Set			
Context			
Positive Context			
Reactions			
Positive Reactions			
Convert to Positive Environment			
Convert to Positive Set			
Convert to Positive Context			
Convert to Positive Reactions			
Statistics			

Table 3.20: Available nodes in ReactionSystemsGUI with inputs and outputs in color.

Continued on next page

(Continued)

Target			
Target of Positive RS			
Run			
Run of Positive RS			
Loop			Applies function <code>lollipops_only_loop_named</code> to RS
Loop of Positive RS			
Create Experiment			Creates experiment from string
Frequency			
Frequency of Positive RS			
Limit Frequency			
Limit Frequency of Positive RS			
Fast Frequency			
Fast Frequency of Positive RS			
Graph System			Creates digraph of RS
Graph Positive System			
Create Dot file			Creates Dot file from a graph
Create Dot file of Positive System			
Create GraphML file			Creates GraphML file from a graph
Create GraphML file of Positive System			
Display node function			
Display edge function			
Color node function			
Color edge function			
Bisimilarity Kanellakis & Smolka			
Bisimilarity Kanellakis & Smolka for Positive RS			
Bisimilarity Paige & Tarjan			

Table 3.20: Available nodes in ReactionSystemsGUI with inputs and outputs in color.

Continued on next page

(Continued)

Bisimilarity Paige & Torjan for Positive RS		
Bisimilarity Paige & Tarjan (ignore labels)		
Bisimilarity Paige & Torjan (ignore labels) for Positive RS		
Create relabeling edge function		
Create relabeling edge function for Positive RS		
Trace		
Positive Trace		
Slice Trace		
Positive Slice Trace		
Trace to string		
Positive trace to string		
Overwrite context entities		
Overwrite reaction entities		
Overwrite context entities of Positive System		
Overwrite reaction entities of Positive System		
Create Grouping Function		
Group Nodes		
Create Grouping Function for Positive System		
Group Nodes of Positive System		

Table 3.20: Available nodes in ReactionSystemsGUI with inputs and outputs in color.

All the types can be seen in table 3.19 that are used for the node's input and output. Each type has a distinct color associated that is used to color the connectors between nodes. All nodes can be seen in table 3.20. The second column holds the color of the inputs types used; the third column holds the color of the outputs types.

Chapter 4

Development

4.1 ReactionSystems

4.1.1 Entities and Translator

Entities are declared in the file `element.rs` and the `Translator` struct is implemented in the file `translator.rs`.

Entities have type `IdType` and are represented as `u32`. Representing arbitrarily named entities with integers has the immediate benefit of faster code execution, but need additional support for the encoding and decoding. Also it does not permit easy merging of different systems. This is because two elements with the same string might be assigned to a different integer and would need to be re-encoded. The `ReactionSystemsGUI` solves this problem by having only one `Translator` class for all entities and systems.

Positive RS have the property that if all the entities are declared in the initial state, in all subsequent states the entities will all be defined either positive or negative. This property can be exploited in the representation of a Positive RS, however the implementation disregards this fact and simply assigns either positive or negative to each positive entity.

The struct `Translator` is formed by two maps, one from strings to `IdType` and the inverse, and by a counter for the last used id. It is essential for this class to be serializable, so that the state of `ReactionSystemsGUI` might save it when necessary. The struct is also used to form the structure `Formatter`, which is used to format all structures that implement `PrintableWithTranslator`.

For example the implementation of `PrintableWithTranslator` for `Set` is the following:

```
1 impl PrintableWithTranslator for Set {
2     fn print(
3         &self,
4         f: &mut fmt::Formatter,
5         translator: &Translator,
6     ) -> fmt::Result {
```

```

7   write!(f, "{")?;
8   let mut it = self.iter().peekable();
9   while let Some(el) = it.next() {
10     if it.peek().is_none() {
11       write!(f, "{}", Formatter::from(translator, el))?;
12     } else {
13       write!(f, "{}, ", Formatter::from(translator, el))?;
14     }
15   }
16   write!(f, "}")
17 }
18 }

```

The structure `Translator` is only borrowed because it is never modified when printing, so only one is needed for all of the print. On lines 11 and 13 instead of directly printing `el`, we first construct another `Formatter` struct and require only for that struct to implement `std::fmt::Display`. This gives modularity and flexibility to the display system.

4.1.2 Set

The structure `set`, implemented in the file `set.rs`, is a key component for all functions in the library. It is realized as a binary tree set[26]. Binary trees were chosen instead of hash sets for various reasons: binary trees support hashing of the whole tree, hash sets do not; the penalty for retrieval of individual elements is offset by the performance gain for set operations like union or intersection.

4.1.3 Reaction

A reaction is a collection of sets, reactants, inhibitors and products for RS and just reactants and products for Positive RS. Since converting between reactions and positive reactions is meaningless for single reactions, we provide a method called

```
into_positive_reactions(reactions : [reactions]) → [positive reactions]
```

that takes a vector of reactions and calculates the prohibiting set and minimizes. The code is available in the file `reactions.rs`.

