



Universiteit Leiden

Opleiding Informatica

Checking for Compatibility
in Team Automata

Name: Lau Bannenberg
Date: 28/08/2015
1st supervisor: Dr. H.C.M. Kleijn
2nd supervisor: Dr. H.J. Hoogeboom

BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Checking for Compatibility in Team Automata

BSc Computer Science Thesis

Lau Bannenberg

August 28, 2015

Supervisors **Dr. H. C. M. Kleijn & Dr. H. J. Hoogeboom**

Abstract

The Team Automata model allows us to reason about the communication and collaboration of systems consisting of both human and machine components. Recently [2] provided a rigorous definition for correct interaction among components. In this thesis we will provide means to test systems according to their compatibility requirements. We provide both a theoretical foundation and an implementation in Python. We look at the complexity and scalability of compatibility testing and examine several methods to manage the effort required, including the possibility of splitting the workload into distributed tasks.

Contents

1	Introduction	5
1.1	Research Goals	5
1.2	Organization of this paper	6
2	Preliminaries	7
2.1	Notation	7
2.2	Definitions and theorems from earlier works	8
2.2.1	(Reactive) Transition Systems	8
2.2.2	Computations	8
2.2.3	Systems of RTSs	8
2.2.4	Synchronizations	9
2.2.5	Compatibility	11
2.2.6	Associativity of Synchronous Product	12
3	Fixed Notation	15
4	Main Contributions	17
4.1	Determinism	17
4.2	Finiteness	17
4.3	Detecting livelock	17
4.4	The straightforward approach: exhaustive search for directly incompatible states	19
4.5	The backwards approach, starting in <i>DIS</i>	21
4.6	Pruning the search space	23
4.6.1	Pruning unreachable states from the statespace	25
4.6.2	Pruning useless transitions	25
4.6.3	Pruning useless actions	25
4.6.4	Reducing components	26
4.7	Constructing bottom-up from “correct” components	26
4.7.1	Using Lemma 30	30
5	Software	37
5.1	Other software	39
6	Experiments	41
7	Conclusions	45
7.1	Future Questions	45
	References	47
	Index	48
A	Code	50
B	Automata test sets	70
B.1	Boss/Employee	70
B.2	Coffee	70
B.3	Settlement	71
B.4	Token Ring	72
B.5	Wedding	74
B.6	Passing the Ball	75

1 Introduction

We live in an era of systems that talk to each other. Subsystems designed by different people are used to assemble new systems, which will in turn be used in other systems. But with so many components interacting, can we predict if our systems will always be interacting correctly?

One interesting model for interacting components was introduced by Ellis in [4]. He introduced the idea of *team automata*. These resemble finite state automata, but with some key differences. The most significant of these is that team automata can perform some actions together, as a way of modeling communication (one component sending an instruction to another) and collaboration (components working together on an action as equals). A rigorous definition for the various ways components can act together was given by [7].

Another major property of team automata is that when we construct a team from these automata, the resulting team automaton can also be used as a component in further constructions. We can then construct a system as the product of teaming up subsystems that have subsystems of their own. This allows division of effort among developers and re-use of components.

Team automata can represent machine components of course, but also people interacting with machines, or even people interacting with each other without any machines at all. In our examples we will use a dysfunctional workplace, a coffee machine and customer, a groom who cannot make up his mind to marry, and a game of passing the ball around to show and test various aspects of team automata. We also have a few other cases that were useful for testing but are not ideal examples of any specific property. These are all included in the appendix.

Team automata can be used to reason about the correctness of the interaction of components. Carmona and Cortadella first defined the idea of Input/Output Compatibility for pairs of automata in [1]. Two main communication failures were to be avoided: message loss and deadlock. Message loss occurs when one automaton wants to send a signal and the other automaton is not ready to receive it. Deadlock occurs if both automata require input from one another before they can proceed themselves, leading to endless waiting. They gave a rigorous definition for when an assemblage of automata could be called compatible. From this definition they suggested a decision procedure for determining compatibility, but did not elaborate on it.

In [2] Carmona and Kleijn lifted several of the restrictions of the model from [1], most significantly that compatibility was defined between a pair of automata. They instead produced a definition of compatibility for a system composed of an arbitrary number of automata. In addition, they went into more detail about the consequences of iterated team-building on compatibility. This was partially in an effort to arrive at an algorithm to determine whether a system was compatible. However, such an algorithm was not found.

That will be the objective of this thesis: to find an algorithm to determine the compatibility of a system of any number of automata, and analyze the complexity and scalability of the algorithm. Besides an abstract algorithm, we will also produce an implementation in Python.

1.1 Research Goals

The prime goal is to find an algorithm to decide whether a set of n components forms a compatible system.

- Analyze the algorithmic complexity of the solution.
- Try to keep the complexity manageable.
- Implement the algorithm and experiment with it. See if anything noteworthy happens.
- Analyze whether there is an advantage to analyzing a team through analyzing subteams and combining them, instead of analyzing the whole team at once.
- While we try several alterations to team automata to improve our algorithm, we must ensure the results remain valid. Specifically, we require the following properties to remain the same in any altered automaton:
 1. The *behaviour* exhibited by the alternate automaton must be the same as the original automaton. To an outside observer everything they can and can't do appears the same.

2. The alternate automaton is internally compatible if and only if the original automaton would have been compatible.
 3. If the alternate automaton is used in an iterative process of team-building, it will have the same effects on its environment as the original. We will not attempt any optimization that will hinder us if we later decide to use the result as an ingredient for an even greater system.
- We will develop such an algorithm in theory, and also provide a Python implementation. In this implementation we will try to include such features as might be useful to a developer who wants to iteratively develop a system by designing and testing subsystems before integrating them.
 - We aim to stay inside the limits already created by earlier works, notably [2], so that our results will be applicable in any case that meets their specifications.

1.2 Organization of this paper

We assume the reader is familiar with finite state automata and the basics of language theory, as well as sets. Because team automata are less well-known, we will introduce them and repeat some of the results from other authors to provide the necessary grounding for our own work.

We start in the **preliminaries** by introducing some notation and then revisit the definitions and theorems from (particularly) [2] that we build on. This section includes some examples of simple team automata to outline just what communication failures we wish to detect.

In our **main contributions** we develop a theoretical foundation for our algorithms, present our algorithms and analyze them. Then we study the possibilities for further improvement, first by looking at sharpening the theory.

In the **software** section we implement the algorithms in Python and explain their usage.

In the **experiments** section we do a few introductory experiments with heuristics to optimize our algorithms.

After **conclusions** follows an **appendix** with Python source code and specification of the components used in examples and experiments.

2 Preliminaries

2.1 Notation

Definitions and theorems will usually be introduced informally followed by a formal description. When a keyword is defined it will be ***bold and italicized***. Later references to it will be *italicized*.

We use the word ‘component’ to refer to a *reactive transition system* that is used as a building block in a *team automaton*. Since team automata can be used as the building blocks for new automata, components might have components of their own.

Definition 1 (*alphabet*).

An **alphabet** Σ is a finite set of symbols. Σ^* denotes a set of strings of 0 or more symbols from Σ .

In this thesis we will commonly use Σ to denote an alphabet of *actions* of a *transition system*. Each symbol $a \in \Sigma$ represents one action of that TS.¹ We deviate from one common style of notation of languages: we will not use superscripts to indicate repetition, except with the Kleene star. For example, a^2 does not mean aa , but $(Q\Sigma)^*$ does indicate a concatenation of 0 or more strings consisting of an element from Q followed by an element of Σ .

When we take the Cartesian product of sets, we can recover elements from the ingredient sets with the **projection** function:

Definition 2 (*projection*).

Let $V = V_1 \times V_2 \times \dots \times V_n$ be sets, $v = (v_1, v_2, \dots, v_n)$ and $\forall i \in \{1, \dots, n\} : v_i \in V_i$, then:

$$\text{proj}_i(v) = v_i$$

We will use subscripted indices to identify components (which are tuples) and the items that belong to them: $Q_{\mathcal{T}_i}$ is the set of states belonging to component \mathcal{T}_i . Usually we will use a shorthand, where we write Q_i instead of $Q_{\mathcal{T}_i}$. If the shorthand would prove confusing, we will use the full subscript instead. We will use superscripted indices to identify the individual items belonging to some set: q_i^1 and q_i^2 are states of Q_i . Although sets of states are not ordered (q^2 may be reachable before q^1), we will normally call the initial state q^0 (components start in an initial state).

When describing *transitions* from one state to another, we often describe it in the form $t = (q, a, q')$, where q is the current state, a is the action, and q' is the next state.

For subsets, we will use $A \subseteq B$ to indicate that A is a subset of B, possibly equal to B. If A is a proper subset of B we will write $A \subsetneq B$.

In pseudocode, we will use the following notations:

- $\text{varname} \leftarrow \text{value}$ assignment
- $[a^1, a^2, \dots, a^n]$ is a list can contain duplicate elements. We can query it for membership with $\text{element} \in \text{listname}$, and look at the first and last elements with list.head and list.tail . We can add and remove elements from the head and tail: $\text{listname.prepend}(\text{element})$, $\text{listname.append}(\text{element})$, listname.remove and $\text{listname.removelast}$; and an empty list can be represented as $[]$. When we access the i -th element of a list we write $\text{listname}[i]$.
- $\{a^1, a^2, \dots, a^n\}$ is a set with unique elements. We can query it for membership with $\text{element} \in \text{setname}$ and add and remove elements with $\text{setname.add}(\text{element})$ and $\text{setname.remove}(\text{element})$. An empty set is represented as \emptyset .
- Functions are written as $\text{functionname}(\text{argument 1}, \text{argument 2}, \dots, \text{argument } n)$.

For example, $\text{func}([a], b)$ is a call to function func with as arguments a list containing just one element (a) and another element (b) that is not in a list.

¹This alphabet could be empty, but that generates trivial automata that cannot change state nor in any way influence other automata. Therefore they have no impact on the compatibility of systems they are included in. It is worth noting that the material in this thesis applies to empty-alphabet automata just as much as it does to normal automata. This becomes relevant because some pruning techniques can remove actions from automata, possibly leaving their alphabets empty.

2.2 Definitions and theorems from earlier works

This thesis builds heavily on earlier works, and uses definitions, theorems and proofs presented there. They are repeated here, with attribution:

2.2.1 (Reactive) Transition Systems

From [2]

Reactive transition systems strongly resemble (finite) state automata. Their main difference is that in an RTS actions are classified on the individual automaton level as input, output or internal. One automaton might send “button!” as an output action and another receive “button?” as input signal. Output actions represent a component taking action to influence its environment or to direct another component. Input actions represent the automaton reacting to such stimuli. Internal actions are hidden from the environment, happening entirely within the “black box”.

There are two other noteworthy differences: the Team Automaton model does not prohibit infinite sets of states, although in practice this is rarely used. Also, there is no notion of a set of accepting states. An RTS often represents an ongoing process, unlike a finite state automaton that is trying to finish recognize a word.

Definition 3 (*transition system*).

A **transition system** or **TS** is a tuple $\mathcal{A} = (Q, \Sigma, \delta, I)$:
 Q is a set of **states**, possibly infinite.
 Σ is an *alphabet* of **actions**, with $Q \cap \Sigma = \emptyset$
 $\delta \subseteq Q \times \Sigma \times Q$ is a set of (labeled) **transitions**.
 $I \subseteq Q$ is a set of initial states.

A TS \mathcal{A} is **reactive** (a **reactive transition system** or **RTS**) if its set of actions Σ is the union of three pairwise disjoint sets of input actions Σ^{inp} , output actions Σ^{out} and internal actions Σ^{int} . The union of input and output actions is the set of external actions Σ^{ext} .

If \mathcal{A} is an RTS as above, then it can also be represented, with the explicit partition of its actions, as $(Q, (\Sigma^{inp}, \Sigma^{out}, \Sigma^{int}), \delta, I)$.

Graphically, output and input actions may be suffixed with the symbol ! and ? respectively to indicate their role, and internal actions can be indicated with the @ suffix.

For a TS $\mathcal{A} = (Q, \Sigma, \delta, I)$ and an action $a \in \Sigma$, the set of a -transitions in \mathcal{A} is $\delta_a = \delta \cap (Q \times \{a\} \times Q)$.

Action a is said to be **enabled** in \mathcal{A} at state $q \in Q$, denoted by $a \text{ EN}_{\mathcal{A}} q$, if there exists a q' such that $(q, a, q') \in \delta$. We say that q is **input-enabled** if $\Sigma^{inp} \subseteq \{a \mid a \text{ EN}_{\mathcal{A}} q\}$.

2.2.2 Computations

From [2]

Definition 4 (*computation*).

A **computation** of \mathcal{A} is a (possibly infinite) sequence $\gamma = q^0 a^1 q^1 a^2 q^2 \dots$ where $q^0 \in I$ and $q^i \in Q$ for all $1 \leq i$, and $(q^i, a^{i+1}, q^{i+1}) \in \delta$, for all $0 \leq i$.

A state q^i of \mathcal{A} is **reachable** if it occurs in any computation of \mathcal{A} ; in other words there is a series of transitions beginning in the initial state that eventually reaches q^i . All states that are not reachable are **unreachable**.

2.2.3 Systems of RTSs

From [2]

A *team automaton* is constructed from a finite number of components which are RTSs. These have to be composable with each other. For that to be true, their internal actions have to be truly internal;

the internal actions cannot also appear in the alphabet of other components of the system-to-be. This is because internal actions are private, not available for synchronization.²

Definition 5 (*composable system*).

A collection $\mathcal{S} = \{\mathcal{A}_i | 1 \leq i \leq n\}$ of component automata (RTSs) $\mathcal{A}_i = (Q_i, \Sigma_i, \delta_i, I_i)$, with Σ_i partitioned into internal, input and output actions $(\Sigma_i^{int}, \Sigma_i^{inp}, \Sigma_i^{out})$, respectively, forms a **composable system** if $\Sigma_i^{inp} \cap \bigcup_{j=1, j \neq i} \Sigma_j = \emptyset$ for all $1 \leq j \leq n$.

Given a composable system, a TA is constructed by selecting *synchronizations* on common actions of its components, while retaining the local effect of each action. The state space Q of such a TA is $\prod_{i=1}^n Q_i$. This means that each of a team automaton's states is a product of the states of its components. When it executes an action involving some of its components, those parts of the global state change. Components that do not participate in the action do not change state.

Definition 6 (*team automata*).

Let \mathcal{S} be a composable system as specified above. Let $a \in \bigcup_{i=1}^n \Sigma_i$.

The **complete transition space** of a in \mathcal{S} is

$\Delta_a(\mathcal{S}) = \{(p, a, p') | p, p' \in \prod_{i=1}^n Q_i, (\text{proj}_i(p), a, \text{proj}_i(p')) \in \delta_i \text{ for some } i \leq 1 \leq n,$
and for all $1 \leq i \leq n$ either $(\text{proj}_i(p), a, \text{proj}_i(p')) \in \delta_i$ or $\text{proj}_i(p) = \text{proj}_i(p')\}$.

A **team automaton** or **TA** over \mathcal{S} is an

RTS $\mathcal{T} = (\prod_{i=1}^n Q_i, \Sigma, \delta, \prod_{i=1}^n I_i)$ with $\delta_a \subseteq \Delta_a(\mathcal{S})$ for all $a \in \Sigma^{ext}$ and $\delta_a = \Delta_a(\mathcal{S})$ for all $a \in \Sigma^{int}$.

Its set of actions is

$\Sigma = \bigcup_{i=1}^n \Sigma_i$ with $\Sigma^{int} = \bigcup_{i=1}^n \Sigma_i^{int}$, $\Sigma^{out} = \bigcup_{i=1}^n \Sigma_i^{out}$ and $\Sigma^{inp} = \bigcup_{i=1}^n \Sigma_i^{inp} - \bigcup_{i=1}^n \Sigma_i^{out}$.

Let $\mathcal{S} = \{\mathcal{A}_i | 1 \leq i \leq n\}$ be a *composable system* as specified before. \mathcal{S} represents a collection of interacting components. Let $a \in \bigcup_{i=1}^n \Sigma_i$ be an action of \mathcal{S} . Then $\text{dom}(\mathcal{S}, a) = \{i | a \in \Sigma_i\}$ is the **domain** of a in \mathcal{S} . The **output domain** of a in \mathcal{S} is $\text{outdom}(\mathcal{S}, a) = \{i | a \in \Sigma_i^{out}\}$. Similarly, $\text{inpdom}(\mathcal{S}, a) = \{i | a \in \Sigma_i^{inp}\}$ is the **input domain** of a in \mathcal{S} . We say that an external action a is **output-domain enabled** (**input-domain enabled**) at p if $a \text{ EN}_{\mathcal{A}_i} \text{proj}_i(p)$ for all $i \in \text{outdom}(\mathcal{S}, a)$ (for all $i \in \text{inpdom}(\mathcal{S}, a)$). It is **domain enabled** if it is both input-domain and output-domain enabled.

An action is said to be **communicating** in \mathcal{S} if it has a non-empty output domain and a non-empty input domain. We write $\Sigma^{com} = \bigcup_{i=1}^n \Sigma_i^{out} \cap \bigcup_{i=1}^n \Sigma_i^{inp}$. External actions which are not communicating in \mathcal{S} may become so in an extension of \mathcal{S} when components have been added. We say that \mathcal{S} is **complete** if each of its external actions is communicating.

2.2.4 Synchronizations

From [7], [2]

As shown before, the transitions of a TA \mathcal{T} over a composable system \mathcal{S} form a subset of that system's complete transition space:

$$\delta_a \subseteq \Delta_a(\mathcal{S}) \text{ for all } a \in \Sigma^{ext} \text{ and } \delta_a = \Delta_a(\mathcal{S}) \text{ for all } a \in \Sigma^{int}$$

The latter means that δ inherits all the transitions based on internal actions of all its components. But let us focus on the former now. It defines that δ_a can contain any of the possible combinations of transitions using a from the components of \mathcal{T} , but not necessarily all or even any of them.

For example, let us take \mathcal{T} to be the TA over $\mathcal{S} = \{\mathcal{A}_1, \mathcal{A}_2\}$. Furthermore, $a \in \Sigma_1^{out}$ and $a \in \Sigma_2^{inp}$, and it is involved in the transitions $(q, a, q') \in \delta_{\mathcal{A}_1}$ and $(r, a, r') \in \delta_{\mathcal{A}_2}$. Now $\delta_{\mathcal{T}_a}$ could include a transition that moves both \mathcal{A}_1 and \mathcal{A}_2 to the next state. This would then be $((q, r), a, (q', r'))$. That would represent \mathcal{A}_1 sending a signal and \mathcal{A}_2 receiving it. But it could also include a transition that moves \mathcal{A}_1 but ignores \mathcal{A}_2 , so $((q, x), a, (q', x))$. That would represent \mathcal{A}_1 sending a signal but \mathcal{A}_2 not responding to it.

²If two components happen to have internal actions with the same name, renaming those actions can make them composable. This could occur, for example, if a system is constructed with several identical components, such as a car with two front wheels. Although the wheels have the same sort of internal actions, those are still distinct actions, not shared – otherwise they would not be internal actions.

We call such a choice of component-transitions to use in a transition for the new TA a **synchronization**. It is possible to select synchronizations ad-hoc, picking some and not others, but normally we will use some rule that selects all synchronizations that meet some criterium, and excludes all others. Originally introduced in [4], this concept was rigorously developed in [7].

One obvious rule is that if an output-action happens in one or more components, all components able to process that action as input must also participate in the action. This would represent for example workers responding to the instructions of supervisors. Another, slightly less obvious rule is to require all components that possess a to participate in every execution of a by the team; if even one of them is not currently ready, a cannot execute. If they are all available, and a is executed, then they all participate. This is sometimes called “action-indispensible” when the rule is applied to a single action, or “synchronous product” if it applies to all actions.

Definition 7 (*synchronous product*).

Let the set of transitions $\chi^{\mathcal{S}}$ be defined by

$$\chi_a^{\mathcal{S}} = \{(p, a, p') \in \Delta_a(\mathcal{S}) \mid (\text{proj}_i(p), a, \text{proj}_i(p')) \in \delta_i \text{ for all } 1 \leq i \leq n \text{ such that } a \in \Sigma_i\} \text{ for all } a \in \Sigma$$

The **synchronous product** of a composable system \mathcal{S} , denoted by $\chi(\mathcal{S})$, is the TA over \mathcal{S} with $\chi^{\mathcal{S}}$ as its set of transitions.

In [2] *compatibility* was defined specifically for synchronous product teams. In addition, it was shown there that many other synchronization schemes can be algorithmically converted to a synchronous product system while preserving their compatibility properties. Therefore in the rest of this work, we will only focus on determining compatibility for synchronous product automata.