4.1.4 Process, Choices and Environment

Context processes, available in `process.rs`, have been implemented as trees. Each pointer to the next process is an `Arc`[25] so that they may be used in concurrent applications, like ReactionSystemsGUI. There is no need for interior mutability, so no mutex or semaphore is used. The name of variables used to identify environment processes are converted like entities from strings to integers and they are handled by `Translator`, since there no reason was found to distinguish them.

The structure `Choices` is available in `choices.rs`; `Environment` is available in file `environment.rs`.

`Environment` has been implemented as a binary tree like sets, in order to be able to hash them; even tho no set operations are needed, the performance penalty is small enough.

4.1.5 System

Systems are implemented in the file `system.rs`. Systems are composed by an environment, a set of initial entities, a process and a vector of reaction rules. Two other private fields are used: `context_elements` and `products_elements`. They hold the set of entities that concern context and the ones that concern the products, such that their union is equal to all the entities available to the system and their intersection is the empty set. These two fields are not public since their computation may be particularly expensive, but is not needed for most of the calculations. So it would be wasteful to compute when creating the system and would be unwieldy to cache the result in every function that uses the results. The choice was to make `System` as a structure with interior mutability. This property is checked by the Rust compiler and forbids one from using the structure in hash maps or binary trees. But since we know that these two fields are completely determined by the other four, we ignore them when calculating the hash and assure the compiler of their stability in the file `clippy.toml`, where it is specified that both `System` and `PositiveSystem` are to be ignored.

Since the automatic assignment to context or product element can be erroneous, nodes to overwrite these values are available in `ReactionSystemsGUI`.

The two key functions `to_transition_iterator` and `to_slicing_iterator` specify that they return an iterator, a lazy structure with a `next` method for obtaining the following value. This is to allow for a more efficient implementation in cases where not all states are needed.

4.1.6 Label

Labels have been implemented in the file `label.rs`. Since their primary function is to hold redundant but useful data for other computations, they do not need any algorithms to be implemented directly in their interface.

The structure for a label is:

```
pub struct Label {
    pub available_entities: Set,
    pub context: Set,
    pub t: Set,
    pub reactants: Set,
    pub reactants_absent: Set,
    pub inhibitors: Set,
    pub inhibitors_present: Set,
```

```
    pub products: Set,
}
```

where t is defined as $t := \text{available_entities} \cup \text{context}$. Since t can be uniquely derived from other fields it is ignored when calculating equality or the hash of the label. Positive labels have a similar structure, with `PositiveSet` instead of `Set` in all of the fields.

4.1.7 Graph

Graphs for RS and Positive RS are declared as

```
pub type SystemGraph = Graph<System, Label, Directed, u32>;
```

and

```
pub type PositiveSystemGraph =
    Graph<PositiveSystem, PositiveLabel, Directed, u32>;
```

in the file `graph.rs`, where `Graph` is from the library `petgraph`[5]. This was done to leverage the traits provided already by the external library. `Graph<N, E, Ty, Ix>` takes four generic parameters:

- Associated data N for nodes and E for edges, called weights. The associated data can be of arbitrary type;
- Edge type Ty that determines whether the graph edges are directed or undirected;
- Index type Ix , which determines the maximum size of the graph.

The index type was chosen to be `u32` to balance performance with maximum size of the graph.

The library already provides methods to export the graphs in Dot and GraphML formats, but the Dot export did not meet all the requirements and has been partially rewritten in `dot.rs`. The biggest difference is in the function `graph_fmt`, which has been simplified and made more ergonomic for specifying color of text and background.

As described in subsection 3.1.9, four structures for specifying the display properties of the Dot and GraphML format have been designed. The implementation closely follows the design description, but results in a lot of boilerplate code that can be seen in the file `format_helpers.rs`, helped slightly by custom macros.

The four structures — `NodeDisplay`, `EdgeDisplay`, `NodeColor`, and `EdgeColor` — all have the `generate` and `generate_positive` methods, which convert the relative structure into an executable function that can be used when creating Dot or GraphML documents. No unified trait has been defined since the functions returned have different types and the use for this trait may be limited.

4.1.8 Slicing Trace

Since traces are only lists of states, often no type associated with them is provided; some trace types are present in `trace.rs`. Of particular interest is the structure `SlicingTrace<S, R, Sys>`. Instead of using traits, it was more convenient to use generic type parameters for the slices structures. For both RS and Positive RS the method `slice` faithfully implements the algorithm described in section 2.3.3. A new slice structure is returned because often the previous slice might get reused as input to other slicings. This occurs in a minor performance penalty if only one slice is requested.

4.1.9 Bisimilarity and Bisimulation

The algorithms described in section 2.4 are implemented in the files in the folder `bisimilarity/src`. They are implemented for arbitrary graphs that satisfy some traits defined in the library `petgraph`. For example from file `bisimilarity_kanellakis_smolka.rs`:

```
pub fn bisimilarity<'a, G>(graph_a: &'a G, graph_b: &'a G) -> bool
where
    G: IntoNodeReferences + IntoEdges,
    G::NodeId: std::cmp::Eq + std::hash::Hash,
    G::EdgeWeight: std::cmp::Eq + std::hash::Hash + Clone,
```

The generic parameter `G` has to satisfy `IntoNodeReferences + IntoEdges` but is not constrained to be a `Graph` and could be for example a `StableGraph` or a `GraphMap`. In this way code portability is maximized.