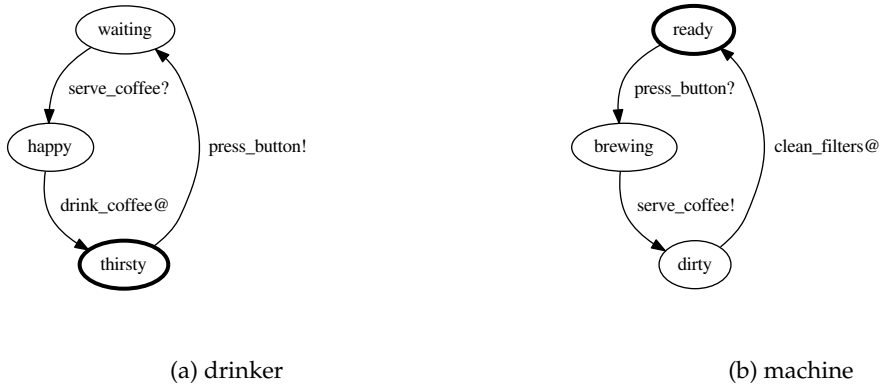


Figure 1: A coffee drinker and a coffee machine.

In Figure 1 we see a coffee consumer and a machine. Their initial states are marked with the thicker line around the state. The consumer starts out thirsty, presses a button (for him an output action), and waits to receive coffee (input). He drinks it (an internal action) and then he becomes thirsty again. Meanwhile the coffee machine is waiting for orders when its button is pushed (input); it then serves coffee (output), and has to clean its filters (internal). When we take the *synchronous product* of these two components we get Figure 2, which shows us only the reachable states of the system.

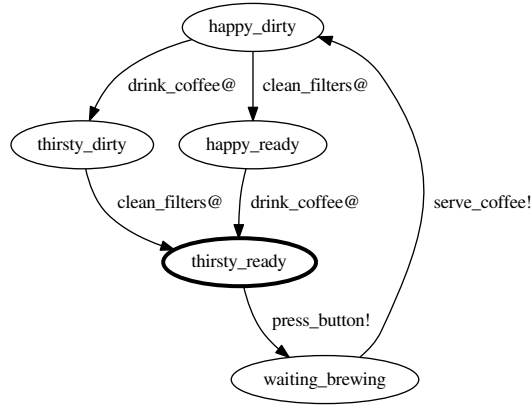


Figure 2: The synchronous product of the coffee drinker and the machine. Only reachable states and useful transitions are shown.

2.2.5 Compatibility

From [2]

An *RTS* \mathcal{A} has a **livelock** if there is an infinite sequence q^1, q^2, q^3, \dots of states in Q and an infinite sequence a^1, a^2, a^3, \dots of internal actions in Σ such that $(q^i, a^i, q^{i+1}) \in \delta$ for all $i \geq 1$. A **livelock-free** RTS will always ultimately execute an external action or terminate.

As defined in [2], a composable system is compatible if there is a **compatibility relation** \mathcal{R} between the states of the components of the new TA;

- The initial states are in the relation.
- Next-states reached by executing actions are in the relation.
- Receptiveness: for every communicating action, if all components that have it as an output action can currently execute (“send”) it, then all components that have it as an input action must also be ready to execute (“receive”) it.
- Deadlock-freeness: whenever the components of the input domain for a communicating action a are waiting for input, the system must be enabled for an action. (a does not need to be possible just yet, but *some* action must be.)

(Also, the team will eventually take communicating actions, not just an endless series of internal actions. Therefore livelock-freeness is also required.)

It should be noted that this system can still end up in a situation where no more actions are possible, even though the above criteria hold. No internal action is enabled in any component, and for every external action the output domain and input domains are both not enabled. In this case, the system has reached a **terminating** state. A TA capable of reaching a terminating state is also called terminating. Termination is not considered a communications failure.

When applied to a system using *synchronous product*, this leads to the following definition:

Definition 8 (*compatibility*).

Let \mathcal{S} be a *composable system* as before. Then $\mathcal{R} \subseteq \prod_{i=1}^n Q_i$ is a compatibility relation for \mathcal{S} if $\prod_{i=1}^n I_i \subseteq \mathcal{R}$ and for all $p \in \mathcal{R}$ the following conditions are satisfied.

1. *Non-communicating Progress*: For all $a \in \bigcup_{i=1}^n \Sigma_i - \Sigma_{com}$, if $a \text{ en}_{\mathcal{A}_i} \text{proj}_i(p)$ for all $i \in \text{dom}(\mathcal{S}, a)$, then $p' \in \mathcal{R}$, whenever $(p, a, p') \in \chi_a^{\mathcal{S}}$.
2. *Receptiveness*: For all $a \in \Sigma_{com}$, if a is output-domain enabled at p , then a is input-domain enabled at p , and $p' \in \mathcal{R}$ whenever $(p, a, p') \in \chi_a^{\mathcal{S}}$.

3. *Deadlock-freeness*: If some action $a \in \Sigma_{com}$ is input-domain enabled at p , then there are $b \in \bigcup_{i=1}^n \Sigma_i$ and $p' \in \prod_{i=1}^n Q_i$ such that $(p, b, p') \in \chi^S$.

S is said to be **compatible** if each of the component automata \mathcal{A}_i is *livelock-free* and there exists a compatibility relation for S . If a relation \mathcal{R} fulfills conditions 1, 2, or 3, then it is called **progressive**, **receptive** or **deadlock-free**, respectively.

As examples of a *incompatible* systems, let us first consider the coffee machine and drinker from the previous section. Although the graph seems to suggest this system works fine, it does experience message loss. If the drinker finishes his coffee before the machine performs the ‘clean_filters@’ action, then he “can’t” press the button. The definition of *synchronous product* requires that all components with the ‘press button’ action participate in it and the machine is not ready yet. What we have here is an example of message loss: it turns out this system fails the *receptiveness* requirement for compatibility.

As an example of deadlock, consider Figure 3 which shows an employee who refuses to work harder until he gets a raise, and a boss who refuses to give a raise until the employee works harder. Getting a raise is an input for the employee and output for the boss, and working harder is output for the employee and input for the boss.




Figure 3: A deadlocked workplace

We should point out that in both the coffee and the workplace example, only reachable states were shown. If the employee somehow started out happy, then the system would not be deadlocked. And because synchronous product forbids actions unless all action-owners participate, we do not get to see what happens when the drinker pushes a button when the machine is not ready.

As an example of livelock, consider Figure 4 which shows us a groom-component (used in the “wedding” case found in the appendix). The groom has the possibility of endlessly hesitating before saying “I do.” Such eternal doubting is clearly a communication failure at such an occasion.

2.2.6 Associativity of Synchronous Product

(The following results from [2] will not be used until section 4.7, but to clearly separate our own work from that of others, we introduce them here.)

We say that two RTSs $\mathcal{A}_1, \mathcal{A}_2$ are **action-complementary** if they have no common output/input actions, e.g., $\Sigma_1^{inp} \cap \Sigma_2^{inp} = \emptyset$ and $\Sigma_1^{out} \cap \Sigma_2^{out} = \emptyset$.

Furthermore, \mathcal{A}_2 **collaborates** with \mathcal{A}_1 if for every state of the Cartesian product, $(p_1, p_2) \in Q_1 \times Q_2$ and every $a \in (\Sigma_1^{inp} \cap \Sigma_2^{inp}) \cup (\Sigma_1^{out} \cap \Sigma_2^{out})$, a is enabled at p_2 whenever a is enabled at p_1 .

Theorem 9. [2]

If $S = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ is *compatible* then $S' = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ is *compatible*.

We repeat the proof of the following lemma here as a prelude to our discussion of a possible variant on it in 4.7.

Lemma 10. [2]

If $S = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ is compatible and $\{\mathcal{A}_1, \mathcal{A}_2\}$ is *receptive* and *deadlock-free*, then $S' = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ is *compatible* provided that \mathcal{A}_3 *collaborates* with $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$.




Figure 4: A groom who can't make up his mind.

Proof. Since internal actions do not communicate, $\chi(\mathcal{A}_1, \mathcal{A}_2)$ is livelock-free if and only if \mathcal{A}_1 and \mathcal{A}_2 are both livelock-free.

Let \mathcal{R} be a compatibility relation for \mathcal{S} . We will prove that $\mathcal{R}' = \{(p_1, p_2, p_3) \mid ((p_1, p_2), p_3) \in \mathcal{R} \text{ and } (p_1, p_2) \text{ reachable in } \chi(\{\mathcal{A}_1, \mathcal{A}_2\})\}$ is a compatibility relation for \mathcal{S}' .

Clearly, \mathcal{R}' includes $I_1 \times I_2 \times I_3$.

Let $p = (p_1, p_2, p_3) \in \mathcal{R}'$ with $((p_1, p_2), p_3) \in \mathcal{R}$ and (p_1, p_2) reachable in $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$.

Non-communicating progress: Let a be a non-communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an action of \mathcal{A}_i for $i = 1, 2, 3$. Let $(p, a, p') \in \chi_a^{\mathcal{S}}$ for some $p' = (p'_1, p'_2, p'_3)$.

We have to prove that $p' \in \mathcal{R}'$ and (p'_1, p'_2) reachable in $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$.

Observe first that a is also a non-communicating action of \mathcal{S} . Moreover, in case a is an input (output) action of \mathcal{A}_1 or \mathcal{A}_2 or both, it is enabled at (p_1, p_2) in $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$. Consequently, a is enabled at $((p_1, p_2), p_3)$ and $((p_1, p_2), p_3), a, ((p'_1, p'_2), p'_3) \in \chi_a^{\mathcal{S}}$ independent of whether it is an input or output or internal action of \mathcal{S}' . Then by the non-communicating progress property of \mathcal{R} it follows that $((p'_1, p'_2), p'_3) \in \mathcal{R}$ and hence $p' \in \mathcal{R}'$.

Receptiveness: Let a be a communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an output action of \mathcal{A}_i for $i = 1, 2, 3$.

We have to prove that a is also enabled for those \mathcal{A}_i for which it is an input action and, moreover, if $(p, a, p') \in \chi_a^{\mathcal{S}'}$ for some $p' = (p'_1, p'_2, p'_3)$, then $p' \in \mathcal{R}'$.

1. If a is not a communicating action of \mathcal{S} , then it must be the case that a is a communicating action of $\{\mathcal{A}_1, \mathcal{A}_2\}$ and a is not an input action of \mathcal{A}_3 . The receptiveness of $\{\mathcal{A}_1, \mathcal{A}_2\}$ implies that

a is enabled at (p_1, p_2) . Thus, a is enabled at the \mathcal{A}_i for which it is an input action. Further, $((p_1, p_2), p_3), a, ((p'_1, p'_2), p'_3) \in \chi_a^{\mathcal{S}}$ and the non-communication progress property for \mathcal{R} yields $((p'_1, p'_2), p'_3) \in \mathcal{R}$. Hence $p' \in \mathcal{R}'$.

2. If a is a communicating action of \mathcal{S} , then either a is a communicating action of $\{\mathcal{A}_1, \mathcal{A}_2\}$ and an input action of \mathcal{A}_3 ; or a is an output (input) action of either \mathcal{A}_1 or \mathcal{A}_2 or both, and an input (output) action of \mathcal{A}_3 . In the first case we use again the receptiveness of $\{\mathcal{A}_1, \mathcal{A}_2\}$ together with the receptiveness of \mathcal{R} , to conclude that a is enabled at those \mathcal{A}_i for which it is an input action and $((p_1, p_2), p_3), a, ((p'_1, p'_2), p'_3) \in \chi_a^{\mathcal{S}}$ with $((p'_1, p'_2), p'_3) \in \mathcal{R}$. Hence in this case $p' \in \mathcal{R}'$. For the latter case, the receptiveness of \mathcal{R} suffices to finish the proof.

Deadlock-freeness: Let a be a communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an input action of \mathcal{A}_i for $i = 1, 2, 3$.

We have to prove that there exists an action b in \mathcal{S} such that $(p, b, p') \in \chi^{\mathcal{S}'}$ for some $p' = (p'_1, p'_2, p'_3)$.

1. If a is an input action of $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ or an input action of \mathcal{A}_3 , then the deadlock-freeness of \mathcal{S} implies that there exist b and $((p'_1, p'_2), p'_3)$ such that $((p_1, p_2), p_3), b, ((p'_1, p'_2), p'_3) \in \chi^{\mathcal{S}}$ and we are done.
2. The remaining case is that a is communicating in $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and not an input action of \mathcal{A}_3 . Now we use the deadlock-freeness of $\{\mathcal{A}_1, \mathcal{A}_2\}$ to infer that there exists an action c and (p'_1, p'_2) such that $((p_1, p_2), c, (p'_1, p'_2)) \in \chi^{\{\mathcal{A}_1, \mathcal{A}_2\}}$. Again we have different cases to distinguish:

If c is not an action of \mathcal{A}_3 , we are done: $(p, c, (p'_1, p'_2, p'_3)) \in \chi^{\mathcal{S}'}$.

The remaining cases are that (i) either c is an input action of $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and an output action of \mathcal{A}_3 or (ii) c is an output action of $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and an input action of \mathcal{A}_3 or (iii) c is an output (input) action of $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and an output (input, respectively) action of \mathcal{A}_3 .

- (i) As before, the deadlock-freeness of \mathcal{S} implies that there exist $((p'_1, p'_2), p'_3)$ such that $((p_1, p_2), p_3), c, ((p'_1, p'_2), p'_3) \in \chi^{\mathcal{S}}$ and we are done.
- (ii) As before, the receptiveness of \mathcal{R} implies that there exist b and $((p'_1, p'_2), p'_3)$ such that $((p_1, p_2), p_3), b, ((p'_1, p'_2), p'_3) \in \chi^{\mathcal{S}}$ and we are done.
- (iii) Due to the fact that \mathcal{A}_3 collaborates with $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$, there exists $((p'_1, p'_2), p'_3)$ such that $((p_1, p_2), p_3), b, ((p'_1, p'_2), p'_3) \in \chi^{\mathcal{S}}$ and we are done.

□

Lemma 11. [2]

The following implications hold:

- If $\{\mathcal{A}_1, \mathcal{A}_2\}$ is *complete*, then $\mathcal{A}_1, \mathcal{A}_2$ are *action-complementary*.
- If $\mathcal{A}_1, \mathcal{A}_2$ are *action-complementary*, then \mathcal{A}_1 *collaborates* with \mathcal{A}_2 .

Theorem 12. [2]

If $\mathcal{S} = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ is *compatible* and $\{\mathcal{A}_1, \mathcal{A}_2\}$ is *receptive* and *deadlock-free*, then $\mathcal{S}' = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ is *compatible* if either:

- \mathcal{S} is *complete*
- $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and \mathcal{A}_3 are *action-complementary*
- \mathcal{A}_3 *collaborates* with $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$

Corollary 13. [2]

If $\mathcal{S} = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ is a *compatible* and *complete* system and $\{\mathcal{A}_1, \mathcal{A}_2\}$ is *compatible*, then $\mathcal{S}' = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ is *compatible*.

3 Fixed Notation

We now fix the following notation for the remainder of this thesis:

- $\mathcal{A} = (Q, \Sigma, \delta, I)$ is a *reactive transition system*.
- \mathcal{S} is the *composable system* over $\{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$.
- $\mathcal{X}(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n)$ is the synchronous product of the *composable system* $\{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$
- $\mathcal{T}^i = \mathcal{X}(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_i)$ where $1 \leq i \leq n$ and $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_i \in \mathcal{S}$.
- $\mathcal{T} = \mathcal{X}(\mathcal{S}) = \mathcal{X}(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n)$.

Informally, \mathcal{A} and \mathcal{T} are both RTSs, however, \mathcal{T} will only be used to refer to an RTS created as the *synchronous product* of other RTSs. If we want to describe a property that holds for any RTS, we will refer to \mathcal{A} . Likewise, we will use \mathcal{A}_i to refer to the i -th component of \mathcal{T} .

\mathcal{T}^i represents a team over the first i components of \mathcal{S} . As \mathcal{S} has n components, $\mathcal{T} = \mathcal{T}^n$.

4 Main Contributions

4.1 Determinism

A deterministic TS has at most one initial state and for every state, every enabled action has exactly one next state.

Definition 14 (*determinism*).

A TS is **deterministic** if the following conditions are both true:

1. $|I| \leq 1$
2. $\forall p \in Q, \forall a \in \Sigma : (p, a, q) \in \delta \wedge (p, a, q') \in \delta \implies q = q'$

If a TS is not deterministic it is **nondeterministic**.

For a deterministic TS we can represent the transitions of δ as a function:

$$\delta(q, a) = \begin{cases} q' & \text{if } (q, a, q') \in \delta \\ q & \text{otherwise} \end{cases}$$

The results of this thesis are applicable to nondeterministic automata as well, with the following adaptations. The transition function $\delta(\text{current}, \text{action})$ would map to a set of next states rather than a single one. Everywhere in our algorithms where we evaluate a current state and action to determine the next state, we instead need to determine all possible next states and handle them. Lastly, we need to run our algorithms for all possible initial states.

So, while there is no theoretical barrier to using nondeterministic automata, for the purpose of clarity and ease of implementation in programming, we will restrict ourselves to deterministic automata.

4.2 Finiteness

The team automata concept allows automata with infinite states, but those pose additional problems for determining compatibility. Because compatibility is defined as a relation of states which all have certain properties, we might have to evaluate infinitely many states to decide whether the relation in fact exists. We will focus on the case of team automata with finite states. The algorithms used in this thesis are not guaranteed to terminate in the case of team automata with infinite states. In fact, they terminate only when only a finite number of states is actually reachable.

4.3 Detecting livelock

Definition 8 requires that components be livelock-free.

Lemma 15 (*persistence of livelock-freeness*).

Each of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ is livelock-free $\iff \mathcal{T}$ is livelock-free.

Proof. Suppose \mathcal{T} has a livelock. This means that starting in some state of \mathcal{T} we can embark on an infinite series of transitions using only internal actions. However, the internal actions of the components of \mathcal{T} do not affect each other's state. One component cannot enable a livelock in another component that did not exist in that component in isolation. So at least one of those components must have a livelock as well.

Now suppose \mathcal{A}_i has a livelock. Since the state of the other components of \mathcal{T} does not matter to the internal transitions \mathcal{T} inherits from \mathcal{A}_i , this livelock continues to exist in \mathcal{T} . □

Lemma 15 makes it possible to prevent livelock during a correct-by-construction bottom-up approach. We can verify for the initial components that they are livelock-free, and thus be certain that they will never cause a livelock when used later-on. We have chosen to insist on livelock-free components in all teams, so that we only need to test for livelocks on entirely new components, not components

created by synchronization. This is algorithmically cheaper (see below) and matches our iterative correct-by-construction philosophy.

It should be noted that we use the definition of livelock from [2], not [1]. A significant difference is that [1] requires livelock-freeness only in reachable parts of the transition space, while [2] requires it for the entire transition state. As this thesis is based strongly on [2] we will follow that definition and requirement. However, later on we will discuss pruning away unreachable states and transitions (see section 4.6), and when applied to an RTS with no unreachable states there is no practical difference between those definitions.

Algorithm 1 TryLivelock(list, action)

```

1: current_state ←  $\delta(\text{list.tail}, \text{action})$ 
2: if current_state ∈ list then
3:   return true
4: end if
5: list.append(current_state)
6: for all  $a \in \Sigma^{int}$  such that  $\exists \text{next\_state} : (\text{current\_state}, a, \text{next\_state}) \in \delta \wedge \text{next\_state} \notin \text{safe}$  do
7:   if TryLivelock(list,  $a$ ) then
8:     return true
9:   end if
10: end for
11: safe.add(current_state)
12: list.removeLast
13: return false

```

Algorithm 2 Detect Livelock

```

1: global safe ←  $\emptyset$ 
2: for all  $q \in Q - \text{safe}$  do
3:   for all  $a \in \Sigma^{int}$  with  $\exists q' : (q, a, q') \in \delta$  do
4:     if TryLivelock( $[q]$ ,  $a$ ) then
5:       return true
6:     else
7:       safe.add(q)
8:     end if
9:   end for
10: end for
11: return false

```

Algorithm 2 calls Algorithm 1 on each state of the RTS \mathcal{A} that has outgoing internal actions, provided that state has not been checked yet. Algorithm 1 then performs a depth-first search of the transition space starting in that state of \mathcal{A} using only internal actions. Because (we required that) \mathcal{A} has finite states, either the algorithm finds a repeat state, or terminates after reaching every other state but not finding a cycle. When TryLivelock exhausts its search space without finding a livelock, it returns False to DetectLivelock; and when DetectLivelock has exhausted its search space without finding a livelock it returns False to the caller. If a livelock is found by TryLivelock then it passes a True result upstream which eventually informs the caller that a livelock does exist in this component.

Since Algorithm 2 and Algorithm 1 maintain a global set called “safe” of states already found to be uninvolved in any livelocks, no state needs to be examined more than once. Unless a livelock is detected, every state will in fact be visited once. Again in the worst case, in every state except one, all internal actions are enabled, and lead towards other states in such a way that the internal transition space remains acyclic. In such a case, the complexity of this algorithm will be $|Q| \times |\Sigma^{int}|$ as it checks for enabled transitions in every state (TryLivelock line 6 and DetectLivelock line 3).

Since the size of the state space of a team of components $|Q| = |Q_1| \times |Q_2| \times \dots \times |Q_n|$ is roughly exponential, and internal action alphabets do not overlap, the complexity of checking a team together would be:

$$\left(|Q_1| \times |Q_2| \times \dots \times |Q_n|\right) \times \left(|\Sigma_1^{int}| + |\Sigma_2^{int}| + \dots + |\Sigma_n^{int}|\right)$$

The complexity of checking the components separately is only:

$$\left(|Q_1| \times |\Sigma_1^{int}|\right) + \left(|Q_2| \times |\Sigma_2^{int}|\right) + \dots + \left(|Q_n| \times |\Sigma_n^{int}|\right)$$

Clearly this is an argument in favor of preprocessing.

4.4 The straightforward approach: exhaustive search for directly incompatible states

To determine if \mathcal{T} is *compatible*, we can use a straightforward algorithm to find a compatibility relation \mathcal{R} for \mathcal{T} . If we can construct \mathcal{R} , then \mathcal{T} is compatible. If we cannot construct \mathcal{R} , \mathcal{T} must be incompatible.

The definition of *compatibility* requires that there exists a compatibility relation; multiple relations are possible but at the very least the initial state must be in the relation, and any state directly reachable from a state in the relation must also be in the relation. Therefore, the smallest compatibility relation contains exactly the set of *reachable* states.

If we find a reachable state q that violates compatibility, through message loss or deadlock, then that state is not in the relation. We call this *directly incompatible* (see Definition 16). But since steps are not allowed to lead outside the relation, all states that can eventually reach q must also not be in the relation. And since q was reachable from the initial state, the initial state is also not in the relation. Therefore, no compatibility relation exists for such a system.

Definition 16 (*directly incompatible state*).

State $q \in Q$ is a **directly incompatible** state if one of the following holds:

1. $\exists a \in \Sigma^{com}$ is output-domain enabled in q , but not input-domain enabled in q .
2. $\exists a \in \Sigma^{com}$ is input-domain enabled in q , but there is no action that is enabled in q .

We will call the set of directly incompatible states **DIS**.

This then leads to the straightforward algorithm for determining the compatibility of \mathcal{S} : start in the initial state and find all states immediately reachable from there; put them in a queue. Visit each of them, finding more reachable states. In each of those states, test for direct incompatibility. If that is found, \mathcal{S} is incompatible. If no direct incompatibility is found and all reachable states have been visited, the set of visited states \mathcal{R} is now a minimal compatibility relation for \mathcal{S} .

Algorithm 3 will decide compatibility for finite, deterministic composable systems.³ If it finds a contradiction it will terminate early⁴, otherwise it has to run until all reachable states have been checked. However, is it a practical solution? Let us examine the complexity.

A reasonable measure of complexity is the number of times the enabledness of an action in a state is checked (lines 12–13 and 28). That is both one of the most frequent and costly computations of this algorithm. This is done for every action, in every reachable state of Q , unless the algorithm terminates early. Good cases include few reachable states and quickly-found contradictions. Bad cases include high amounts of reachable states and ironically, compatible systems that require the full exhaustive search to confirm.

Since $|Q| = |Q_1| \times |Q_2| \times \dots \times |Q_n|$ and $|\Sigma| = |\Sigma_1 \cup \Sigma_2 \cup \dots \cup \Sigma_n|$, we get a worst-case complexity of $O(|Q_i|^n)$, where $|Q_i|$ is average. Such exponential complexity may pose an obstacle to practical use of this algorithm. In the following sections, we will look at other approaches to try to handle this.

³The algorithm could easily be adjusted to also do so for nondeterministic systems as discussed before.

⁴This is not a technical necessity; the algorithm could instead just output a reference to the offending state, and continue searching, so that all offending states are found in one run of the algorithm.

Algorithm 3 Exhaustive State-Space Search

```
1: for all components  $\mathcal{A}_1, \dots, \mathcal{A}_n$  do
2:   Use Algorithm 2 to verify that the component is livelock-free; if not, stop.
3: end for

4:  $I \leftarrow [I_1 \dots I_n]$ 
5:  $Todo \leftarrow [I]$ 
6:  $Done \leftarrow \emptyset$ 

7: while  $Todo \neq []$  do
8:    $current\_state \leftarrow Todo.head$ 
9:    $some\_progress \leftarrow \mathbf{false}$ 
10:   $awaiting\_input \leftarrow \mathbf{false}$ 

11:  for all  $a \in \Sigma^{com}$  do
12:    if  $a$  is output domain enabled then
13:      if  $a$  is input domain enabled then
14:         $next\_state \leftarrow \delta(current\_state, a)$ 
15:        if  $next\_state \notin Todo \wedge next\_state \notin Done$  then
16:           $Todo.append(next\_state)$ 
17:        end if
18:         $some\_progress \leftarrow \mathbf{true}$ 
19:      else
20:        CONTRADICTION: MESSAGE LOSS
21:      end if
22:    else
23:      if  $a$  is input domain enabled then
24:         $awaiting\_input \leftarrow \mathbf{true}$ 
25:      end if
26:    end if
27:  end for

28:  for all  $a \notin \Sigma^{com}$  with  $\wedge \forall i : a \in \Sigma_i \implies a \in EN_{\mathcal{A}_i} current\_state[i]$  do
29:     $next\_state \leftarrow \delta(current\_state, a)$ 
30:    if  $next\_state \notin Todo \wedge next\_state \notin Done$  then
31:       $Todo.append(next\_state)$ 
32:    end if
33:     $some\_progress \leftarrow \mathbf{true}$ 
34:  end for

35:  if  $some\_progress = \mathbf{false} \wedge awaiting\_input = \mathbf{true}$  then
36:    CONTRADICTION: DEADLOCK
37:  end if
38:   $Done.add(Todo.head)$ 
39:   $Todo.removefirst$ 
40: end while
```

4.5 The backwards approach, starting in *DIS*

Conceivably, only a small portion of the total state space of the system has compatibility problems. If this section is not reachable, then the system is fine. Since the previous approach experiences high complexity in systems that are compatible or close to it (having few states with direct incompatibility), it might be worthwhile to focus on the bad apples. If we can identify the states where a communication failure takes place, we can then try work backwards, to see if there is a path going to those states from the initial state.

We now define ultimately incompatible states as those that are directly incompatible, as well as those that can through a series of transitions arrive at a directly incompatible state.

Definition 17 (*ultimately incompatible state*).

A state q is **ultimately incompatible** if one of the following holds:

1. $q \in DIS$
2. $\exists a, \exists q' : (q, a, q') \in \delta \wedge q' \in DIS$
3. $\exists a, \exists q' : (q, a, q') \in \delta \wedge q'$ is ultimately incompatible.

We will call the set of ultimately incompatible states **UDIS**.

Theorem 18 (*states in the compatibility relation are not in UDIS*).

If \mathcal{R} is a *compatibility relation*, then $q \in \mathcal{R} \implies q \notin UDIS$

Proof. Follows directly from Definition 8 which requires that the states of \mathcal{R} do not suffer message loss or deadlock (i.e. not to be directly incompatible), and requires the same from all states reachable from it; versus Definition 17 which requires states in UDIS to have a path to a state that is directly incompatible. \square

Lemma 19 (*I not in UDIS*).

All components of \mathcal{S} are livelock-free and $I \notin UDIS \implies \mathcal{S}$ is *compatible*.

Proof. Follows directly from Definitions 8 and 17. \square

Lemma 19 suggests another straightforward method to decide the compatibility of \mathcal{S} ;

1. Verify that all components of \mathcal{S} are livelock-free.
2. Compute *DIS*.
3. Compute *UDIS*.
4. Verify that $I \notin UDIS$

Step 1 we accomplish with Algorithm 2, and step 4 is also simple. The intermediate steps are more complicated.

Let us consider the complexity of Algorithm 4. It performs two separate loops, computing states where receptiveness fails (1–6) and where there the system is deadlocked (7–15). To compute the states where receptiveness fails we must for every communicating output action a compute the states of the output domain where it is enabled, and then the states of the input domain where it is not enabled. To do so we need, for every component i , to go through Q_i to see if a is enabled there, so we are now already looking at a complexity of $|\Sigma^{out} \cap \Sigma^{inp}| \times |Q_1| \times |Q_2| \times \dots \times |Q_n|$.

Next we have the complexity of looking for deadlocks. Now we do this $|\Sigma^{inp}|$ times, and each such time first compute all enabled states of the input domain, and then for all possible states of the output domain, test *all* actions for enabledness. This gets us a complexity of $|\Sigma^{inp}| \times |Q_1| \times |Q_2| \times \dots \times |Q_n| \times |\Sigma|$.

Algorithm 4 Compute DIS

Let $\mathcal{S}_{a!}$ be the output domain for a and let $\mathcal{S}_{a?}$ be the input domain for a .

```
1:  $Q^{loss} \leftarrow \emptyset$ 
2: for all  $a \in \Sigma^{out} \cap \Sigma^{com}$  do
3:    $Q_{a!} \leftarrow \{q \mid q \in \mathcal{S}_{a!} \wedge a \text{ EN}_{\mathcal{S}_{a!}} q\}$  // Add  $q$  if the output domain is enabled there.
4:    $Q_{a?} \leftarrow \{q \mid q \in \mathcal{S}_{a?} \wedge \neg(a \text{ EN}_{\mathcal{S}_{a?}} q)\}$  // Add  $q$  if the input domain is not enabled there.
5:    $Q^{loss}.add(Q_{a!} \times Q_{a?})$ 
6: end for

7:  $Q^{lock} \leftarrow \emptyset$ 
8: for all  $a \in \Sigma^{inp} \cap \Sigma^{com}$  do
9:    $Q_{a?} \leftarrow \{q \mid q \in \mathcal{S}_{a?} \wedge a \text{ EN}_{\mathcal{S}_{a?}} q\}$  // Add  $q$  if the input domain is enabled. there.
10:  for all  $p \in Q_{a?} \times Q_{\mathcal{S}_{a!}}$  do
11:    if  $\neg(\exists b, \exists p' : (p, b, p') \in \delta)$  then
12:       $Q^{lock}.add(p)$  // Input expected, but no actions are enabled.
13:    end if
14:  end for
15: end for

16:  $DIS \leftarrow Q^{loss} \cup Q^{lock}$ 
```

Algorithm 5 Compute UDIS

```
1:  $UDIS \leftarrow \emptyset$ 
2:  $todo \leftarrow [q \mid q \in DIS]$ 
3: while  $todo \neq []$  do
4:    $current\_state \leftarrow todo.head$ 
5:   for all  $\{p \mid p \in Q \wedge p \notin UDIS \wedge p \notin todo \wedge \exists a : (p, a, current\_state) \in \delta\}$  do
6:      $todo.append(p)$ 
7:   end for
8:    $UDIS.add(current\_state)$ 
9:    $todo.removefirst$ 
10: end while
11: return  $UDIS$ 
```

The worst-case complexity of Algorithm 5 is disappointing. In the worst case we have to search, for every state found to be in UDIS, every state not yet in UDIS (or *todo*) to see if there is an action transitioning to the current state. Suppose we have a graph where the states essentially form a line from start to end, and the very last state is directly incompatible. Furthermore, every state is connected to the next one with every action (except for the final state; all we need to know about its outgoing transitions is that they result in direct incompatibility). Then for every state, beginning with the last one and working back to the beginning, we need to look up all the state/action combinations of all earlier states. This results in a complexity of:

$$\sum_{i=1}^{|Q|-1} (|Q| - i) \times |\Sigma|$$

This complexity is ostensibly not so bad, but keep in mind that $|Q| = |Q_1| \times |Q_2| \times \dots \times |Q_n|$. So even without the added complexity of computing DIS, the worst-case complexity of the backwards approach is no better than that of the exhaustive forwards approach given earlier.

Given the difficulty of computing DIS and the worst-case complexity of computing UDIS, this approach is not feasible.

4.6 Pruning the search space

As the previous sections have shown, the size of the state space tends to explode when automata are teamed up. We will now investigate if it is possible to limit this explosion. We will examine what parts of a team automaton can be pruned without changing its actual functioning. To be admissible, a pruning method must meet these criteria:

1. The computations of the pruned automaton are the same as those of the original. (It has the same behaviour.)
2. The compatibility of the resulting automaton has to be the same as that of the original, regardless of what automata it is teamed up with.

We do not optimize towards specific teamings with other components; we want the same results from a teaming with original or pruned automaton in every situation.

With these criteria, our pruned automaton remains a correct substitute in all future teams.

We have already discussed unreachable states of individual components. Likewise, a transition can be *useless* if it can never be used, and an action can be *useless* if there are no useful transitions using that action.

Definition 20 (*usefulness*).

A transition $t \in \delta$ is **useful** if $t = (p, a, q)$ for some reachable state p . A transition that is not useful is **useless**.

An action a is useful if there is at least one useful $t \in \delta$ such that $t = (p, a, q)$. An action that is not useful is useless.

It would seem like we could prune away all these unreachable and useless parts of an automaton to make it smaller and faster to process. But we have to be sure that doing so does not change the way it functions, both in isolation, and in the context of a team.

Definition 21 (*behaviour*).

The **behaviour** $\Gamma(\mathcal{A})$ of a *TS* \mathcal{A} is the set of all *computations* of \mathcal{A} .

We now introduce *active computations*. These can be understood as the parts of a team's computations that a given component of that team actively plays a role in.

Definition 22 (*active computation*).

If $\gamma = q^0 a^1 q^1 a^2 q^2 \dots a^n q^n$ is a computation of \mathcal{T} , and \mathcal{A}_i is the i -th component of \mathcal{T} , then γ_i^n is the **active computation** of \mathcal{A}_i corresponding to γ , defined thus:

1. $\gamma_i^0 = \text{proj}_i(q^0)$
2. assume $\gamma_i^j = \omega p^j$ with $\omega \in (Q_i \Sigma_i)^*$, $p^j \in Q_i$, and $0 \leq j < n$
then $\gamma_i^{j+1} = \begin{cases} \gamma_i^j a^{j+1} \text{proj}_i(q_{j+1}) & \text{if } (p^j, a^{j+1}, \text{proj}_i(q_{j+1})) \in \delta_i \\ \gamma_i^j & \text{otherwise} \end{cases}$

Definition 23 (*active behaviour*).

The **active behaviour** $\Gamma_i^T(\mathcal{T})$ of a component \mathcal{A}_i of \mathcal{T} is the set of the active computations of \mathcal{A}_i corresponding to the computations of \mathcal{T} .

So, how does the active behaviour of a component compare to its behaviour?

Theorem 24 (*active behaviour of a component in a team is a subset of its solitary behaviour*).

$$\Gamma_i^T(\mathcal{T}) \subseteq \Gamma(\mathcal{A}_i)$$

Proof. All the active computations of \mathcal{A}_i as defined by Definition 22 are also normal computations of \mathcal{A}_i as defined by Definition 4. Therefore all active computations of \mathcal{A}_i are also present in $\Gamma(\mathcal{A}_i)$. \square

This does not immediately appear to be a very interesting result, but in the case of *synchronous product*, we can expand it a little, and then it does become very relevant to pruning;

Theorem 25 (*component's active behaviour in a team is a subset of active behaviour in subteams containing it*).

Let $1 \leq i \leq j < k \leq n$. Then the following holds:

$$\Gamma_i^T(\mathcal{T}^k) \subseteq \Gamma_i^T(\mathcal{T}^j)$$

Proof. Suppose $\gamma_i \in \Gamma_i^T(\mathcal{T}^i)$, but $\gamma_i \notin \Gamma_i^T(\mathcal{T}^j)$. Then γ_i must contain a sequence $(q, a, r) \in \delta_i$ where $\text{proj}_i(q)$ is not a *reachable* state of \mathcal{A}_i in any computation of \mathcal{T}^j . However, q was a reachable state of \mathcal{A}_i in isolation, because it occurs in γ_i . Thus q must have become unreachable as a result of applying Definition 7 (*synchronous product*). This necessarily implies that some other component \mathcal{A}_x of \mathcal{T}^j is not enabled for action a' whenever \mathcal{A}_i is enabled for an action a' with $(p, a', q) \in \delta_i$. But if $\mathcal{A}_x \in \mathcal{T}^k$ then $\mathcal{A}_x \in \mathcal{T}^k$ and therefore $\gamma_i \notin \Gamma_i^T(\mathcal{T}^k)$. This proves that $\Gamma_i^T(\mathcal{T}^k) - \Gamma_i^T(\mathcal{T}^j) = \emptyset$.

The reverse does not necessarily hold however; it is possible that $\gamma_i \in \Gamma_i^T(\mathcal{A}^j)$ while $\gamma_i \notin \Gamma_i^T(\mathcal{A}^k)$. All that is required is that \mathcal{A}_y with $j < y \leq k$ causes q to become unreachable in \mathcal{A}_i as shown above. Therefore it is possible that $\Gamma_i^T(\mathcal{A}^j) \neq \Gamma_i^T(\mathcal{A}^k)$ \square

Theorem 25 has some interesting implications for iterative construction of team automata. Once part of the behaviour of a component becomes inaccessible in a subteam, it will remain inaccessible in all teams that include that subteam.

Theorem 26 (*propagation of uselessness*).

Let $1 \leq i \leq j < k \leq n$. Then the following holds:

1. Let $p \in Q_i$.
If $\nexists q : q \in Q_{T^j}$ so that $\text{proj}_i(q) = p$ and q is reachable in \mathcal{T}^j ,
then $\nexists r : r \in Q_{T^k}$ so that $\text{proj}_i(r) = p$ and r is reachable.
2. Let $t = (p, a, p') \in \delta_i$.
If $\nexists t' : t' = (q, a, q') \in \delta_{T^j}$ so that $(\text{proj}_i(q), a, \text{proj}_i(q')) = t$ and t' is useful,
then $\nexists t'' : t'' = (r, a, r') \in \delta_{T^k}$ so that $(\text{proj}_i(k), a, \text{proj}_i(k')) = t$ and t'' is useful.
3. If action a is useless in \mathcal{T}^j then it is also useless in \mathcal{T}^k .

Proof.

1. If $\text{proj}_i(q)$ is unreachable in \mathcal{A}^j then it is not part of any computation in $\Gamma_i^T(\mathcal{A}^j)$ By Theorem 25, it is therefore also not part of $\Gamma_i^T(\mathcal{A}^k)$, so it is also unreachable in \mathcal{A}^k .
2. If t is useless in \mathcal{T}^j then by Definition 20 its current state is not reachable in \mathcal{T}^j . Then according to (1) its current state is also unreachable in \mathcal{T}^k , and so t is also useless in \mathcal{T}^k .
3. If a is useless in \mathcal{T}^j then a exists in Σ_i of some component \mathcal{A}_i of \mathcal{T}^j , but \mathcal{A}_i has no useful transition using a as its action. According to (2) no useless transition of \mathcal{A}_i can become useful in \mathcal{T}^k . Then under Definition 7 no synchronization on a can take place in \mathcal{T}^k , because $a \in \Sigma_i$, and therefore it must be part of any synchronization on a , but no useful transition from \mathcal{A}_i can be found.

□

4.6.1 Pruning unreachable states from the statespace

As we have shown in Theorem 26, once a state becomes unreachable it will never become reachable again. We can therefore safely remove it from the component without changing either the behaviour of the component or the behaviour of any teams using that component. Note that Algorithm 3 implicitly does this, because its *Done*-set of states contains only reachable states.

4.6.2 Pruning useless transitions

Theorem 26 likewise informs us that it is safe to remove useless transitions. Since these transitions do not show up in behaviour, they can be pruned from a component without changing its behaviour.