4.1.10 Assert

As described in 3.1.11, a custom language has been developed for the purpose of modifying the graphs. The code is available in the folder `assert/src`. The implemented language can be seen as just one function that will be executed on each node or edge of the graph. The return value of the function will be used to group or relabel the input values. For this purpose a structured statically-typed interpreted language with global variables declarations. The choices make evaluating the language very lightweight and since the programs are usually very short, its not detrimental to the user experience.

Typechecking is done only over operator arguments, range definition and return statements. For example the program

```
node {
    if node.system.SystemEntities > {p0} then {
        return 0;
    }
    if node.system.SystemEntities > {p1} then {
        return false;
    }
}
```

```

if node.system.SystemEntities > {p2} then {
    return 2;
}
return 3;
}

```

would return an error since return statements don't agree on type returned. But the program

```

node {
    if node.system.SystemEntities > {p0} then {
        return 0;
    }
    if node.system.SystemEntities > {p1} then {
        return 1;
    }
    if node.system.SystemEntities > {p2} then {
        return 2;
    }
}

```

would satisfy the typechecker, even tho not all applications will return a value. This error will only be caught executing.

4.1.11 Grammar

The code for the unified grammar is available in the folder `grammar/src` and the code for the separated grammar is available in the folder `grammar_separated/src`. The parser generator code has been placed in separate workspaces so that compilation time may be reduced. The Rust compiler sequentially compiles each file in the same workspace, and if one file is modified, all other files must be re-linked. By separating into different workspaces the computation is parallelized and modifying a file results only in the workspace being recompiled. The workspaces for the grammar are particularly slow to compile and required this treatment.

LALRPOP library allows for user specific errors to be declared. Only two have been employed, `NumberTooBigUsize` and `NumberTooBigi64`, since the default error messaging was adequate. Custom error display has been implemented in file `helper.rs` which creates error messages with color and that highlight the erroneous part of the input. For example the specification of the example in subsection 2.1.1 is the following:

```

Environment: []
Initial Entities: {p1,p3}
Context: [{}.{inc}.{inc}.{dec}.{dec,inc}.nil]
Reactions: (
    ...
)

```

If we omit the last dot in the context:

```
Context: [{}.{inc}.{inc}.{dec}.{dec,inc}nill]
```

we obtain the following error message:

```
Unrecognized token "nill" between positions 82 and 86.  
Expected: (".")  
Line 3 position 40 to 44:  
3 |Context: [{}.{inc}.{inc}.{dec}.{dec,inc}nill]  
|  
|  
|
```

During parsing the symbols are immediately encoded with `Translator` when appropriate. This was made possible by passing state parameter, `Translator`, to the parser.

The code for the unified grammar is mirrored in the separate grammar, where every parser is available to use. This was done because the code generated by the parser generator could grow very large for very nested public functions. By reducing the number of public functions, this problem is mitigated.

4.2 ReactionSystemsGUI

To build an application with egui, a struct that implements `eframe::App` is needed. The methods implemented this way will be called by the internal engine when the GUI will need to be repainted. Since the update function will be called numerous times every second, it is important that expensive calculations be cached for subsequent frames. Since all the heavy computations concern only the RS, a cache is developed that assigns to each node's output the cached value. The cache also has hashes for the previous inputs that speed up comparisons when deciding if a value is still valid or should be replaced.

```
struct CacheInternals {  
    values: HashMap<OutputId, BasicValue>,  
    hash_values: HashMap<OutputId, u64>,  
    hash_inputs: HashMap<OutputId, (u64, Vec<u64>)>,  
    last_output: Option<BasicValue>,  
}
```

where `OutputId` is the type of the output of the nodes, `BasicValue` is the type of the possible values computed in the nodes and hashes are stored as `u64`. `hash_inputs` contains both a hash and a list of hashes. The first refers to the xor of the latter, and is used to quickly check if all the inputs are unchanged. Any interaction with the node structure or with the text fields invalidates the appropriate entries in the cache.

Every time an update to the values of the nodes is requested, a new thread is started so that the GUI thread can resume. Then every other node that is connected to the

inputs of the node that is focused, and for which the output needs to be calculated, is scanned and added to a queue if the outputs are not cached. Finally, for each node, the function `process_template` creates the outputs and populates the cache. When the last node has been worked out, the thread terminates and the GUI thread displays the result.

The library defines which types are possible outputs of a node in the structure `BasicDataType` and `BasicValue`, then declares the types of nodes in the structure `NodeInstruction`. 29 types and 72 instructions have been implemented. Each type has a color associated with it, that is used to paint endpoints with the type and connecting curves between nodes. Nodes are organized in categories and can be added with right click of the mouse. The canvas can be zoomed and panned, helping the user organize the nodes.