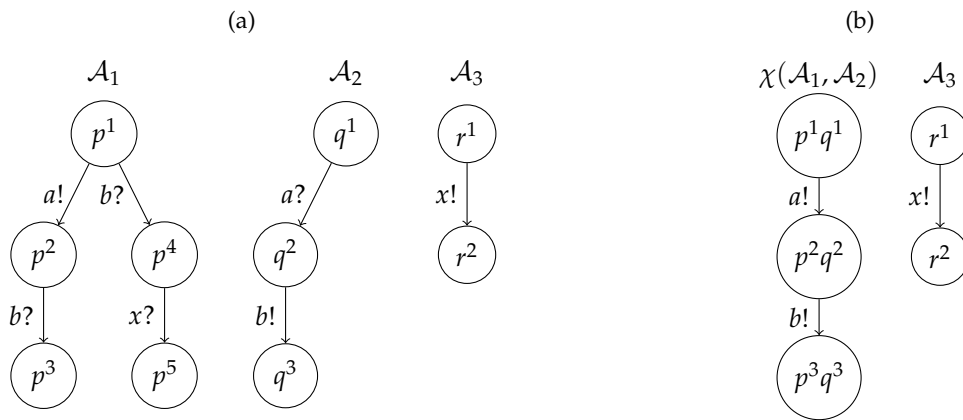
Because useless transitions are defined as originating in unreachable states, pruning unreachable states necessarily implies the pruning of useless transitions, because those transitions' originating state disappears; if $t = (q, a, q')$ and we prune q then we must also prune t .

4.6.3 Pruning useless actions

If an internal action is or becomes useless (because no transitions using it remain) then that action can be pruned as well, because its absence will not make any difference for the components behaviour or compatibility, in isolation or in a team.

It is *not* always safe to prune away useless external actions. Consider for example Figure 5; in $\mathcal{X}(\{\mathcal{A}_1, \mathcal{A}_2\})$ action $x?$ is useless. If we keep it, $\{\mathcal{X}(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ is not compatible, because receptiveness fails for $x!$. If we had discarded $x?$ then $\{\mathcal{X}(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ would have been compatible.

Figure 5: Useless actions still influence compatibility



The cause for this is the definition of *synchronous product*: if we want to execute a , all components with a in their alphabet need to participate in the execution. Because external actions that are currently noncommunicating might become so in the future, pruning them is not safe.

4.6.4 Reducing components

We can now formulate a way to reduce a component to its minimum size, without changing anything about its functioning, either in isolation or in a team. The result is a component which has only reachable states and only useful transitions. It can still have useless external actions, for reasons explained in section 4.6.3.

Definition 27. The **Reduced Reactive Transition System** (or **RRTS**) corresponding to the *RTS* \mathcal{A} is another *RTS* $\mathcal{A}' = \langle Q', \Sigma', \delta', I' \rangle$:

- $Q' = \{q \in Q \mid q$
- $\Sigma' = \Sigma^{int'} \cup \Sigma^{out'} \cup \Sigma^{inp'}$
- $\Sigma^{int'} = \{a \in \Sigma^{int} \mid a \text{ is useful}\}$
- $\Sigma^{out'} = \Sigma^{out}$
- $\Sigma^{inp'} = \Sigma^{inp}$
- $\delta' = \{t \in \delta \mid t \text{ is useful}\}$
- $I' = I$

Clearly, there is only one RRTS corresponding to a given RTS. However, many RTSs might be reduced to the same RRTS. We can use Algorithm 6 to obtain an RRTS. The complexity of Algorithm 6 is based on the number of times the inner loop on lines 8–20 is executed. This is at worst $|Q| \times |\Sigma|$ because every action might be enabled in every state.

When we wish to explicitly denote that an RTS is an RRTS, we will indicate this with ρ . Likewise, we can indicate a variant of synchronous product that while creating a team using synchronous product, also reduces the resulting team, and we will write that as $\chi\rho()$.

Algorithm 6 ReduceRTS($\langle Q, (\Sigma^{int}, \Sigma^{out}, \Sigma^{inp}), I, \delta \rangle$)

```

1: Todo  $\leftarrow$  [initial_state]
2:  $Q' \leftarrow \emptyset$ 
3:  $\Sigma^{out'} \leftarrow \Sigma^{out}$ 
4:  $\Sigma^{inp'} \leftarrow \Sigma^{inp}$ 
5:  $\delta' \leftarrow \emptyset$ 
6: while Todo  $\neq$  [] do
7:    $q \leftarrow$  Todo.head

8:   for all  $a \in \Sigma$  do
9:     if  $\exists q' \in Q : (q, a, q') \in \delta$  then
10:       $\delta' \leftarrow (q, a, q')$ 
11:      if  $q' \notin Q' \wedge q' \notin$  Todo then
12:        Todo.append( $q'$ )
13:      end if
14:      if  $a \in \Sigma^{int}$  then
15:         $\Sigma^{int'} \leftarrow a$ 
16:      end if
17:    end if
18:  end for

19:   $Q'.add$ (Todo.head)
20:  Todo.removefirst
21: end while
22: return  $\langle Q', (\Sigma^{int'}, \Sigma^{out'}, \Sigma^{inp'}), I', \delta' \rangle$ 

```

4.7 Constructing bottom-up from “correct” components

The results of [2] referenced in section 2.2.6 demonstrate that (if certain conditions hold), it is possible to assemble a team in steps rather than all at once, while preserving its compatibility properties. Can we

use such an approach to cut down on the complexity of using Algorithm 3? If we have a set of automata that we want to form into a team, can we do so in steps rather than all at once, to reduce the workload?

For this to be even possible, we need to study the results of [2] and evaluate how practical their usage is. To begin with, Lemma 10 defines the relation between the compatibility of subteams and the final team only in a pairwise fashion; it provides no guidance on how to assemble a final team from three or more subteams, only from two subteams. This is because of the requirement that the second component *collaborates* with the first. This collaboration requirement is problematic for us, because it is not easy to test without actually building up the state space. So, to know if two automata collaborate and are therefore worth testing for compatibility (by computing their synchronous product) we have to compute their synchronous product.

The following lemma, theorem and corollary explore ways to circumvent this collaboration requirement. They use the fact that *action-complementary* components automatically collaborate, and that if a pair of automata form a complete system, they must be action-complementary. These results are not satisfactory however. It is true that all action-complementary pairs of components collaborate, because they have no shared actions that can fail the test. But relying on action-complementarity means we can no longer handle collaborating pairs that do share actions. To be able to use the much simpler test of action-complementarity we sacrifice coverage of the whole problem. Completeness is an even narrower set of automata, so we lose even more coverage.

But even the collaborating approach is not really good enough. We would also like to cover systems that cannot be subdivided into collaborating pairs. Can we do that? We need to take a look at where the requirement for collaboration in Lemma 10 comes from. The proof of the Lemma is quite long, but uses collaboration only near the end, to handle the last case.

We will now propose a variation on Lemma 10 that dispenses with the collaboration requirement. This produces a more generic rule that we can use to write a recursive algorithm for determining the compatibility of \mathcal{S} .

Lemma 28 (*correspondence of transitions*).

Let $1 < k < n$
 Let $\mathcal{T} = \chi(\{\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\}), \mathcal{A}_{k+1}, \dots, \mathcal{A}_n\})$
 Let $\mathcal{T}' = \chi(\{\mathcal{A}_1, \dots, \mathcal{A}_n\})$
 Then:

$$(((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)) \in \delta_{\mathcal{T}} \iff ((p_1, \dots, p_n), a, (p'_1, \dots, p'_n)) \in \delta_{\mathcal{T}'}$$

Proof. Observe that if $t = (((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n))$ is a transition of \mathcal{T} then for all $i = 1, \dots, n$, (p_i, a, p'_i) is a transition of \mathcal{A}_i or a is not an action of \mathcal{A}_i .

Likewise, if $t' = ((p_1, \dots, p_n), a, (p'_1, \dots, p'_n))$ is a transition of \mathcal{T}' , then for $i = 1, \dots, n$, (p_i, a, p'_i) is a transition of \mathcal{A}_i or a is not an action of \mathcal{A}_i .

Recall that by Definition 7 every component \mathcal{A}_i of a team made with synchronous product that has a in its alphabet participates in every transition using a or else a cannot be executed at that state.

If \mathcal{T} can execute t , then all participating components are enabled for a in that state of \mathcal{T} . As a consequence, all those components are also able to participate in the state of \mathcal{T}' that has the same local states for the components.

If \mathcal{T} executes t , then \mathcal{T} has $a \in \Sigma_i, \mathcal{A}_i$ participates. If \mathcal{T}' executes a as well, then \mathcal{A}_i would participate as well. Likewise if \mathcal{T} shadows \mathcal{T}' . \square

Lemma 29 (*correspondence of states*).

Let $1 < k < n$
 Let $\mathcal{T} = \chi\rho(\{\chi\rho(\{\mathcal{A}_1, \dots, \mathcal{A}_k\}), \mathcal{A}_{k+1}, \dots, \mathcal{A}_n\})$
 Let $\mathcal{T}' = \chi\rho(\{\mathcal{A}_1, \dots, \mathcal{A}_n\})$
 Then $(p_1, \dots, p_k), p_{k+1}, \dots, p_n \in Q_{\mathcal{T}} \iff (p_1, \dots, p_n) \in Q_{\mathcal{T}'}$

Proof. Follows directly from Lemma 28 and Definition 27 as follows:

Observe that $(I_1 \times \dots \times I_k) \times I_{k+1} \dots \times I_n \in Q_{\mathcal{T}}$ and $(I_1 \times \dots \times I_n) \in Q_{\mathcal{T}'}$

Assume $(p_1, \dots, p_k), p_{k+1}, \dots, p_n \in Q_{\mathcal{T}}$. Because \mathcal{T} and \mathcal{T}' are reduced, $(p_1, \dots, p_k), p_{k+1}, \dots, p_n$ is reachable. If we apply Lemma 28 we know that (p_1, \dots, p_n) is therefore also reachable in \mathcal{T}'

Then if $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)) \in \delta_{\mathcal{T}}$, therefore $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in Q_{\mathcal{T}}$ and according to Lemma 28 $((p_1, \dots, p_n), a, (p'_1, \dots, p'_n)) \in \delta_{\mathcal{T}'}$, and therefore $(p'_1, \dots, p'_n) \in Q_{\mathcal{T}'}$.

The reverse implication works in the same way. \square

For a set of components $\{\mathcal{A}_2, \dots, \mathcal{A}_n\}$, we say that it **collaborates on output** with \mathcal{A}_1 if for every state of the Cartesian product, $(p_1, p_2, \dots, p_n) \in Q_1 \times Q_2 \times \dots \times Q_n$, and every $a \in (\Sigma_1^{out} \cap (\Sigma_2^{out} \cup \dots \cup \Sigma_n^{out}))$, a is enabled at (p_2, \dots, p_n) whenever a is enabled at p_1 .

Lemma 30 (correspondence of compatibility).

Let $1 < k < n$.

Let $\mathcal{S} = \{\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\}), \mathcal{A}_{k+1}, \dots, \mathcal{A}_n\}$ be compatible;

Let $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ is *receptive* and *deadlock-free*;

If $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$ collaborates with every output action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$, then $\mathcal{S}' = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ is *compatible*.

Proof. Since \mathcal{S} is compatible, by Definition 8 each of $\chi(\{\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\}), \mathcal{A}_{k+1}, \dots, \mathcal{A}_n\})$ is livelock-free. And by Lemma 15 that means that $\mathcal{A}_1, \dots, \mathcal{A}_n$ are all livelock-free. And again by Lemma 15 that means $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_n\})$ is livelock-free.

Let \mathcal{R} be a compatibility relation for \mathcal{S} . We will prove that

$$\mathcal{R}' = \{(p_1, \dots, p_n) \mid ((p_1, \dots, p_k), p_{k+1}, \dots, p_n) \in \mathcal{R} \text{ and } (p_1, \dots, p_k) \text{ reachable in } \chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})\}$$

is a compatibility relation for \mathcal{S}' .

Clearly, \mathcal{R}' includes $I_1 \times \dots \times I_n$.

Let $p = (p_1, \dots, p_n) \in \mathcal{R}'$ with $((p_1, \dots, p_k), p_{k+1}, \dots, p_n) \in \mathcal{R}$ and (p_1, \dots, p_k) reachable in $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$.

Non-communicating progress: Let a be a non-communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an action of \mathcal{A}_i for $i = 1, \dots, n$. Let $(p, a, p') \in \chi_a^{\mathcal{S}'}$ for some $p' = (p'_1, \dots, p'_n)$.

We have to prove that $p' \in \mathcal{R}'$ and (p'_1, \dots, p'_k) reachable in $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$.

Observe first that a is also a non-communicating action of \mathcal{S} . Moreover, in case a is an input (output) action of one or more of $\mathcal{A}_1, \dots, \mathcal{A}_n$, it is enabled at (p_1, \dots, p_n) in $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_n\})$. Consequently, a is enabled at $((p_1, \dots, p_k), p_{k+1}, \dots, p_n)$ and $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)) \in \chi_a^{\mathcal{S}}$ independent of whether it is an input or output or internal action of \mathcal{S}' . Then by the non-communicating progress property of \mathcal{R} it follows that $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \mathcal{R}$ and hence $p' \in \mathcal{R}'$.

Receptiveness: Let a be a communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an output action of \mathcal{A}_i for $i = 1, \dots, n$.

We have to prove that a is also enabled for those \mathcal{A}_i for which it is an input action and, moreover, if $(p, a, p') \in \chi_a^{\mathcal{S}'}$ for some $p' = (p'_1, \dots, p'_n)$, then $p' \in \mathcal{R}'$.

If a is not a communicating action of \mathcal{S} , then it must be the case that a is a communicating action of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ and a is not an input action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. The receptiveness of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ implies that a is enabled at (p_1, \dots, p_k) . Thus, a is enabled at the \mathcal{A}_i for which it is an input action. Further, $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)) \in \chi_a^{\mathcal{S}}$ and the non-communication progress property for \mathcal{R} yields $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \mathcal{R}$. Hence $p' \in \mathcal{R}'$.

If a is a communicating action of \mathcal{S} , then we have cases:

1. a is a communicating action of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ and an input action of one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$; in this case we use again the receptiveness of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ together with the receptiveness of \mathcal{R} to conclude that a is enabled at those \mathcal{A}_i for which it is an input action and $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), a, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)) \in \chi_a^{\mathcal{S}}$ with $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \mathcal{R}$. Hence in this case $p' \in \mathcal{R}'$.

2. a is an output action of one or more of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ but not an input action of any of them, and an input (and possibly output) action of one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. In this case the receptiveness of \mathcal{R} suffices to finish the proof.
3. a is an input (but not output) action of one or more of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$, and an output (and possibly input) action of one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. In this case the receptiveness of \mathcal{R} suffices to finish the proof.
4. a is not an action of any of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$, but is an input action in one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$ and also an output action in one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. Again the receptiveness of \mathcal{R} suffices to finish the proof.

Deadlock-freeness: Let a be a communicating action of \mathcal{S}' and assume that a is enabled at p_i whenever a is an input action of \mathcal{A}_i for $i = 1, \dots, n$.

We have to prove that there exists an action b in \mathcal{S} such that $(p, b, p') \in \chi^{\mathcal{S}'}$ for some $p' = (p'_1, \dots, p'_n)$.

1. If a is an input action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ or an input action of one or more of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$, then the deadlock-freeness of \mathcal{S} implies that there exist b and $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)$ such that $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), b, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \chi^{\mathcal{S}}$ and we are done.
2. The remaining case is that a is communicating in $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ and not an input action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. Now we use the deadlock-freeness of $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ to infer that there exists an action c and (p'_1, \dots, p'_k) such that $((p_1, \dots, p_k), c, (p'_1, \dots, p'_k)) \in \chi^{\{\mathcal{A}_1, \dots, \mathcal{A}_k\}}$. Again we have different cases to distinguish:

If c is not an action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$, we are done: $(p, c, (p'_1, \dots, p'_n)) \in \chi^{\mathcal{S}'}$.

The remaining cases are that:

- (a) c is an input action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ and an output (and possibly input) action of some or all of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$.
As before, the deadlock-freeness of \mathcal{S} implies that there exist $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)$ such that $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), c, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \chi^{\mathcal{S}}$ and we are done.
- (b) c is an output action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ and an input action of some or all of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$, but not an output action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$.
As before, the receptiveness of \mathcal{R} implies that there exist b and $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)$ such that $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), b, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \chi^{\mathcal{S}}$ and we are done.
- (c) c is an input action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ and an input action of some or all of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$, and not an output action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$. Again, the deadlock-freeness of \mathcal{S} implies that there exist $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)$ such that $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), c, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \chi^{\mathcal{S}}$ and we are done.
- (d) c is an output action of $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ and an output action of some or all of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$ and possibly an input action of any of $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$.
Because of the fact that $\mathcal{A}_{k+1}, \dots, \mathcal{A}_n$ collaborates with $\chi(\{\mathcal{A}_1, \dots, \mathcal{A}_k\})$ for output actions, there exists $((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n)$ such that $((p_1, \dots, p_k), p_{k+1}, \dots, p_n), c, ((p'_1, \dots, p'_k), p'_{k+1}, \dots, p'_n) \in \chi^{\mathcal{S}}$ and we are done.

□

Compared to Lemma 10 this lemma shows a few improvements; it narrows the scope of collaboration that is required to only collaboration on output actions, and it generalizes it to systems with a “leader” of length ≥ 1 , with ≥ 1 “followers”. However, we were not able to drop the collaboration completely. Still, this new lemma also enables us to break down systems of $n > 3$ into smaller compatibility testing jobs instead of having to test everything at once.

4.7.1 Using Lemma 30

The implication of Lemma 30 is that we can take some elements out of \mathcal{S} , provided the rest collaborates on output with them; see if the chosen elements are compatible, and if so replace them in \mathcal{S} with their synchronous product. For example, let's take a system $E = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$ which represents five people passing a ball around in a clockwise manner, shown in Figure 6. They each have an input action to catch a ball from their right neighbour, and an output action to throw it to their left neighbour.




Figure 6: Five players passing around a ball: $E = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$

It is immediately obvious that in this system, each of the components enjoys the output collaboration of the other components, because they do not share any output actions. We will now demonstrate how Lemma 30 allows us to decide the compatibility of our example system step by step. First we pick the first two players and test them for compatibility using Algorithm 3. Having found that they are compatible, we then replace \mathcal{A}_1 and \mathcal{A}_2 with $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ so that $E' = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$. Shown in Figure 7:




Figure 7: Five players passing around a ball: $E' = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$

We don't know yet if E' is compatible, but if it is, then by Lemma 30 so is E . Note that in this intermediate step, there is a state where both components have a ball. So we take two components out of E' and repeat the test. From all the possible choices, we happen to prefer $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ and \mathcal{A}_3 . They are compatible and so $E'' = \{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \mathcal{A}_5\}$. Shown in Figure 8:

This system now has a component that has many more states than the previous components. In fact it has $(|Q_1| \times |Q_2|) \times |Q_3| = 8$ states. And we are not done yet. We add \mathcal{A}_4 , to obtain Figure 9:

Even bigger. But then we add in the fifth component (Figure 10):

Now we trace our way back to the original system:

- The final application of our algorithm tells us that $E''' = \{\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \mathcal{A}_5\})\}$ is compatible.

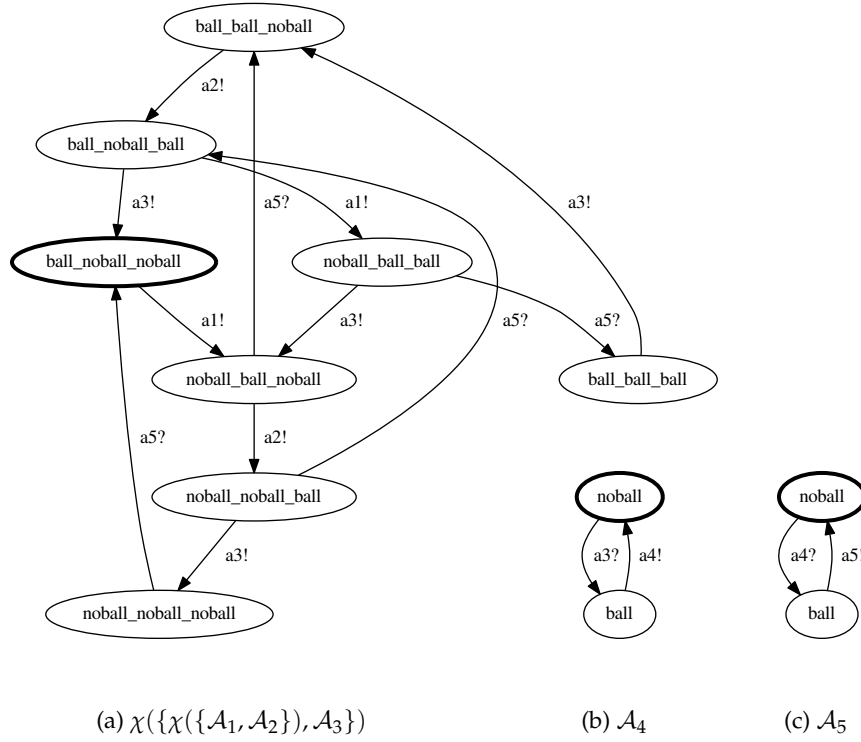




Figure 8: Five players passing around a ball: $E'' = \{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \mathcal{A}_5\}$

- Since we knew that $\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4\}$ is compatible and \mathcal{A}_5 collaborates with $\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4\})$ on output, by Lemma 30, that means that $E'' = \{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \mathcal{A}_5\}$ is compatible.
- Then because we know that $\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \mathcal{A}_5\}$ and $\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}$ are compatible and that \mathcal{A}_4 and \mathcal{A}_5 collaborate on output with $\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\})$, we know that $E' = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$ is compatible.
- And then because we know that $\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$ and $\{\mathcal{A}_1, \mathcal{A}_2\}$ are compatible and that $\mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5$ collaborate on output with $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$, we can conclude that $E = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$ is compatible.