An peculiar node is the “String to SVG” one. It takes a Dot file as string as an input and outputs an SVG value. The string is first parsed as an Dot file using the library layout[21], then the resulting graph is converted to a tree that represents an SVG. Then it is converted into string to be able to be parsed again by the library resvg[23]. Finally an image buffer is allocated and the tree is rendered on the pixel map. Since egui library is not optimized to display arbitrary images, the pixel map is then converted to texture so that it may be cached more easily. To save on space the texture is not serialized and is recomputed when needed. The result can be either displayed on screen or saved as a PNG image.

The code for the render of SVG files is implemented in `svg.rs`.

The entry point for the native application is in the file `main.rs` and the entry point for the web application is `web.rs`. To interface with WebAssembly, only three functions are strictly needed: `new`, `start` and `destroy`. These functions are translated to wasm and used as bindings for JavaScript.

To build for web first we invoke the Rust compiler with the command

```
cargo build -p "reaction\_systems\_gui" --release --all-features
--lib --target wasm32-unknown-unknown
```

that builds for the target wasm32. Then using wasm-bindgen[11] we create the appropriate bindings with the command

```
wasm-bindgen "[..]/reaction\_systems\_gui.wasm" --out-dir docs
--no-modules --no-typescript
```

As an additional step we optimize using `wasm-opt` from the library binaryen[3] with

```
wasm-opt "[..]/reaction\_systems\_gui\_bg.wasm" -O2 --fast-math
-o "[..]/reaction\_systems\_gui\_bg.wasm"
```

The code can then be served statically and used in a HTML canvas. Bash scripts are provided that automates this process: `build_web.sh` and `start_server.sh`.

Chapter 5

Testing and Validation

5.1 Tests

During the development of the library some tests were developed in order to test behavior in the changing code. They can be run with `cargo test`. Tests don't cover all code branches, but have been written for pieces of code that might break more easily. Tests are usually present in a separate file as the structure declaration and have the suffix “`_test.rs`” in their name, so that they might be easily recognized.

For example in the workspace `bisimilarity` tests have been written for the algorithms implemented.

```
#[test]
fn bisimilar_paige_tarjan_3() {
    use petgraph::Graph;
    let mut graph_b = Graph::new();

    let node_b_1 = graph_b.add_node(1);
    let node_b_2 = graph_b.add_node(2);
    graph_b.add_edge(node_b_1, node_b_2, 1);
    let node_b_3 = graph_b.add_node(3);
    graph_b.add_edge(node_b_2, node_b_3, 2);
    let node_b_4 = graph_b.add_node(4);
    graph_b.add_edge(node_b_3, node_b_4, 2);

    let mut graph_c = Graph::new();

    let node_c_1 = graph_c.add_node(1);
    let node_c_2 = graph_c.add_node(2);
    graph_c.add_edge(node_c_1, node_c_2, 1);
    let node_c_3 = graph_c.add_node(3);
    graph_c.add_edge(node_c_2, node_c_3, 2);
    let node_c_4 = graph_c.add_node(4);
```

```

graph_c.add_edge(node_c_3, node_c_4, 2);
let node_c_5 = graph_c.add_node(5);
graph_c.add_edge(node_c_1, node_c_5, 1);
graph_c.add_edge(node_c_5, node_c_3, 2);

assert!(bisimilarity(&&graph_b, &&graph_c));
assert!(bisimilarity_ignore_labels(&&graph_b, &&graph_c))
}

```

The macro call `#[test]` instructs rust to treat the function as a test and thus generate the respective code only when building for the test suite. Then the two graphs are constructed; finally the results are tested with the `assert!` macro.

The structure of the two graph can be seen in figure 5.1.

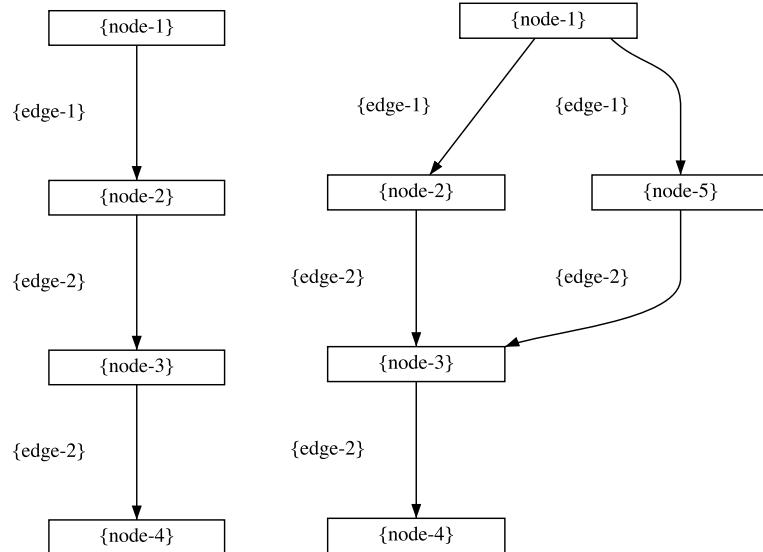


Figure 5.1: Graphs of tests

The behavior of the two systems is the same by ignoring the labels on the edges and by considering them, so they are always bisimilar, so the test should not panic.

The graphs have been generated using ReactionSystemsGUI using the nodes displayed in figure 5.2. Note that to differentiate the two edges exiting `node-1` additional entities `left` and `right` are used, but are not displayed because the function for displaying the edges masks only for `edge-1` and `edge-2`.

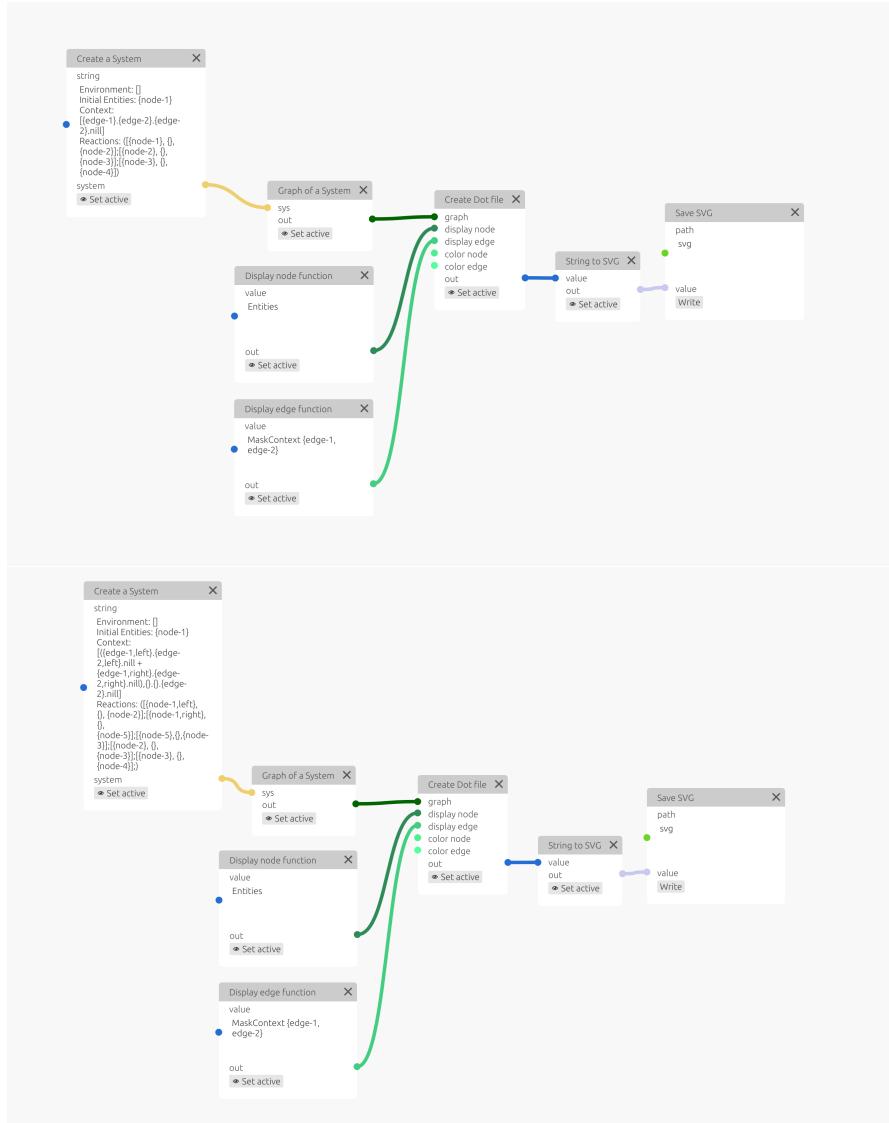


Figure 5.2: Generating graphs using ReactionSystemsGUI

In addition to automated tests, some example inputs are provided in the folder `testing`. The extension `.system` symbolizes system and associated instructions; the extension `.experiment` symbolizes an experiment, see 3.1.13.

These examples were also used to do manual integration testing.

The example `target.system` generates as output:

After 6 steps we arrive at state:

{b}

After 6 steps we arrive at state:

{b}

Environment: [x = {a}.y, y = ({a}.{a, b}.nill + {b}.nill)]

Initial Entities: {a, b}

Context: $[(\{a,b\}.\{a\}.\{a,c\}.x + \{a,b\}.\{a\}.\{a\}.nill)]$

Reactions: ([{a,b}], {c}, {b})

Target > Print,

Target (Limit: 7) > Print,

Target (Limit: 6) > Print,

Target (Limit: 5) > Print

After 6 steps we arrive at state:

{b}

After 5 steps we arrive at state:

{}

The output is correct since, as we can see from the graph of the system in figure 5.4 that the system has as leftmost production six states and the end state has entities set $\{b\}$.