Lemma 30 doesn't force us to assemble the system in this particular order; we could have also tried to combine the \mathcal{A}_3 with \mathcal{A}_4 and then combined $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$ with $\chi(\{\mathcal{A}_3, \mathcal{A}_4\})$, as shown in Figure 11 and Figure 12.

Although Figure 9 and Figure 12 look different, they have the same amount of states and the remaining synchronization in Figure 12 results in a TA identical to Figure 10. But we have one more variation to try: $\hat{E}'' = \{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_5\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\})\}$ (Figure 13)



(a) $\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4\})$

(b) \mathcal{A}_5

Figure 9: Five players passing around a ball: $E''' = \{\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4\}), \mathcal{A}_5\}$

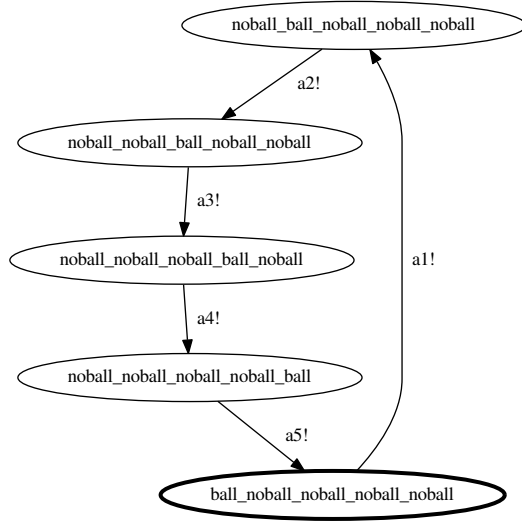
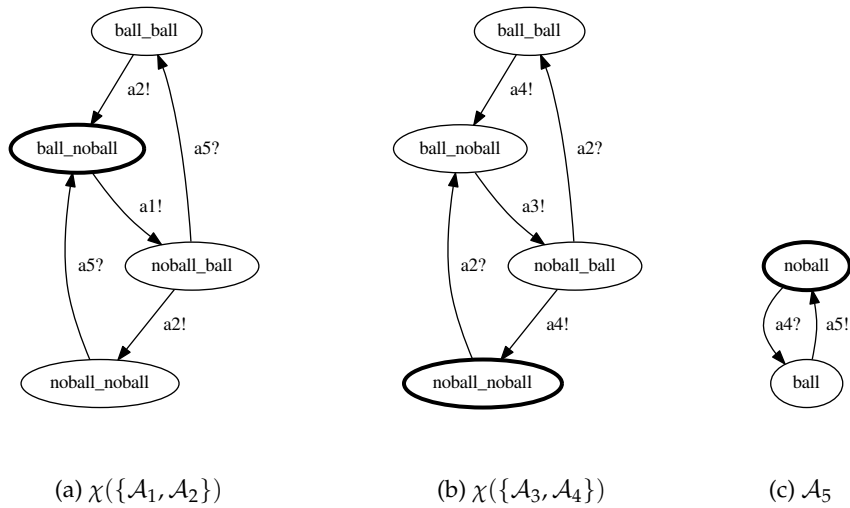


Figure 10: Five players passing around a ball: $\chi(\{\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4, \}, \mathcal{A}_5\})$

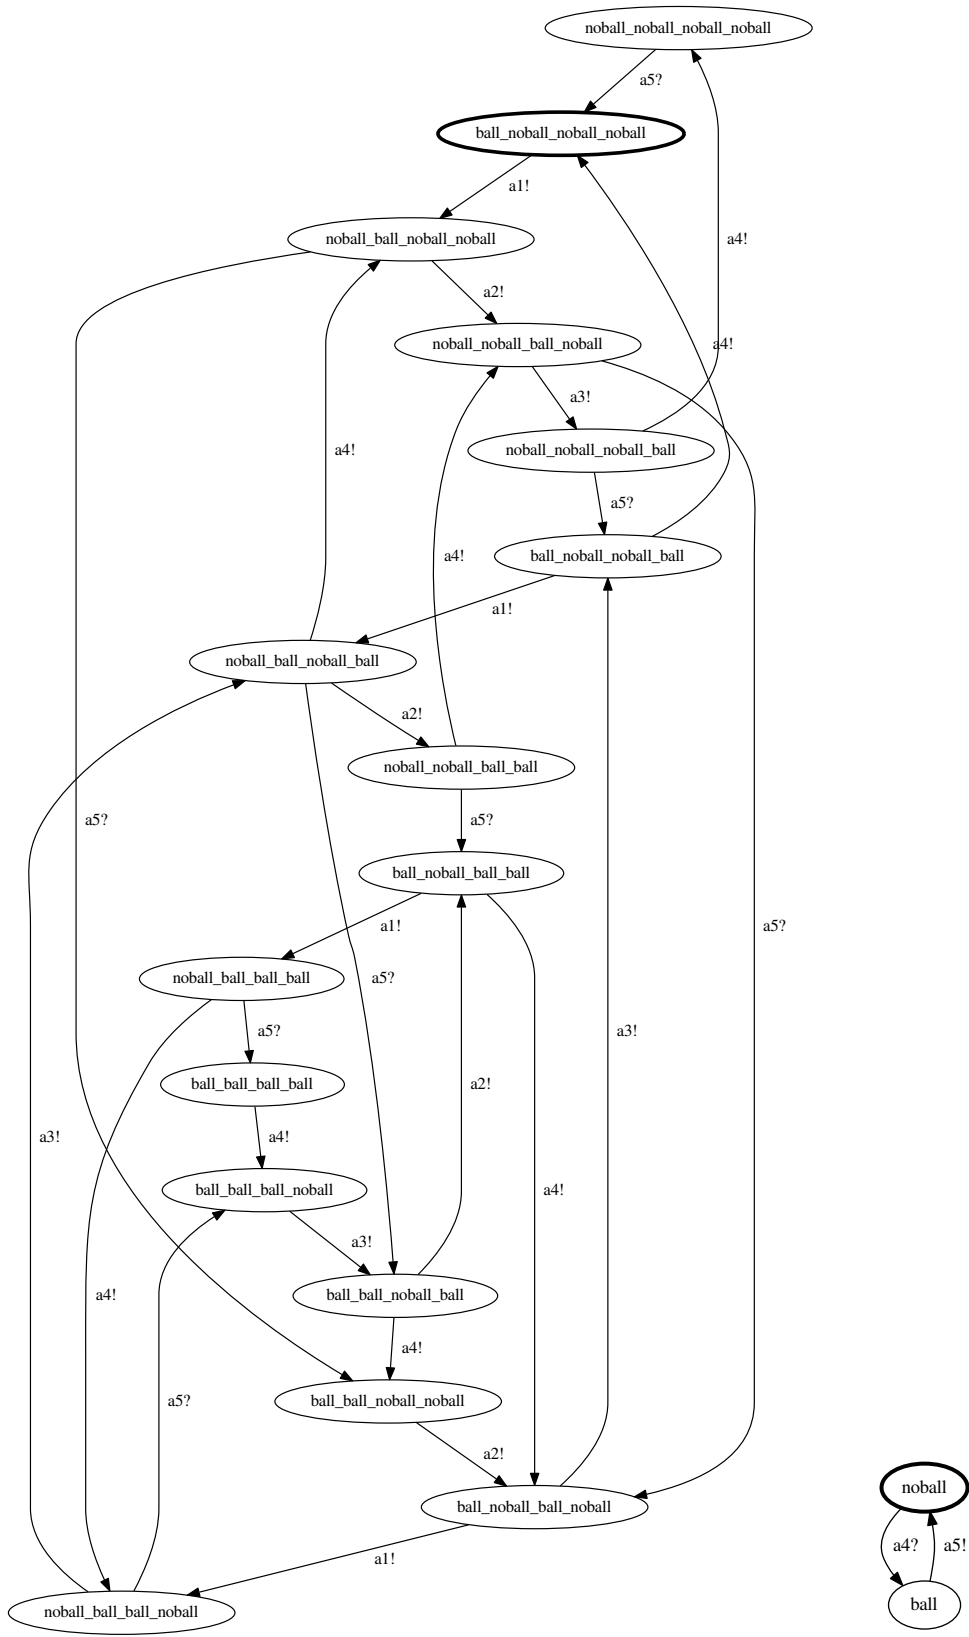


(a) $\chi(\{\mathcal{A}_1, \mathcal{A}_2\})$

(b) $\chi(\{\mathcal{A}_3, \mathcal{A}_4\})$

(c) \mathcal{A}_5

Figure 11: Five players passing around a ball: $\hat{E}'' = \{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\}), \mathcal{A}_5, \}$



(a) $\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\})\})$

(b) \mathcal{A}_3

Figure 12: Five players passing around a ball: $\hat{E}'' = \{\chi(\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \chi(\{\mathcal{A}_4, \mathcal{A}_5\}))\}, \mathcal{A}_3\}$

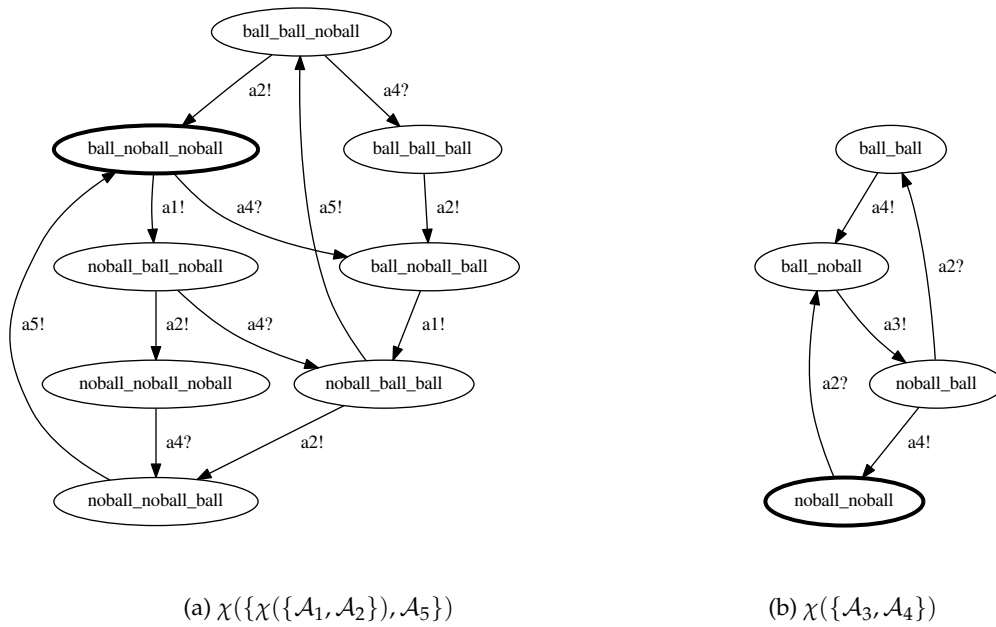


Figure 13: Five players passing around a ball: $\hat{E}'' = \{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_5\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\})\}$

So what is the use of all these different ways of nesting an automaton? Is one method somehow better? We ran these through the Python implementation (see below) which records the effort performed at each synchronization. “Effort” is the amount of tests to see if component \mathcal{A}_i is enabled for action a at state p performed by Algorithm 3, as discussed in section 4.4.

nesting	effort	picture
$\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$	50	Figure 6
$\{\chi(\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_3\}), \mathcal{A}_4\}), \mathcal{A}_5\}$	283	Figure 9
$\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\}), \mathcal{A}_5\}$	163	Figure 12
$\{\chi(\{\chi(\{\mathcal{A}_1, \mathcal{A}_2\}), \mathcal{A}_5\}), \chi(\{\mathcal{A}_3, \mathcal{A}_4\})\}$	107	Figure 13

Clearly, the uncomplicated “all at once” approach did well. But even for such a simple system, the other approaches show a significant difference in performance. In the Experiments section we will experiment with different nestings of sub-products.

5 Software

We have written a Python 2.7 implementation using the results from the earlier sections. We chose Python as a programming language because it offers all the features we need to implement the algorithms in a natural way: object orientation; good data structures for tuples, sets, lists, dictionaries and strings; imperative and functional programming capabilities; good interface with Graphviz for visual rendering of results. Aside from that, Python is widely available and easy to read. We chose object-orientation as a programming paradigm because it is a natural fit for the very object-like team automata.

The software assumes deterministic automata with finite sets of states.

The software starts with the TA class (in `TA.py`), which implements team automata as objects having a name, sets of internal, input and output actions, states, transitions and an initial state as attributes. This class also provides various member functions for manipulating it, such as adding or removing transitions, checking for livelock, generating a corresponding RRTS and so forth. Of particular use is the `screendump()` function which prints a human-readable rendering of the automaton to the screen. Apart from those, the TA class also has two sets that start out empty: `message_losses` and `deadlocks`. These are used by the synchronization algorithm to record any *directly incompatible* states encountered.

The TA module also governs the file I/O of automata to `.out` files. The file format is intended to be compatible with an earlier software project by Casanova (see 5.1), and that software uses the `.out` suffix on output files.

TAs are stored in text files, which look like this:

```
1 .model drinker
2 .inputs serve_coffee
3 .outputs press_button
4 .internal drink_coffee
5 .state graph
6     thirsty press_button waiting
7     waiting serve_coffee happy
8     happy drink_coffee thirsty
9 .marking {thirsty}
10 .end
```

`../source_code/cases/drinker.out`

Model, inputs, outputs, internal, state graph, marking and end are reserved words. The marking is a set containing one state. Only standard alphabet characters and underscore are permitted, and whitespace is a delimiter. It is not exactly user-friendly but necessary to maintain compatibility with the software of Casanova. While interacting with our software, we can abstract most of this to a more pleasant format.

The next module is `synchronize.py` which contains our `synchronous_product_effort(components = [], find_all_contradictions=True)` function. This function computes the synchronous product of the list of components given, as well as the amount of effort required. These are then returned as a `(result, effort)` tuple. If any incompatibilities were encountered, these will be found in `result.message_losses` or `result.deadlocks`. If both sets remain empty then the components were compatible.

The actual working of `synchronous_product_effort` is based on Algorithm 3; starting in the initial state it explores a known-reachable state, checks it for direct incompatibility, determines what transitions are enabled, adds those transitions and their destinations to `result`, and goes on to explore any other states discovered to be reachable but as yet unexplored, until all reachable states have been explored.

`synchronous_product_effort` by default does not stop if it discovers a directly incompatible state. Rather, it records it and goes on to explore the remaining reachable states. In this way, when the algorithm concludes it has computed *DIS* so that a developer can use it to debug. By giving it `find_all_contradictions=False` as an argument this is changed so that it aborts a synchronous product if it finds a contradiction.

Our next layer of programming is `scheduler.py`. This module provides a `scheduler` class which processes schedule files. It first recursively calls schedulers to resolve any sub-schedules (nested synchronous products), before taking on the system of automata to test for compatibility.

The scheduler will process the system by selecting two components and taking their synchronous product to see if they are compatible. If they are, then (justified by Lemma 30) it removes those two components from the system and replaces them with their product. If they are not component, the product is discarded and a different pair is selected from a queue of eligible pairs. When no more eligible pairs remain, all remaining components are sent to `synchronous_product_effort` together to

see if they are compatible. If they are, the system was compatible and the result is a TA corresponding (via Lemma 28 and Lemma 29) to the synchronous product of the original system. Otherwise, the system was not compatible, but the corresponding product is still produced along with its sets of message loss and deadlock states, for debugging purposes.

The schedule file is a text file containing one line describing the system, and looks like this:

$$(a1, a2, (a3, a4))$$

Which corresponds to $\{\mathcal{A}_1, \mathcal{A}_2, \chi\rho(\{\mathcal{A}_3, \mathcal{A}_4\})\}$. When executed by the scheduler, it will first run another scheduler to replace (a3, a4) with their synchronous product, without caring about their compatibility. Then it will proceed to test the compatibility of (a1, a2, Xa3a4).

As mentioned above, the scheduler maintains a queue of eligible pairs for nested synchronization. This queue is formed by, for every possible pair, computing a fitness score. When the time comes to try a new pair, the fittest pair that satisfies the output collaboration requirement of Lemma 30 is used. If it is compatible, all other pairs that overlap with it are dropped and the result of the synchronization is compared with all remaining components to evaluate the fitness of those new possible pairs.

So how is fitness computed? The scheduler uses the `analyze(A, B)` function as an alias for a strategy function. If a strategy argument was given to the scheduler, any calls to `analyze(A, B)` are actually executed by a matching `strategy_[strategyname](A, B)` function instead which records the fitness of (A, B), if appropriate.

Built-in strategies include assigning high fitness to pairs which have relatively many shared actions, FIFO and LIFO, and pairs which could result in large or small $|Q|$ for their synchronous product. This part of the module is intended to be easy to expand with further strategies that come to mind.

When the scheduler class constructor is called, its parameter `self.contradictions` is set. If it is set to True, then the scheduler will use `synchronous_product_effort` with the `find_all_contradictions` parameter set to True; and False if `self.contradictions` is False.

Also, components loaded from file are first reduced. After reduction components are checked for livelocks. Justified by Lemma 15, if `self.contradictions` is set to False, if it detects a livelock the scheduler will quit immediately. If set to True, the livelock in the initial component will be reported but otherwise allowed. Any products of synchronization that retain those livelocks are still considered incompatible however, so they will not be used. If a livelock should happen to be pruned away in the (reduced!) synchronous product, then the final result may still be compatible.

A separate command-line module is available to use the scheduler. It is called `scheduler_commandline.py` and accessed from the `/source_code/cases` directory by calling `python ../scheduler_commandline [strategy=optional strategy argument] [complete=optional completeness argument] schedule=[mandatory schedule name]`. If no strategy is chosen it will use FIFO. If completeness is not specified then `self.contradictions` in the scheduler class instance will be set to False. To set it to true use `complete=True` as an argument.

We can produce a visualization of an automaton with the `graph.py` module. It contains the `makeGraph(subject = TA.TA())` function, which takes a single TA object as an argument. This module can also be accessed from a commandline with `graph_commandline.py` using the name of the automaton file as an argument. An optional argument can be given to specify the file format, otherwise the format is Scalable Vector Graphics (SVG).

The programs allow for automated experimentation; a simple example of this can be found in the `experiment.py` file.

Files, contents:

- **TA.py** – the team automaton class.
- **synchronize.py** – function to compute the synchronous product of a set of n team automata.
- **scheduler.py** – runs nested compatibility testing according to a strategy.
- **scheduler_commandline.py** – run a schedule from the command line.

Called from `/source_code/cases` as:

```
python ../scheduler_commandline [strategy=optional strategy argument] [complete=optional completeness argument] schedule=[mandatory schedule name]
```

- **graph.py** – makes graphical output.

- **graph_commandline.py** – command-line interface for `graph.py`.

Called from `/source_code/cases` as:

```
python ../graph_commandline [optional file format argument]
[mandatory filename argument]
```

- **experiment.py** – example use of the other modules to test a collection of schedules of automata at the press of a button.

Called from `/source_code/cases` as:

```
python ../experiment.py
```

Dependencies:

- Python 2.7 modules:

- `functools`
- `graphviz`
- `itertools`
- `math`
- `operator`
- `os`
- `pyparsing`
- `random`

- `Graphviz`

Tested on Ubuntu 14.04 LTS

5.1 Other software

As far as we are aware, only one earlier software implementation of team automata exists, written by Jonas Casanova and found on `sourceforge.net` under the name **synchroteam**. [3] It is available under the GNU General Public License version 2.0 (GPLv2).