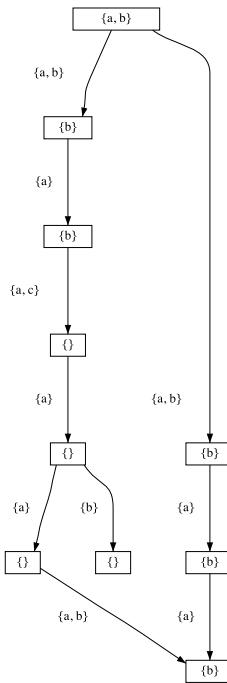


Figure 5.4: Graph of example 5.3a, nodes have available entities as label, edges have the entities provided by the context as labels.

A small Perl script is provided that can convert systems created for the Prolog version

of the program into the syntax described in 3.1.7. The script was used to convert systems that model mutual exclusion (MEX) available in the folder `testing/mex`.

A MEX system is composed of n looping processes, which only one can be in the critical section. Each process is identified by `out_i` if is out of the critical section, `req_i` if has requested to enter the critical section and `in_i` if is in the critical section. Without the token `act_i` no process can change state:

```
[{out_1}, {act_1}, {out_1}];  
[{req_1}, {act_1}, {req_1}];  
[{in_1}, {act_1}, {in_1}];
```

And `act_i` is required to change state:

```
[{out_1, act_1}, {}, {req_1}];
```

Any subset of processes can be activated, the requests are preserved:

```
[{req_1, act_1, act_2}, {}, {req_1}];  
[{req_1, act_1, act_3}, {}, {req_1}];  
...  
[{req_1, act_1, act_n}, {}, {req_1}];
```

Entering the critical section is handled by two other entities: `lock`, which symbolizes that a process is in the critical section, and `done`, which symbolizes that the critical section has been exited. `lock` remains until done:

```
[{lock}, {done}, {lock}];
```

No other processes have to have the lock for a process to enter the critical section:

```
[{req_1, act_1}, {lock, act_2, ..., act_n}, {in_1, lock}];
```

The execution of these examples is more computationally expensive with increasing n . The execution of `mex10.system` took 78791.202 milliseconds to run the instruction 5.1 which generate the graph of the system, converts it to dot format and saves it.

```
Digraph > Dot
  | Entities
  | Context
  | ! "white"
  | ! "black"
  > Save("out.dot")
```

Instruction 5.1: Instruction for MEX systems.

`mex5.system` takes instead 109.148 milliseconds to run.

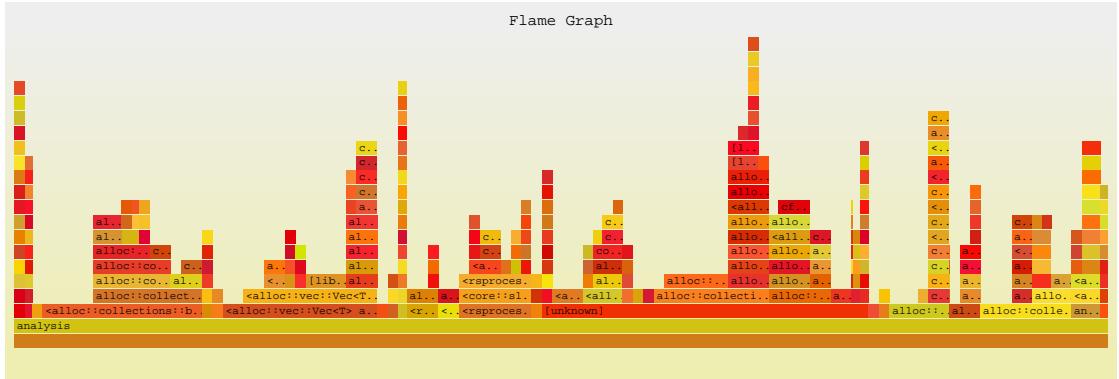


Figure 5.5: Flamegraph of MEX RS with 5 processes.

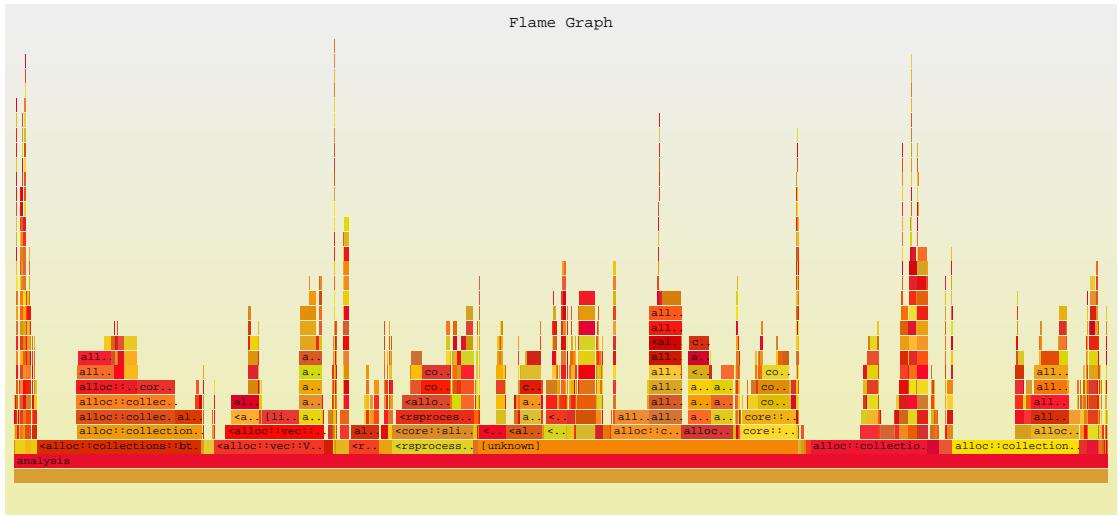


Figure 5.6: Flamegraph of MEX RS with 10 processes.

Performance has been analyzed using perf[17] and flamegraph[15].

Both of the flamegraphs share distinct features, but the one regarding MEX RS with 10 processes has finer resolution since the execution time is much longer. In decreasing order, some of the that took most samples are:

- `unknown` that took 29% of samples,
- `alloc::collections::btree::map::IntoIter::dying_next` that took 14% of samples,
- `<alloc::collections::btree::map::Iter as core::iter::traits::iterator::Iterator>::next` that took 11% of samples,
- `alloc::collections::btree::append::bulk_push` that took 11% of samples,
- `alloc::collections::btree::map::BTreeMap::bulk_build_from_sorted_iter` that took 5,1% of samples,
- ...

From the perf data we can gather that parsing took less than 3% of total time, and that `unknown` refers to the methods belonging to System (section 4.1.5). This behavior is expected since most of the computation is carried by the structure that generates the graph.

5.2 Validation

During development key issues identified from previous projects where performance and usability. The biggest problem with Prolog software is exceeding the stack limit and thus running out of memory. This problem is completely solved by using Rust. Another problem was that of performance. On dot file generation a 2 to 7 times performance improvement is seen, depending on the model simulated.

Usability has been taken into account both for an end user and for a programmer that intend to expand the system: grammar follows general rules largely compatible with previous projects; the grammar is decoupled from the internal representation and thus permits greater maintainability and expandability; the use of traits permits more modularity and the coexistence of multiple types of RS in the same project; domain specific languages allow more efficient and intuitive instruction specification; the graphical user interface presents instructions and methods over reaction systems in a more intuitive way that with just a command line interface; the node system allows greater modularity and for easy additions of new instructions; easy SVG generation reduce the time spent between different software and speeds up the end user's tasks; saving the state of the application allows for lower friction when switching between projects; the web GUI provides a platform agnostic interface for quick development. Thus the original goal to develop a more user-friendly and developer-friendly reaction system modeler has been met.

Chapter 6

Conclusion

Reaction Systems were originally conceived as a theoretical framework to model biochemical processes in living cells. Applying RS theory to practical modeling tasks has been non-trivial. The behavior of a RS can exhibit complex dynamics due to the non-monotonic role of inhibitors. In practice the number of entities and reactions can be large, making manual reasoning unfeasible without proper automated support.

This thesis contributes to bridging the gap between RS theory and practical use by providing a concrete software realization of the formalism. The software is available in the repositories ReactionSystems[19] and ReactionSystemsGUI[20].

6.1 Summary of Contributions

- **New RS Modeling Platform:** This thesis introduced a new software platform for Reaction Systems modeling, analysis and design, implemented in the Rust programming language for high performance and reliability. The software provides both a command-line interface (CLI), native graphical user interface (GUI) and a web based one. This enables users to model RSs either through text-based commands or in an interactive visual environment.
- **Comprehensive Feature Set:** The tool supports a rich set of features: simulation of RS, bisimulation of RSs, trace slicing, graph generation with multiple output formats like Dot, GraphML and SVG, loop analysis, automated conversion between types of RS. These features collectively offer a comprehensive toolkit for analyzing RS behavior.
- **Performance and Design:** The implementation in Rust allows to maximize performance, ensure safety and permit compositionality of future expansions. This eliminates prior memory issues with implementations in Prolog and dramatically improves execution speed. For example generating state-graph visualizations (Dot graphs) is now up to 7 times faster than before. Beyond performance, the software's architecture was designed for clarity and extensibility: the RS grammar is kept independent of internal data structures to simplify maintenance, and core components use Rust traits to support modular extension.

Future work may extend current models by a more in-depth performance analysis and optimization; may extend the methods available in the GUI; may provide new Reaction Systems types, like Quantitative Reaction Systems[14] or Multiset Reaction Systems[6]; may provide methods to convert different structures like boolean networks into RS; or may provide additional test coverage of the current code.