It computes products of automata; not only synchronous products but it also allows the specification of other synchronization types. It does not however determine compatibility. We have aimed to maintain compatibility with `synchroteam` by using the same file format for components.

We were inspired by its use of `Graphviz` to automatically produce visualizations from component files, and adapted the technique to use Python's `Graphviz` module.

6 Experiments

We have done some initial experiments with our Python implementation, to get a first impression of the effect of nested compatibility checking using different strategies. We used `experiment.py` to run our cases through `scheduler.py` using strategies detailed below. In all cases did we use the same setting for our synchronous product/compatibility checking function `synchronous_product_effort(components = [], find_all_contradictions=True)`: we set it to stop as soon as the first contradiction was found.

We used the following strategies:

- **action similarity:** $fitness = \frac{|\Sigma_1 \cap \Sigma_2|}{|\Sigma_1 \cup \Sigma_2|}$

The idea being that components that share relatively many actions are important to each other and that it is therefore a good idea to combine those first.

- **action dissimilarity:** $fitness = \frac{-|\Sigma_1 \cap \Sigma_2|}{|\Sigma_1 \cup \Sigma_2|}$

This strategy was inspired as the opposite of action similarity, to see if anything interesting happens.

- **first in first out:** $fitness = last\ assigned\ fitness - 1$

The oldest pair is always the favourite for synchronization testing. All original components will be used (if possible) before any products of previous synchronizations are used.

In the passing the ball example, this would have resulted in a series of Figures 6, 7, 11, 13, 10.

- **last in first out:** $fitness = last\ assigned\ fitness + 1$

Newly produced components will always be favourite to be paired with second-youngest components. This corresponds to Figures 6–10.

- **Q minimization:** $fitness = -|Q_1| \times |Q_2|$

The idea being that we want to keep our components as small as possible so that synchronizing them is algorithmically cheaper.

In the passing the ball example, this would have resulted in a series of Figures 6, 7, 11, 13, 10.

- **Q maximization:** $fitness = |Q_1| \times |Q_2|$

Again an anti-strategy, to see if there are situations where going for the big components is better.

This corresponds to Figures 6–10.

- **all at once:** synchronize all components together at once, instead of pairwise. The goal of all the other strategies is to beat this one.

This corresponds to Figure 6 going directly going to Figure 10.


It should be noted that with *action (dis)similarity* and *Q minimization/maximization* multiple pairs might get the same fitness score. In that case the program breaks ties in a nondeterministic manner.

The cases we used were the components already mentioned earlier (coffee machine/drinker, workplace, passing the ball) as well as a wedding (using the groom, but with his livelock removed), a token ring system that was also used as an example in [2] as well as an abbreviated version of it (tokenRing2), and a settlement process inspired by [6]. All of these are included as component specifications in the appendix.


The results are as follows:

	Compatible?	action sim	action dissim	fifo	lifo	Q min	Q max	all at once
workplace	False	4	4	4	4	4	4	4
coffeeTime	False	20	20	20	20	20	20	20
tokenRing	False	2500	5157	421	1916	2637	4324	17
tokenRing2	True	633	2185	397	297	696	2185	252
settlement	False	176	222	222	217	222	261	132
wedding	True	430	430	320	430	430	430	320
passing	True	292	174	64	64	174	222	50

The results for “workplace” and “coffeeTime” are trivial: when a system contains only two components every strategy will achieve the same score.



Relative effort compared to 'all at once'



Several facts immediately jump out. Clearly, no nested compatibility checking strategy was able to beat the default method of checking all components together. The strategies *action dissimilarity* and *Q max* consistently performed poorly.

FIFO did surprisingly well, notably outperforming *Q minimization*. This could be a random result; the performance of *FIFO* and *LIFO* is influenced by the order in which components are presented in the original schedule. (A1, A2, A3, A4) could score differently than (A3, A1, A2, A4). Still, it is plausible that *FIFO* is simply a good strategy. Teams have a tendency to experience exponential growth, so making many few-generation products should work better than one many-generation product that keeps adding components to its output. Which is why the occasionally good performance of *LIFO* is notable. An explanation may be that schedules tend to be somewhat organized: a developer will list related components together.

Another significant finding is that no single strategy was the best one every time. This leads to the idea that perhaps different strategies are ideal for different kinds of components. A system where many components do not share actions with more than a few other components might benefit more from *action similarity* perhaps. If components tend to have spare transitions between their states, perhaps a different strategy dominates than if they are internally well-connected. And a system which is supposed to follow a single path from start to finish may benefit from a different strategy than one that branches into different paths.

In any case, we should ask ourselves whether there is any point to these strategies given that “all at once” outperformed all of them. An answer may be that Yes, there is a point: distribution. The all-at-once approach is far harder to implement in a distributed manner than some of these strategies. However, to pick the right strategy, a better insight is required in how to classify styles of team automata so that the best heuristic can be picked for nesting the work.

7 Conclusions

We have succeeded at our primary goal, to find and implement an algorithm that determines whether a system of n components is *compatible*. It meets the criteria we established: it delivers a result, consisting of a compatible/not-compatible verdict, as well as an actual team over the system.

Our methods maintain the integrity of the original input with regards to current and future behaviour; to an outside observer our reduced system appears to behave exactly the same as the original system would have. The only things we remove are unreachable states, useless transitions and useless internal actions. To ensure that the team has the same effect in future systems, we maintain useless input and output actions. In all cases, our results remain consistent with the basis inherited from [2], so that all our results are applicable in an environment according to their criteria.

While producing an algorithm was the primary goal, making sure it was a useful algorithm was of course also a priority. Analysis of its complexity revealed it to be exponential. So we sought to find improvements, to try to reduce its exponential complexity. We considered an alternative algorithm which turned out more unfeasible. We examined pruning methods and a divide and conquer approach. None of these succeeded in reducing the complexity to a lower order. It seems likely to us that this is simply not possible – finding out if states with certain properties are reachable closely resembles the question of computing the intersection of the languages of two finite automata, which is known to be an NP-complete problem ([5], page 266). So the question turns to finding ways to make the best of a bad situation. And it turns out there are some promising leads there.

We have been able to reduce the requirements for using a nested approach to compatibility checking and generalize the proof that it is possible. This makes it theoretically possible to split up a compatibility checking task into multiple smaller tasks so that we can distribute the workload, taking advantage of concurrency.

However, this requires a strategy to select an order in which to try subteams. We have demonstrated that several strategies are possible. The strategies we implemented are by no means the only possible ones however; we have provided a framework to which new strategies can be added. An interesting finding is that no nesting strategy consistently outperformed all the others. This suggests that we could discover which strategies work best in which situations.

7.1 Future Questions

A major open question is how useful nested checking is, considering the much better performance of checking all components together at once. Determining if the non-nested checking forms an upper bound on the performance of nested strategies would be worthwhile.

A topic related to nested checking is the lack of a typology of team automata. The model can represent many different things, but we have no standard of what to expect. We cannot say that it is normal for a component to have states sparsely or densely connected with transitions, or to have large or small alphabets. But we could develop a typology to identify certain categories of components which have typical combinations of such properties. The value here is that we can then work on picking a strategy that best matches the system we are working on.

A hindrance when testing team automata is the lack of a major test set. We looked at the possibility of generating a test set with examples of various properties. We believe this is in fact a very interesting avenue for future research because it posits the question: how do we generate an automaton that will have a property X that we desire – such as compatibility!

Since checking for incompatibility turns out to be so costly, developing methods to write models guaranteed to be compatible is perhaps a more rewarding route to take. Lemma 30 hints at where this may lead: developers may be instructed to ensure that the component they are responsible for will collaborate on certain actions, be always receptive to other actions etcetera. Obviously, models which guarantee collaboration on output are of interest here.

An entirely different direction of research would be to examine how well the notion of compatibility serves us in practice. We examined the use of team automata as applied to a project of Pieter Kwantes e.a. in [6] to translate Business Process Modeling Notation diagrams into statements that can be logically checked. At first, team automata appeared like an obvious tool for this, but translating the compliance requirements was rather an ad-hoc process. It was possible to model failure to follow a settlement protocol as a failure of compatibility, but this required quite a bit of engineering the model.

Related to that, we have some lingering doubts about the definition of livelock. There is a discrepancy between the definitions for livelock used in [2] and [1]; the latter is only concerned with reachable

livelocks. But is this really the sum of possible livelocks? Consider a system where two components need to both be ready before they can execute an action that the rest of the system needs in order to progress. However, component 1 is ready and component 2 is not. On the clock signal (an input for both), the ready component becomes unready and the unready component becomes ready. This system could still meet all the criteria for compatibility, but it looks like a two-component livelock in all but name.

References

- [1] Josep Carmona and Jordi Cortadella. Input/output compatibility of reactive systems. In *Fourth International Conference on Formal Methods in Computer-Aided Design*, pages 360–377, 2002.
- [2] Josep Carmona and Jetty Kleijn. Compatibility in a multi-component environment. *Theoretical Computer Science*, 484:1–15, 2013.
- [3] Jonas Casanova. Synchroteam: team automata synchronizer. <http://synchroteam.sourceforge.net/>, 2013. [Online; accessed 28-August-2015].
- [4] Clarence A. Ellis. Team automata for groupware systems. In *GROUP '97 Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge*, pages 415–424, 1997.
- [5] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- [6] Pieter M. Kwantes, Pieter Van Gorp, Jetty Kleijn, and Arend Rensink. Towards compliance verification between global and local process models. In Francesco Parisi-Presicce and Bernhard Westfechtel, editors, *Graph Transformation*, volume 9151 of *Lecture Notes in Computer Science*, pages 221–236. Springer International Publishing, 2015.
- [7] Maurice H. ter Beek, Clarence A. Ellis, Jetty Kleijn, and Grzegorz Rozenberg. Synchronizations in team automata for groupware systems. *Computer Supported Cooperative Work*, 12:21–69, 2003.

Index

- action-complementary, **12, 14, 27**
- actions, **7, 8**
- active behaviour, **24, 24**
- active behaviour of a component in a team is a subset of its solitary behaviour, **24**
- active computation, **23, 23**
- active computations, **23**
- alphabet, **7, 7, 8**

- behaviour, **5, 23, 23**

- collaborates, **12, 12, 14, 27**
- collaborates on output, **28**
- communicating, **9**
- compatibility, **10, 11, 11, 19**
- compatibility relation, **11, 21**
- compatible, **12, 12, 14, 19, 21, 28, 45**
- complete, **9, 14**
- complete transition space, **9**
- component's active behaviour in a team is a subset of active behaviour in subteams containing it, **24**
- composable system, **9, 9, 11, 15**
- computation, **8, 8**
- computations, **23**
- correspondence of compatibility, **28**
- correspondence of states, **27**
- correspondence of transitions, **27**

- deadlock-free, **12, 12, 14, 28**
- Definitions
 - active behaviour, **24**
 - active computation, **23**
 - alphabet, **7**
 - behaviour, **23**
 - compatibility, **11**
 - composable system, **9**
 - computation, **8**
 - determinism, **17**
 - directly incompatible state, **19**
 - projection, **7**
 - synchronous product, **10**
 - team automata, **9**
 - transition system, **8**
 - ultimately incompatible state, **21**
 - usefulness, **23**
- determinism, **17**
- deterministic, **17**
- directly incompatible, **19, 19, 37**
- directly incompatible state, **19**
- DIS, **19, 21, 37**
- domain, **9**
- domain enabled, **9**

- enabled, **8**

- I not in UDIS, **21**

- input domain, **9**
- input-domain enabled, **9**
- input-enabled, **8**

- Lemmas
 - correspondence of compatibility, **28**
 - correspondence of states, **27**
 - correspondence of transitions, **27**
 - I not in UDIS, **21**
 - persistence of livelock-freeness, **17**
- livelock, **11**
- livelock-free, **11, 12**

- nondeterministic, **17**

- output domain, **9**
- output-domain enabled, **9**

- persistence of livelock-freeness, **17**
- progressive, **12**
- projection, **7, 7**
- propagation of uselessness, **24**

- reachable, **8, 19, 24**
- reactive, **8**
- reactive transition system, **7, 8, 15**
- receptive, **12, 12, 14, 28**
- receptiveness, **12**
- Reduced Reactive Transition System, **26**
- RRTS, **26**
- RTS, **8, 11, 26**

- states, **8**
- states in the compatibility relation are not in UDIS, **21**
- synchronization, **10**
- synchronizations, **9**
- synchronous product, **10, 10–12, 15, 24, 25**
- synchroteam, **39**

- TA, **9**
- team automata, **9**
- team automaton, **7, 8, 9**
- terminating, **11**
- Theorems
 - active behaviour of a component in a team is a subset of its solitary behaviour, **24**
 - component's active behaviour in a team is a subset of active behaviour in subteams containing it, **24**
 - propagation of uselessness, **24**
 - states in the compatibility relation are not in UDIS, **21**
- transition system, **7, 8, 8**
- transitions, **7, 8**
- TS, **8, 17, 23**

- UDIS, **21, 21**

ultimately incompatible, **21**
ultimately incompatible state, **21**
unreachable, **8**
useful, **23, 26**
usefulness, **23**
useless, **23, 23**

A Code

```
#!/usr/bin/env python
2
class TA:
4
    def __init__(self, filename=None):
6        """Initialize from file or blank (no argument)

8        File should be formatted thus:

10        .model MODENAME
11        .inputs INPUT1 INPUT2
12        .outputs OUTPUT1 OUTPUT2
13        .internal INTERNAL1 INTERNAL2
14        .state graph
15        CURRENTSTATE1 ACTION1 NEXTSTATE1
16        CURRENTSTATE2 ACTION2 NEXTSTATE2
17        .marking INITIALSTATE
18        .end

20        CAPITALS indicate variables. lowercase indicates keywords.
21        """
22        # create blank TA
23        if filename is None:
24            self.name = "Anon"
25            self.states = set()
26            self.input_actions = set()
27            self.output_actions = set()
28            self.internal_actions = set()
29            self.initial_state = ""
30            self.transitions = {} # (current state, action) : next state
31            self.message_losses = set()
32            self.deadlocks = set()
33            return
34        # create TA from file
35        f = open(filename, "rb")
36        reading_transitions = False
37        self.transitions = {}
38        self.states = set()
39        self.message_losses = set()
40        self.deadlocks = set()
41        for line in f:
42            if ".model" in line:
43                self.name = line.strip().split(" ")[1]
44            elif ".inputs" in line:
45                self.input_actions = set(line.strip().split(" ")[1:])
46            elif ".outputs" in line:
47                self.output_actions = set(line.strip().split(" ")[1:])
48            elif ".internal" in line:
49                self.internal_actions = set(line.strip().split(" ")[1:])
50            elif ".state graph" in line:
51                reading_transitions = True
52            elif ".marking" in line:
53                reading_transitions = False
54                temp = line.strip().split("{")[1].split("}")
55                if "" in temp:
56                    print "Error: multiple initial states."
57                    print line
58                self.initial_state = temp
59            elif ".end" in line:
60                break
61            elif reading_transitions == True:
62                temp = line.strip().split(" ")
63                if len(temp) != 3:
64                    print "Error: weird transition:"
65                    print line
66                if (temp[0], temp[1]) in self.transitions:
67                    print "Error: nondeterministic transition:"
68                    print line
69                else:
70                    self.add_transition(temp[0], temp[1], temp[2])
```

```

72 def get_all_actions(self):
73     """Return the union of input, output and internal alphabets."""
74     return self.output_actions.union(self.input_actions).\
75         union(self.internal_actions)
76
77 def sanity_check(self):
78     """Check whether a TA is properly initialized and whether it conforms
79     to the requirements for composability.
80
81     TA sanity requirements:
82     - It must have a nonempty initial state. This implementation uses
83     deterministic TAs, so exactly one initial state is required.
84     - The initial state must belong to the overall set of states.
85     - Transitions go from states in Q to other states in Q.
86     - Transitions use an action from the Input, Output or Internal
87     alphabet.
88     - Input, Output and Internal alphabets cannot overlap.
89
90     Note that this function does NOT check for livelock-freeness.
91     """
92     if self.initial_state == "":
93         print "Initial state not set."
94         return False
95     if self.initial_state not in self.states:
96         print "Initial state not in Q."
97         return False
98     if not self.input_actions.isdisjoint(self.output_actions):
99         print "Input and output actions overlap"
100        return False
101     if not self.input_actions.isdisjoint(self.internal_actions):
102         print "Input and internal actions overlap"
103         return False
104     if not self.output_actions.isdisjoint(self.internal_actions):
105         print "Output and internal actions overlap"
106         return False
107     for item in self.transitions:
108         if item[0] not in self.states:
109             print "Bad transition ", item, self.transitions[item],
110             print "; starting state not in Q."
111             return False
112         if ((item[1] not in self.input_actions) and
113             (item[1] not in self.output_actions) and
114             (item[1] not in self.internal_actions)):
115             print "Bad transition ", item, self.transitions[item],
116             print "; action not in Sigma."
117             return False
118         if self.transitions[item] not in self.states:
119             print "Bad transition ", item, self.transitions[item],
120             print "; destination state not in Q."
121             return False
122     return True # no insanity detected
123
124 def screendump(self):
125     """Dump the automaton to the screen.
126     """
127     print "Name: " + self.name
128     print "  states:"
129     for item in self.states:
130         print "    " + item
131     print "  input actions:"
132     for item in self.input_actions:
133         print "    " + item
134     print "  output actions:"
135     for item in self.output_actions:
136         print "    " + item
137     print "  internal actions:"
138     for item in self.internal_actions:
139         print "    " + item
140     print "  initial state: " + self.initial_state
141     print "  transitions:"
142     for item in self.transitions:
143         print ("    (" + item[0] + ", " + item[1] + ", " +
144             self.transitions[item] + ")")

```

```

146     if len(self.message_losses) > 0:
147         print "    MESSAGE LOSSES:"
148         for item in self.message_losses:
149             print "        " + item[0] + " > " + item[1]
150     if len(self.deadlocks) > 0:
151         print "    DEADLOCKS:"
152         for item in self.deadlocks:
153             print "        " + item
154
155 def save(self, filename=None):
156     """Writes the automaton to file, using the fileformat from Synchroteam
157
158     Insists on ".out" as file extension; appending if necessary.
159     If no filename is given, the automaton name is used as filename.
160
161     For more information about Synchroteam:
162     http://sourceforge.net/projects/synchroteam/
163     """
164     if filename == None:
165         filename = self.name + ".out"
166     elif not filename.endswith(".out"):
167         filename = filename + ".out"
168     with open(filename, "wb") as d:
169         d.write(".model " + self.name + "\r\n")
170         d.write(".inputs")
171         for item in self.input_actions:
172             d.write(" " + item)
173         d.write("\r\n")
174         d.write(".outputs")
175         for item in self.output_actions:
176             d.write(" " + item)
177         d.write("\r\n")
178         d.write(".internal")
179         for item in self.internal_actions:
180             d.write(" " + item)
181         d.write("\r\n")
182         d.write(".state graph\r\n")
183         for item in self.transitions:
184             d.write("    " + item[0] + " " + item[1] + " " + self.transitions[item])
185             d.write("\r\n")
186         d.write(".marking {" + self.initial_state + "}\r\n")
187         d.write(".end\r\n")
188
189 def rename_internal_actions(self):
190     """Renames internal actions by prepending TA name.
191
192     In a composable systems, component automata are not allowed to share
193     internal actions. This method aims to ensure that all internal actions
194     receive a unique name.
195     """
196     temp = set()
197     for action in self.internal_actions:
198         temp.add(self.name + "-" + action)
199     self.internal_actions = temp
200
201 def is_action_enabled(self, state, action):
202     """Returns true if the action is enabled, false if not.
203
204     Action is enabled if (current state, action, next state) IN transitions
205     """
206     if (state, action) in self.transitions:
207         return True
208     else:
209         return False
210
211 def get_enabled_actions(self, state):
212     """Get all actions that are enabled in the given state.
213     """
214     return {x for x in self.get_all_actions() if self.is_action_enabled(state, x)}
215
216 def get_next_state(self, state, action):
217     """Get the next state if the given action is taken in the given state.
218     """

```

```

218     if (state, action) in self.transitions:
219         return self.transitions[(state, action)]
220     else:
221         print "ERROR:", state, action, "not defined."
222         return None
223
224 def get_directly_reachable_states(self, state):
225     """Get the states that can be directly reached from the current state
226     """
227     return {x for x in self.states if x in {self.get_next_state(state, y)
228         for y in self.get_enabled_actions(state)}}
229
230 def get_reachable_states(self):
231     """Get all states that can be reached eventually
232
233     Starting in the initial state, add the directly reachable states of
234     known reachable states until the set stops growing.
235     """
236     result = {self.initial_state}
237     todo = self.get_directly_reachable_states(self.initial_state)
238     while len(todo) != 0:
239         current = todo.pop()
240         result.add(current)
241         todo |= {x for x in self.get_directly_reachable_states(current) if
242             x not in result}
243     return result
244
245 def get_useful_transitions(self):
246     """Returns the set of all useful transitions.
247
248     A transition is useful if it starts in a reachable state.
249     """
250     return {(x[0], x[1], self.transitions[x]) for x in self.transitions if
251         x[0] in self.get_reachable_states()}
252
253 def get_useful_actions(self):
254     """Returns the set of all useful actions.
255
256     An action is useful if it is used by at least one useful transition.
257     """
258     return {x[1] for x in self.get_useful_transitions()}
259
260 def get_unreachable_states(self):
261     """Returns all unreachable states.
262     """
263     return {x for x in a.states if x not in a.get_reachable_states()}
264
265 def get_useless_transitions(self):
266     """Returns all useless transitions.
267     """
268     return {(x[0], x[1], self.transitions[x]) for x in self.transitions if
269         x[0] not in self.get_reachable_states()}
270
271 def get_useless_actions(self):
272     """Returns all useless actions.
273     """
274     return {x[1] for x in self.get_all_actions() if x not in
275         self.get_useful_transitions()}
276
277 def add_state(self, state):
278     """Add a new state to the TA.
279     """
280     if isinstance(state, str):
281         self.states.add(state)
282     else:
283         print "Error: adding a non-string state"
284
285 def add_internal_action(self, action):
286     """Adds a new internal action to the TA.
287     """
288     if isinstance(action, str):
289         if (action in self.input.actions):
290             print "Error: action is already an input action."

```

```

292         elif (action in self.output_actions):
293             print "Error: action is already an output action."
294         else:
295             self.internal_actions.add(action)
296     else:
297         print "Error: adding a non-string action"
298
299 def add_input_action(self, action):
300     """Adds a new input action to the TA.
301     """
302     if isinstance(action, str):
303         if (action in self.internal_actions):
304             print "Error: action is already an internal action."
305         elif (action in self.output_actions):
306             print "Error: action is already an output action."
307         else:
308             self.input_actions.add(action)
309     else:
310         print "Error: adding a non-string action"
311
312 def add_output_action(self, action):
313     """Adds a new output action to the TA.
314     """
315     if isinstance(action, str):
316         if (action in self.internal_actions):
317             print "Error: action is already an internal action."
318         elif (action in self.input_actions):
319             print "Error: action is already an input action."
320         else:
321             self.output_actions.add(action)
322     else:
323         print "Error: adding a non-string action"
324
325 def add_transition(self, start_state, action, end_state):
326     """Adds a new transition to the TA.
327
328     If start/end state not yet defined, will also be added.
329     Actions cannot be implicitly added because we'd lack information on
330     whether they're input/output/internal.
331     """
332     self.add_state(start_state)
333     self.add_state(end_state)
334     if action not in self.get_all_actions():
335         print "Error: action not defined:", action
336     else:
337         self.transitions[(start_state, action)] = end_state
338
339 def set_initial_state(self, state):
340     """Set the initial state to a new value.
341     """
342     if isinstance(state, str):
343         if (state in self.states):
344             self.initial_state = state
345         else:
346             print "Error: suggested new initial state not in set of states."
347     else:
348         print "Error: new initial state not a string."
349
350 def remove_state(self, state):
351     """Remove a state from the TA's set of states.
352
353     It is an error to remove the initial state.
354     When a state is removed, all transitions starting or ending
355     in that state are also removed.
356     """
357     if state in self.states:
358         if state is self.initial_state:
359             print "Error: trying to remove initial state."
360         else:
361             self.states.remove(state)
362             for t in self.get_enabled_actions(state):
363                 del self.transitions[(state, t)]
364             for t in {t for t in self.transitions if self.transitions[t]

```



```

364         is state}:
365             del self.transitions[t]
366     else:
367         print "Error: trying to remove nonexistent state."
368
369 def remove_transition(self, start_state, action):
370     """Remove a transition.
371     """
372     if (start_state, action) in self.transitions:
373         del self.transitions[(start_state, action)]
374     else:
375         print "Error: that transition doesn't exist."
376
377 def remove_action(self, action):
378     """Remove an action from the alphabet.
379
380     All transitions using that action will also be removed.
381     """
382     if action in self.internal_actions:
383         self.internal_actions.remove(action)
384         for t in {t for t in self.transitions if t[1] is action}:
385             del self.transitions[t]
386     elif action in self.input_actions:
387         self.input_actions.remove(action)
388         for t in {t for t in self.transitions if t[1] is action}:
389             del self.transitions[t]
390     elif action in self.output_actions:
391         self.output_actions.remove(action)
392         for t in {t for t in self.transitions if t[1] is action}:
393             del self.transitions[t]
394     else:
395         print "Error: cannot remove nonexistent action"
396
397 def reduce_this_automaton(self):
398     """Reduces this automaton.
399
400     Removes all unreachable states, which will also remove all useless
401     transitions; useless actions remain.
402     """
403     for s in self.get_unreachable_states():
404         self.remove_state(s)
405
406 def get_enabled_internal_actions(self, state):
407     """Get all internal actions that are enabled in the given state.
408     """
409     return {x for x in self.internal_actions if self.is_action_enabled(
410         state, x)}
411
412 def get_internal_next_states(self, state):
413     """Get all states that can be directly reached by an internal action.
414
415     Intended for livelock checking.
416     """
417     return {x for x in self.states if x in {self.get_next_state(state, y)
418         for y in self.get_enabled_internal_actions(state)}}
419
420 def detect_livelock(self, quiet=False):
421     """Find livelocks anywhere in the automaton.
422
423     Livelocks are found by selecting states with an outgoing internal
424     transition, then for each of them calling a graph exploration function
425     that only travels along internal transitions.
426
427     Note that this function can also detect unreachable livelocks if the
428     automaton contains unreachable states. It is recommended to use this
429     function on automata already reduced with the reduce_automaton or
430     get_reduced_automaton functions.
431     """
432     safe = set()
433     todo = {x for x in self.states if
434         len(self.get_enabled_internal_actions(x)) > 0}
435     for item in todo.difference(safe):
436         safe, live = self.explore_internal_paths(safe, so_far = [item], quiet=quiet)

```

```

438         if live:
439             return True
440     return False

441 def explore_internal_paths(self, safe, so_far, quiet=False):
442     """Explore internal paths, to detect livelock

443     So_far is the states already reached. If a possible next state is also
444     in so_far, that means a livelock exists, and the function returns.
445     Otherwise it recurses on next states until it finds a longer livelock
446     or it exhausts all possibilities.
447     """
448     for x in self.get_internal_next_states(so_far[-1]).difference(safe):
449         if x in so_far:
450             if quiet==False:
451                 print "Livelock found:", so_far, x
452             return (safe, True)
453         else:
454             so_far.append(x)
455             safe, live = self.explore_internal_paths(safe, so_far, quiet)
456             if live:
457                 return (safe, True)
458             so_far.pop()
459     safe.add(so_far[-1])
460     return (safe, False)

461 def get_reduced_version(self):
462     """Returns a reduced version of this automaton.

463     The reduced version will contain:
464     - States IFF they are reachable
465     - Transition IFF they are useful
466     - Internal actions IFF they are useful
467     - All Input and Output actions
468     """
469     RRTS = TA()
470     RRTS.name = "R_" + self.name
471     RRTS.initial_state = self.initial_state
472     RRTS.input_actions = self.input_actions
473     RRTS.output_actions = self.output_actions
474     todo = [self.initial_state]
475     done = set()
476     while len(todo) != 0:
477         for action in self.get_enabled_actions(todo[0]):
478             if action not in RRTS.get_all_actions():
479                 RRTS.add_internal_action(action)
480                 next = self.get_next_state(todo[0], action)
481                 RRTS.add_transition(todo[0], action, next)
482                 if next not in todo and next not in done:
483                     todo.append(next)
484             done.add(todo[0])
485         todo.pop(0)
486     return RRTS

487 def is_automaton_reduced(self):
488     """Tells you if an automaton is already reduced.

489     This operation is as expensive as generating a reduced automaton and
490     should be avoided.
491     """
492     if (len(self.states) == len(self.get_reachable_states()) and
493         len(self.transitions) == len(self.get_useful_transitions())):
494         return True
495     else:
496         return False

497 def is_automaton_compatible(self):
498     """Returns true if compatible according to definition.

499     This requires the following:
500     - No livelocks
501     - No message loss (always receptive)
502     - Deadlock-free

```

```
510     Note: this assumes the automaton was constructed with a synchronous
512     production function that records incompatibilities, or that it is an
514     elementary automaton which has no "history" for incompatibilities.
516     """
518     if self.detect_livelock(quiet=True):
519         return False
520     if (len(self.message_losses) + len(self.deadlocks)) == 0:
521         return True
522     return False
```

../source_code/TA.py

```

1  #! /usr/bin/env python
3  from TA import TA
   import itertools
5
7  def take_transition(current, action, participants, components, result):
   """Follow a transition (current, action)→next
9
   Computes the next-state arrived at by executing 'action', then adds
11  the transition to result.transitions, and if necessary also adds the 'next'
   state to result.states.
13
   The 'participants' are the components of 'result' that take part in the
15  transition.
17
   Returns the 'next' state to the caller.
   """
19  next = tuple(current[s] if s not in participants else
   components[s].get_next_state(current[s], action)
21  for s in range(len(components)))
   result.add_transition("-".join(current), action, "-".join(next))
23  return next
25 def add_if_needed(state, todo, done):
   """Adds the state to 'todo' if it's not been encountered before.
27  """
   if state not in todo and state not in done:
29     todo.append(state)
31 def synchronous_product_effort(components = [], find_all_contradictions=True):
   """Compute the synchronous product of a list of components.
33
   Will also determine message loss and deadlock situations.
35
   Does not pay attention to livelocks. Livelocks will not cause the algorithm
37  trouble. Note that livelocks could have been checked and
   prevented at the component preprocessing stage. If the components are
39  livelock-free, then the synchronous product will also be livelock-free.
   (If the components were not livelock-free, then the synchronous product
41  might still be livelock-free, because it is a Reduced system.)
43
   If the result is livelock-free, and len(result.message_losses) ==
   len(result.deadlocks) == 0, then results.states is a compatibility relation.
45  """
   effort = 0
47  # check sanity of input
   if len(components) == 0:
49     print "ERROR: no components."
     return None
51  elif len(components) == 1:
     return components[0]
53  # check for composability
   for pair in list(itertools.combinations(components, 2)):
55     if pair[0] == pair[1]:
         print "Error: synchronizing duplicate component"
         return None
57     if not pair[0].internal_actions.isdisjoint(pair[1].get_all_actions()):
59         print "ERROR: shared internal action"
         return None
61     if not pair[1].internal_actions.isdisjoint(pair[0].get_all_actions()):
         print "ERROR: shared internal action"
         return None
63  result = TA()
65  result.name = "XR_" + "-".join([c.name for c in components])
   # fill alphabets
67  # internal actions will not be added to 'result' unless they prove useful.
   internal_actions = set().\
69     union(*[c.internal_actions for c in components])
   # external actions will always be added to 'result'.
71  result.output_actions = set().\
     union(*[c.output_actions for c in components])
73  result.input_actions = set().\

```

```

75         union(*[c.input_actions for c in components]).\
76         difference(result.output_actions)
77     # compute implied sub-alphabets and domains
78     external_actions = result.output_actions.union(result.input_actions)
79     communicating_actions = set()
80     shared_input_actions = set() # shared input, but not output
81     shared_output_actions = set() # shared output, but not input
82     noncommunicating_input_actions = set()
83     noncommunicating_output_actions = set()
84     inp_dom = {}
85     out_dom = {}
86     int_dom = {}
87     for action in result.get_all_actions():
88         if action in internal_actions:
89             continue # useful internal actions will be added later
90     inp = {c for c in range(len(components))
91           if action in components[c].input_actions}
92     out = {c for c in range(len(components))
93          if action in components[c].output_actions}
94     Li = len(inp)
95     Lo = len(out)
96     if Li > 0:
97         inp_dom[action] = inp
98     if Lo > 0:
99         out_dom[action] = out
100    if Li * Lo > 0:
101        communicating_actions.add(action)
102    elif Li == 1:
103        noncommunicating_input_actions.add(action)
104    elif Li > 1:
105        shared_input_actions.add(action)
106    elif Lo == 1:
107        noncommunicating_output_actions.add(action)
108    elif Lo > 1:
109        shared_output_actions.add(action)
110    for c in range(len(components)):
111        for action in components[c].internal_actions:
112            int_dom[action] = c
113    # initialize statespace
114    done = set()
115    todo = [tuple(c.initial_state for c in components)]
116    result.initial_state = "-".join(todo[0])
117    result.add_state(result.initial_state)
118    # explore the statespace
119    while len(todo) > 0:
120        # explore one particular state
121        possible_deadlock = "unknown"
122        current = todo.pop()
123        done.add(current)
124        current_state = "-".join(current)
125        for action in communicating_actions:
126            effort += len(out_dom[action])
127            effort += len(inp_dom[action])
128            if all(components[c].is_action_enabled(current[c], action)
129                 for c in out_dom[action]): # output domain enabled
130                if all(components[c].is_action_enabled(current[c], action)
131                     for c in inp_dom[action]): # input domain enabled
132                    next = take_transition(current, action,
133                                         set().union(inp_dom[action],
134                                                      out_dom[action]), components, result)
135                    add_if_needed(next, todo, done)
136                    possible_deadlock = "averted"
137                else: # message loss
138                    result.message_losses.add((current_state, action))
139                    if find_all_contradictions == False:
140                        return (result, effort)
141            else: # output not domain enabled
142                if all(components[c].is_action_enabled(current[c], action)
143                     for c in inp_dom[action]): # input domain enabled
144                    if possible_deadlock == "unknown":
145                        possible_deadlock = "threat"
146        for action in shared_input_actions:
147            effort += len(inp_dom[action])

```

```

147         if all(components[c].is_action_enabled(current[c], action)
148                for c in inp_dom[action]): # input domain enabled
149             next = take_transition(current, action,
150                                   inp_dom[action], components, result)
151             add_if_needed(next, todo, done)
152             possible_deadlock = "averted"
153     for action in shared_output_actions:
154         effort += len(out_dom[action])
155         if all(components[c].is_action_enabled(current[c], action)
156                for c in out_dom[action]): # output domain enabled
157             next = take_transition(current, action,
158                                   out_dom[action], components, result)
159             add_if_needed(next, todo, done)
160             possible_deadlock = "averted"
161     for action in noncommunicating_input_actions:
162         effort += len(inp_dom[action])
163         if all(components[c].is_action_enabled(current[c], action)
164                for c in inp_dom[action]): # input domain enabled
165             next = take_transition(current, action,
166                                   inp_dom[action], components, result)
167             add_if_needed(next, todo, done)
168             possible_deadlock = "averted"
169     for action in noncommunicating_output_actions:
170         effort += len(out_dom[action])
171         if all(components[c].is_action_enabled(current[c], action)
172                for c in out_dom[action]): # output domain enabled
173             next = take_transition(current, action,
174                                   out_dom[action], components, result)
175             add_if_needed(next, todo, done)
176             possible_deadlock = "averted"
177     for action in internal_actions:
178         effort += 1
179         if components[int_dom[action]].is_action_enabled(
180             current[int_dom[action]], action):
181             result.internal_actions.add(action) # useful => include
182             next = take_transition(current, action,
183                                   {int_dom[action]}, components, result)
184             add_if_needed(next, todo, done)
185             possible_deadlock = "averted"
186     if possible_deadlock == "threat": # threat was not averted
187         result.deadlocks.add(current_state)
188         if find_all_contradictions == False:
189             return (result, effort)
190     # post-exploration checks
191     if result.sanity_check():
192         return (result, effort)
193     else:
194         print "Error: resulting TA is not sane."
195         return None

```

../source_code/synchronize.py

```

1 # module to make graphs for Team Automata
2 # strongly inspired by http://matthiaseisen.com/articles/graphviz/
3
4 # external modules
5 import functools
6 import graphviz
7 import os
8
9 # local modules
10 import TA
11
12 # default format
13 my_format = 'svg'
14
15 # shorthand
16 digraph = functools.partial(graphviz.Digraph, format=my_format)
17
18 def changeformat(my_format):
19     """ Change the output format
20     """
21     return functools.partial(graphviz.Digraph, format=my_format)
22
23 def add_nodes(graph, nodes):
24     """ Add nodes to the graph
25
26     Nodes are either entered as a string, or as a tuple.
27
28     When entered as a tuple, it is done in the form (string, {}) and the second
29     element holds style information for graphviz.
30     """
31     for n in nodes:
32         if isinstance(n, tuple):
33             graph.node(n[0], **n[1])
34         else:
35             graph.node(n)
36     return graph
37
38 def add_edges(graph, edges):
39     """ Adds edges between nodes.
40
41     Edges are either entered as a string, or as a tuple.
42
43     When entered as a tuple, it is done in the form (string, {}) and the second
44     element holds style information for graphviz.
45     """
46     for e in edges:
47         if isinstance(e[0], tuple):
48             graph.edge(*e[0], **e[1])
49         else:
50             graph.edge(*e)
51     return graph
52
53 def make_digraph(name, nodes, edges):
54     """ Make a directed graph from a set of nodes and edges connecting them.
55     """
56     add_edges(
57         add_nodes(digraph(), nodes),
58         edges
59     ).render(name)
60
61
62 def makeGraph(subject = TA.TA()):
63     """ Make the graph for a TA.
64
65     The argument should be a TA object.
66     """
67     nodes = list(subject.states.difference(set(subject.initial_state)))
68     nodes.append((subject.initial_state, {'penwidth': '3'}))
69     edges = []
70     for t in subject.transitions:
71         if t[1] in subject.internal_actions:
72             edges.append(((t[0], subject.transitions[t]), {'label': ' '+t[1]+'@ '}))
73

```

```

75     elif t[1] in subject.input_actions:
76         edges.append(((t[0], subject.transitions[t]), {'label': ' '+t[1]+'? '}))
77     elif t[1] in subject.output_actions:
78         edges.append(((t[0], subject.transitions[t]), {'label': ' '+t[1]+'! '}))
79     else:
80         print "Error: transition using action of unknown type:", t, subject.transitions[t]
81 make_digraph(subject.name, nodes, edges)
82 os.remove(subject.name)

```

../source_code/graph.py

```

1 import sys
2 import graph
3 import TA
4
5 # this script is intended to allow us to make graphs from the commandline.
6 # usage:
7 # python graph_commandline.py filename
8 # produces a graph with the TA's actual name as 'name.svg', regardless of the filename
9
10 if len(sys.argv) == 3:
11     graph.digraph = graph.changeformat(sys.argv[1])
12     subject = TA.TA(sys.argv[2])
13 else:
14     subject = TA.TA(sys.argv[1])
15 graph.makeGraph(subject)

```

../source_code/graph_commandline.py


```

1  #! /usr/bin/env python
3  import itertools
4  import random
5  import operator
6  import math
7  from pyparsing import nestedExpr
9  from TA import TA
10 import synchronize
11
12
13
14
15
16
17 class Scheduler:
18
19     def __init__(self, filename=None, strategy=None, find_all_contradictions=True):
20         """ Create a new schedule
21
22         If a filename is given, that schedule will be loaded. Otherwise an empty
23         schedule is returned, which can be used to resolve sub-schedules.
24         """
25         self.components = {}
26         self.combinations = {}
27         self.ordered_components = []
28         self.initial_effort = 0
29         self.set_strategy(strategy)
30         self.contradictions = find_all_contradictions
31         if filename != None:
32             goal = self.read_schedule_file(filename)
33             self.initial_effort = self.resolve_components_effort(goal)
34         else:
35             return
36
37     def read_schedule_file(self, filename):
38         """ Read in a schedule file and parse it
39         """
40         data = ""
41         with open(filename, "r") as myfile:
42             goal=myfile.read().replace('\n', '').replace(' ', ' ') \
43                 .replace(' ', ' ')
44             schedule = nestedExpr('(', ')').parseString(goal).asList()[0]
45             return schedule
46
47     def resolve_components_effort(self, goal):
48         """ Distinguish base components and subschedules in the schedule.
49
50         Base components are loaded. Subschedules are recursively resolved into
51         components, which are then added.
52
53         The effort expended on subschedules is inherited by the superschedule.
54         """
55         total_effort = 0
56         for item in range(len(goal)):
57             if isinstance(goal[item], str):
58                 self.load_component(goal[item])
59             if isinstance(goal[item], list):
60                 subteam = Scheduler(None, self.strategy)
61                 effort = subteam.resolve_components_effort(goal[item])
62                 total_effort += effort
63                 result, effort = subteam.run_effort()
64                 total_effort += effort
65                 result.save()
66                 self.load_component(result.name)
67         return total_effort
68
69     def load_component(self, name):
70         """ Load a component from file and preprocess it.
71
72         Components are reduced and then checked for livelocks. Even if a livelock
73         is found, the loading continues.