Bibliography

- [1] Luca Aceto, Anna Ingolfsdottir, and Jirí Srba. “The algorithmics of Bisimilarity”. In: *Advanced Topics in Bisimulation and Coinduction* (Sept. 2011), pp. 100–172. doi: [10.1017/cbo9780511792588.004](https://doi.org/10.1017/cbo9780511792588.004).
- [2] Christel Baier, Joost-Pieter Katoen, and Kim Guldstrand Larsen. *Principles of model checking*. MIT Press, The, 2016.
- [3] *binaryen: a compiler and toolchain infrastructure library for WebAssembly, written in C++*. Oct. 27, 2025. url: <https://github.com/WebAssembly/binaryen>.
- [4] *BioResolve web page, a Prolog interpreter for Reaction Systems analysis*. <https://pages.di.unipi.it/bruni/LTSRS/>. 2025.
- [5] Agustín Borgna. *petgraph: Graph data structure library. Provides graph types and graph algorithms*. Version 0.8.3. Sept. 30, 2025. url: <https://github.com/petgraph/petgraph>.
- [6] Paolo Bottoni, Victor Mitrana, and Ion Petre. “Multiset Reaction Systems”. In: *Lecture Notes in Computer Science* (2025), pp. 179–193. doi: [10.1007/978-3-031-97274-4_11](https://doi.org/10.1007/978-3-031-97274-4_11).
- [7] Linda Brodo, Roberto Bruni, and Moreno Falaschi. “A logical and graphical framework for Reaction Systems”. In: *Theoretical Computer Science* 875 (July 2021), pp. 1–27. doi: [10.1016/j.tcs.2021.03.024](https://doi.org/10.1016/j.tcs.2021.03.024).
- [8] Linda Brodo et al. “Slicing analyses for negative dependencies in Reaction Systems Modeling Gene Regulatory Networks”. In: *Natural Computing* (Sept. 2025). doi: [10.1007/s11047-025-10046-5](https://doi.org/10.1007/s11047-025-10046-5).
- [9] Daniel Burgener. *lalrpop: convenient LR(1) parser generator*. Version 0.22.2. May 22, 2025. url: <https://github.com/lalrpop/lalrpop>.
- [10] Rance Cleaveland and Oleg Sokolsky. “Equivalence and preorder checking for finite-state systems”. In: *Handbook of Process Algebra* (2001), pp. 391–424. doi: [10.1016/b978-044482830-9/50024-2](https://doi.org/10.1016/b978-044482830-9/50024-2).
- [11] daxpedda. *wasm-bindgen: Easy support for interacting between JS and Rust*. Version 0.2.105. Oct. 27, 2025. url: <https://github.com/wasm-bindgen/wasm-bindgen>.
- [12] John Ellson. *Graphviz is open source graph visualization software*. url: <https://graphviz.org/doc/info/lang.html>.

- [13] Emil Ernerfeldt. *egui: An easy-to-use immediate mode GUI that runs on both web and native*. Version 0.33.2. Nov. 13, 2025. url: <https://github.com/emilk/egui>.
- [14] Victor Mitrana et al. “Quantitative reaction systems”. In: *2025 5th International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET)* (May 2025), pp. 1–6. doi: [10.1109/iraset64571.2025.11008263](https://doi.org/10.1109/iraset64571.2025.11008263).
- [15] Dirkjan Ochtman. *flamegraph: A simple cargo subcommand for generating flamegraphs, using inferno under the hood*. Version 0.6.10. Nov. 7, 2025. url: <https://github.com/flamegraph-rs/flamegraph>.
- [16] Robert Paige and Robert E. Tarjan. “Three partition refinement algorithms”. In: *SIAM Journal on Computing* 16.6 (Dec. 1987), pp. 973–989. doi: [10.1137/0216062](https://doi.org/10.1137/0216062).
- [17] *perf(1) — Linux manual page*. 6.17-2. Nov. 2025.
- [18] *WebAssembly Core Specification*. Version 2.0. W3C, Sept. 14, 2025. url: <https://www.w3.org/TR/wasm-core-2/>.
- [19] Elvis Rossi. *ReactionSystems*. <https://github.com/elvisrossi/ReactionSystems>. 2025.
- [20] Elvis Rossi. *ReactionSystemsGUI*. <https://github.com/elvisrossi/ReactionSystemsGUI>. 2025.
- [21] Nadav Rotem. *layout-rs: A graph visualization program*. Version 0.1.3. Apr. 24, 2025. url: <https://github.com/nadavrot/layout>.
- [22] Grzegorz Rozenberg and Joost Engelfriet. “Elementary Net Systems”. In: *Lecture Notes in Computer Science* (1998), pp. 12–121. doi: [10.1007/3-540-65306-6_14](https://doi.org/10.1007/3-540-65306-6_14).
- [23] Laurenz Stampfl. *resvg: An SVG rendering library*. Version 0.45.1. Apr. 17, 2025. url: <https://github.com/linebender/resvg>.
- [24] The GraphML Team. *GraphML is a comprehensive and easy-to-use file format for graphs*. url: <http://graphml.graphdrawing.org/>.
- [25] The Rust teams. *A thread-safe reference-counting pointer. ‘Arc’ stands for ‘Atomically Reference Counted’*. url: <https://doc.rust-lang.org/std/sync/struct.Arc.html>.
- [26] The Rust teams. *An ordered set based on a B-Tree*. url: <https://doc.rust-lang.org/std/collections/struct.BTreeSet.html>.
- [27] The Rust teams. *Rust Programming Language*. url: <https://rust-lang.org/>.
- [28] trevyn. *egui_node_graph2: A helper library to create interactive node graphs using egui*. Version 0.7.0. Nov. 3, 2024. url: https://github.com/trevyn/egui_node_graph2.