```

```

75     The detect_livelock method of TA will report what livelock was found;
76     this function reports in which component it was found.
77     """
78     temp = TA(name + ".out")
79     temp = temp.get_reduced_version()
80     if temp.detect_livelock():
81         print "    in component:", name
82         if self.contradictions == False:
83             print "Stopping now, because complete search (despite finding a contradiction
already) is OFF."
84             exit(0)
85     self.components["R_"+name] = temp
86     self.ordered_components.append("R_"+name)
87
88 def collaborates_nicely(self, choice):
89     """ Tests if a choice of next job enjoys collaboration of remainder.
90
91     Will return True if the automata of choice share no output action with
92     the remaining components that is not also an input action in some of
93     the remaining components; else False.
94     """
95     left_out = set().union(self.components[choice[0]].output_actions,
96                             self.components[choice[1]].output_actions)
97     right_out = set()
98     for c in self.components:
99         if c == choice[0] or c == choice[1]:
100             continue
101         right_out = right_out.union(self.components[c].output_actions)
102     if len(left_out.intersection(right_out)) > 0:
103         return False
104     return True
105
106 def get_next_job(self, previous, choice):
107     """ Determine what components to try to synchronize next.
108
109     'previous' is the previous synchronization result, which is examined
110     here.
111
112     If 'previous' is compatible:
113         - The components used to make 'previous' are removed from the
114         remaining unsynchronized components
115         - All combinations using the old components are removed from the
116         list of possibilities
117         - 'previous' is added to the set of remaining unsynchronized
118         components
119         - Combinations with the new component are computed and added to the
120         list of possibilities
121     If 'previous' is NOT compatible:
122         - It is not added to the set of remaining unsynchronized components
123         - It's constituent components stay in the set of remaining
124         unsynchronized components
125         - The combination of its constituent components is removed from
126         the list of possible combinations
127
128     If viable combinations remain they're sorted by fitness. One by one they
129     will be tested to see if they are output-action-complementary. The first
130     one meeting this criterium is returned.
131     Else, a last-ditch attempt is made to synchronize all remaining
132     components at once.
133     If that also fails, there is no next job.
134     """
135     # STEP 1: process previous job
136     if isinstance(previous, TA):
137         # if previous pairing turned out incompatible
138         if previous.is_automaton_compatible():
139             # if previous pairing came from combination
140             if (choice[0].name, choice[1].name) in self.combinations:
141                 # dispose of the incompatible combination
142                 del self.combinations[(choice[0].name, choice[1].name)]
143             # else a final attempt turned out incompatible
144         else:
145             return None # no more jobs to try

```

```

147         else:
148             # if previous pairing came from combination
149             if (choice[0].name, choice[1].name) in self.combinations:
150                 # flush obsolete combinations and synchronized components
151                 obsolete = [x for x in self.combinations.keys() if (
152                     x[0] == choice[0].name or
153                     x[1] == choice[0].name or
154                     x[0] == choice[1].name or
155                     x[1] == choice[1].name)]
156                 for item in obsolete:
157                     del self.combinations[item]
158                     del self.components[choice[0].name]
159                     del self.components[choice[1].name]
160                 # add and analyze the new component
161                 self.components[previous.name] = previous
162                 self.ordered_components.append(previous.name)
163                 for component in self.components:
164                     if component != previous.name:
165                         self.analyze(component, previous.name)
166                 # success was the result of final synchronization
167                 else:
168                     return None # no more jobs needed
169
170 # STEP 2: select next job
171 # if viable pairs are still available
172 if len(self.combinations) > 0:
173     # get choices sorted by fitness
174     choices = sorted(self.combinations, key=self.combinations.get,
175                     reverse=True)
176     # now return the highest-fitness possible choice that collaborates nicely
177     for choice in choices:
178         if self.collaborates_nicely(choice):
179             return [self.components[choice[0]], self.components[choice[1]]]
180     # if no pair was found that collaborates nicely, try all at once
181     self.combinations.clear()
182     return list(self.components.values())
183 # elif 3+ components remain but no combinations, try all at once
184 elif len(self.components) > 2:
185     return list(self.components.values())
186 # else give up
187 else:
188     return None
189
190 def run_effort(self):
191     """ Synchronize the schedule.
192
193     This function assumes that all subschedules have already been resolved.
194     It iterates by asking get_next_job for a new job, and submitting the
195     result of the previous job, until no more jobs are forthcoming. Then
196     what remains is the synchronous product of the schedule including
197     information about its compatibility.
198
199     It sends jobs to the 'synchronize.synchronous_product_effort' function
200     which returns the result of synchronization and the effort expended.
201
202     Effort expended on jobs that didn't succeed is still counted.
203     """
204     total_effort = self.initial_effort
205     for pair in list(itertools.combinations(self.components, 2)):
206         self.analyze(pair[0], pair[1])
207     result = None
208     choice = None
209     while True:
210         choice = self.get_next_job(result, choice)
211         if choice == None:
212             break
213         result, effort = synchronize.synchronous_product_effort(choice, self.
214     contradictions)
215         total_effort += effort
216     return (result, total_effort)
217
218 def set_strategy(self, strategy):
219     """ Set a strategy for this schedule.

```

```

219     It associates a strategy function to the analyze(A, B) function, so that
220     any calls to analyze use the associated function.
221     """
222     self.strategy = strategy # record for reuse by recursion over subteams
223     if strategy == "action sim":
224         self.analyze = self.strategy_action_similarity
225     elif strategy == "action dissim":
226         self.analyze = self.strategy_action_dissimilarity
227     elif strategy == "fifo":
228         self.analyze = self.strategy_fifo
229     elif strategy == "lifo":
230         self.analyze = self.strategy_lifo
231     elif strategy == "random":
232         self.analyze = self.strategy_random
233     elif strategy == "Q min":
234         self.analyze = self.strategy_Q_min
235     elif strategy == "Q max":
236         self.analyze = self.strategy_Q_max
237     elif strategy == "all at once":
238         self.analyze = self.strategy_all_at_once
239     else:
240         print "ERROR: unknown strategy."
241         exit(1)
242
243     def strategy_random(self, A1, A2):
244         """ Combinations receive a random fitness value.
245         """
246         self.combinations[(A1, A2)] = random.random()
247
248     def strategy_action_similarity(self, A1, A2):
249         """ Pairs receive higher scores if they have more actions in common.
250         """
251         AA1 = self.components[A1]
252         AA2 = self.components[A2]
253         total = len(set().union(AA1.get_all_actions(), AA2.get_all_actions()))
254         shared = len(AA1.get_all_actions().intersection(AA2.get_all_actions()))
255         self.combinations[(A1, A2)] = shared/float(total)
256
257     def strategy_action_dissimilarity(self, A1, A2):
258         """ Pairs receive higher scores if they have less actions in common.
259         """
260         AA1 = self.components[A1]
261         AA2 = self.components[A2]
262         total = len(set().union(AA1.get_all_actions(), AA2.get_all_actions()))
263         shared = len(AA1.get_all_actions().intersection(AA2.get_all_actions()))
264         self.combinations[(A1, A2)] = -(shared/float(total))
265
266     def strategy_fifo(self, A1, A2):
267         """ First In First Out synchronization of pairs.
268
269         Results of synchronization end up at the back of the queue, so first
270         all base components are used, then all first-generation product
271         components and so forth.
272         """
273         B1 = None
274         B2 = None
275         obsolete = set()
276         if len(self.combinations) < 1:
277             # select the leftmost component that still exists
278             for item in range(len(self.ordered_components)):
279                 if self.ordered_components[item] in self.components.keys():
280                     B1 = self.ordered_components[item]
281                     break
282             else:
283                 obsolete.add(item)
284             # select the second to leftmost component that still exists
285             for item in range(len(self.ordered_components)):
286                 if self.ordered_components[item] in self.components.keys():
287                     if self.ordered_components[item] != B1:
288                         B2 = self.ordered_components[item]
289                     break

```

```

291         else:
292             obsolete.add(item)
293         # get rid of leftmost components that no longer exist
294         for item in sorted(obsolete, reverse=True):
295             del self.ordered_components[item]
296         self.combinations[(B1, B2)] = 1
297
298 def strategy_lifo(self, A1, A2):
299     """ Last In First Out synchronization of pairs.
300
301     Effectively, the result of the latest synchronization is immediately
302     used to synchronize again, with the rightmost base component.
303     """
304     B1 = None
305     B2 = None
306     obsolete = set()
307     if len(self.combinations) < 1:
308         # select rightmost component that still exists
309         for item in reversed(range(len(self.ordered_components))):
310             if self.ordered_components[item] in self.components.keys():
311                 B1 = self.ordered_components[item]
312                 break
313         else:
314             obsolete.add(item)
315         # select second to rightmost component that still exists
316         for item in reversed(range(len(self.ordered_components))):
317             if self.ordered_components[item] in self.components.keys():
318                 if self.ordered_components[item] != B1:
319                     B2 = self.ordered_components[item]
320                     break
321         else:
322             obsolete.add(item)
323         # remove rightmost components that no longer exist
324         for item in sorted(obsolete, reverse=True):
325             del self.ordered_components[item]
326         self.combinations[(B1, B2)] = 1
327
328 def strategy_Q_min(self, A1, A2):
329     """ Take the pair of components with the smallest combined state space
330     """
331     self.combinations[(A1, A2)] = -(len(self.components[A1].states) * \
332                                     len(self.components[A2].states))
333
334 def strategy_Q_max(self, A1, A2):
335     """ Take the pair of components with the largest combined state space
336     """
337     self.combinations[(A1, A2)] = len(self.components[A1].states) * \
338                                     len(self.components[A2].states)
339
340 def strategy_all_at_once(self, A1, A2):
341     """ Do not make pairs at all.
342
343     As a result, all components will be synchronized as if all viable pairs
344     had already been eliminated, i.e. immediately try too synchronize all
345     components together at once.
346     """
347     if len(self.components) == 2:
348         self.combinations[(list(self.components)[0],
349                             list(self.components)[1])] = 1

```

../source_code/scheduler.py

```

1 import sys
import scheduler
3
4 schedule = None
5 strategy = "fifo"
6 complete = False
7 for elem in sys.argv[1:]:
8     if elem.startswith("strategy="):
9         strategy = elem.replace("strategy=", "")
10    if elem.startswith("complete="):
11        if elem.replace("complete=", "") == "True" or elem.replace("complete=", "") == "true":
12            complete = True
13    if elem.startswith("schedule="):
14        schedule = elem.replace(".scd", "").replace("schedule=", "")
15
16 print "SCHEDULE: ", schedule
17 print "STRATEGY: ", strategy
18 print "COMPLETE: ", complete
19
20 if schedule == None:
21     print "ERROR: schedule parameter required. Usage: \"schedule=[schedulingname]\""
22     exit(1)
23
24 worker = scheduler.Scheduler(schedule+".scd", strategy=strategy, find_all_contradictions=
25     complete)
26 result, effort = worker.run_effort()
27 result.save('X'+schedule)
28
29 #result.screendump()
30
31 print "Effort: ", effort
32 if result.is_automaton_compatible():
33     print "This is a compatible system."
34 else:
35     print "This is NOT a compatible system."

```

../source_code/schedule_commandline.py

```

import scheduler
import graph

schedules = ['workplace', 'coffeeTime', 'tokenRing', 'tokenRing2', 'settlement',
            'wedding', 'passing']
strategies = ['action sim', 'action dissim', 'fifo', 'lifo',
            'Q min', 'Q max', 'all at once']

stats = []
find_all_contradictions=False

for schedule in schedules:
    row = [schedule]
    print "\n", schedule.upper()
    for strategy in strategies:
        worker = scheduler.Scheduler(schedule+".scd", strategy=strategy,
        find_all_contradictions=find_all_contradictions)
        result, effort = worker.run_effort(find_all_contradictions=find_all_contradictions)
        result.save('X'+schedule)
        if strategy == strategies[0]:
            row.append(result.is_automaton_compatible())
#         graph.makeGraph(result)
#         result.screendump()
            row.append(effort)
        stats.append(row)

offset = 15

print "".rjust(offset), "Compatible?".rjust(offset),
for strategy in strategies: print strategy.rjust(offset),
print
for row in stats:
    for cell in row:
        print str(cell).rjust(offset),
    print

# output as tex table
with open('experiment.tex', 'w') as f:
    f.write("\begin{myTable}{| X[1.5] | X[1.5] | }")
    for x in range(len(stats[0])-2): f.write("X[r] | ")
    f.write("}\n\\hline\n")
    f.write("&".rjust(offset))
    f.write("\footnotesize Compatible? &".rjust(offset))
    for s in range(len(strategies)-1):
        f.write(("\\footnotesize "+strategies[s]+" &").rjust(offset))
    f.write(("\\footnotesize "+strategies[-1]).rjust(offset))
    f.write("\\\\\\n\\hline\n")
    for row in stats:
        for cell in range(len(row)-1):
            f.write(("\\footnotesize "+str(row[cell])+" &").rjust(offset))
        f.write(("\\footnotesize "+str(row[-1]).rjust(offset))
        f.write("\\\\\\n\\hline\n")
    f.write("\end{myTable}")

# output as CSV
with open('experiment.csv', 'w') as f:
    f.write(", Compatible?,")
    for s in strategies:
        f.write(s + ",")
    f.write("\n")
    for row in stats:
        for cell in row:
            f.write(str(cell) + ",")
        f.write("\n")

```

../source_code/experiment.py

B Automata test sets

B.1 Boss/Employee

```
1 .model boss
2 .inputs work_harder
3 .outputs get_raise
4 .internal
5 .state graph
6     unhappy work_harder happy
7     happy get_raise happy
8 .marking {unhappy}
9 .end
```

../source_code/cases/boss.out

```
1 .model employee
2 .inputs get_raise
3 .outputs work_harder
4 .internal
5 .state graph
6     unhappy get_raise happy
7     happy work_harder happy
8 .marking {unhappy}
9 .end
```

../source_code/cases/employee.out

B.2 Coffee

```
1 .model drinker
2 .inputs serve_coffee
3 .outputs press_button
4 .internal drink_coffee
5 .state graph
6     thirsty press_button waiting
7     waiting serve_coffee happy
8     happy drink_coffee thirsty
9 .marking {thirsty}
10 .end
```

../source_code/cases/drinker.out

```
1 .model machine
2 .inputs press_button
3 .outputs serve_coffee
4 .internal clean_filters
5 .state graph
6     ready press_button brewing
7     brewing serve_coffee dirty
8     dirty clean_filters ready
9 .marking {ready}
10 .end
```

../source_code/cases/machine.out

B.3 Settlement

```
1 .model investor1
2 .inputs placementConfirmation cancellationConfirmation
3 .outputs placementOrder1 cancellationOrder1
4 .internal
5 .state graph
6     start1 placementOrder1 waiting1
7     waiting1 cancellationOrder1 waiting1Cancellation
8     waiting1Cancellation cancellationConfirmation finished1
9     waiting1 placementConfirmation finished1
10    start1 cancellationConfirmation finished1
11    waiting1 cancellationConfirmation finished1
12 .marking {start1}
13 .end
```

../source_code/cases/investor1.out

```
1 .model investor2
2 .inputs placementConfirmation cancellationConfirmation
3 .outputs placementOrder2 cancellationOrder2
4 .internal acceptPlacementConfirmation
5 .state graph
6     start2 placementOrder2 waiting2
7     waiting2 cancellationOrder2 waiting2Cancellation
8     waiting2Cancellation cancellationConfirmation finished2
9     waiting2 placementConfirmation wrong
10    wrong cancellationOrder2 waiting2Cancellation
11    start2 cancellationConfirmation finished2
12    waiting2 cancellationConfirmation finished2
13    wrong acceptPlacementConfirmation finished2
14 .marking {start2}
15 .end
```

../source_code/cases/investor2.out

```
1 .model custodian
2 .inputs placementOrder1 placementOrder2 cancellationOrder1 cancellationOrder2
3 .outputs placementConfirmation cancellationConfirmation
4 .internal
5 .state graph
6     No_No placementOrder1 Yes_No
7     No_No placementOrder2 No_Yes
8     No_Yes placementOrder1 Yes_Yes
9     Yes_No placementOrder2 Yes_Yes
10    Yes_Yes placementConfirmation finished
11    Yes_Yes cancellationOrder1 cancelling
12    Yes_Yes cancellationOrder2 cancelling
13    Yes_No cancellationOrder1 cancelling
14    No_Yes cancellationOrder2 cancelling
15    cancelling cancellationOrder1 cancelling
16    cancelling cancellationOrder2 cancelling
17    cancelling cancellationConfirmation finished
18    cancelling placementOrder1 cancelling
19    cancelling placementOrder2 cancelling
20 .marking {No_No}
21 .end
```

../source_code/cases/custodian.out

B.4 Token Ring

```
1 .model user1
2 .inputs grant1 tick
3 .outputs request1
4 .internal enter1 go1
5 .state graph
6     offline enter1 online
7     offline tick offline
8     online tick online
9     waiting tick online
10    online request1 waiting
11    waiting grant1 online
12    online go1 offline
13 .marking {online}
14 .end
```

../source_code/cases/user1.out

```
1 .model station1a
2 .inputs token1 tick request1
3 .outputs grant1
4 .internal
5 .state graph
6     waiting tick waiting
7     waiting token1 haveToken
8     haveToken request1 willGrant
9     haveToken tick waiting
10    willGrant grant1 haveGranted
11    haveGranted tick waiting
12 .marking {waiting}
13 .end
```

../source_code/cases/station1a.out

```
1 .model station1b
2 .inputs tick token1
3 .outputs token2
4 .internal
5 .state graph
6     waiting tick waiting
7     waiting token1 haveToken
8     haveToken tick mustPass
9     mustPass token2 waiting
10 .marking {waiting}
11 .end
```

../source_code/cases/station1b.out

```
1 .model user2
2 .inputs grant2 tick
3 .outputs request2
4 .internal enter2 go2
5 .state graph
6     offline enter2 online
7     offline tick offline
8     online tick online
9     waiting tick online
10    online request2 waiting
11    waiting grant2 online
12    online go2 offline
13 .marking {online}
14 .end
```

../source_code/cases/user2.out

```
1 .model station2a
2 .inputs token2 tick request2
3 .outputs grant2
4 .internal
5 .state graph
```

```

6   waiting tick waiting
   waiting token2 haveToken
8   haveToken request2 willGrant
   haveToken tick waiting
10  willGrant grant2 haveGranted
   haveGranted tick waiting
12 .marking {waiting}
   .end

```

../source_code/cases/station2a.out

```

1 .model station2b
  .inputs tick token2
3  .outputs token3
  .internal
5  .state graph
   waiting tick waiting
7   waiting token2 haveToken
   haveToken tick mustPass
9   mustPass token3 waiting
  .marking {waiting}
11 .end

```

../source_code/cases/station2b.out

```

1 .model user3
  .inputs grant3 tick
3  .outputs request3
  .internal enter3 go3
5  .state graph
   offline enter3 online
7   offline tick offline
   online tick online
9   waiting tick online
11  online request3 waiting
   waiting grant3 online
   online go3 offline
13 .marking {online}
   .end

```

../source_code/cases/user3.out

```

1 .model station3a
2  .inputs token3 tick request3
  .outputs grant3
4  .internal
  .state graph
6   waiting tick waiting
   waiting token3 haveToken
8   haveToken request3 willGrant
   haveToken tick waiting
10  willGrant grant3 haveGranted
   haveGranted tick waiting
12 .marking {haveToken}
   .end

```

../source_code/cases/station3a.out

```

1 .model station3b
  .inputs tick token3
3  .outputs token1
  .internal
5  .state graph
   waiting tick waiting
7   waiting token3 haveToken
   haveToken tick mustPass
9   mustPass token1 waiting
  .marking {haveToken}
11 .end

```

../source_code/cases/station3b.out

B.5 Wedding

```
1 .model bride
2 .inputs do_you_groom groom_I.do groom_I.dont do_you_bride poll_objections objection
   pronounce_marriage
3 .outputs bride_enters bride_I.do bride_I.dont kiss
4 .internal
5 .state graph
   outside bride_enters bw1
7   bw1 do_you_groom bw2
   bw2 groom_I.do bw3
9   bw2 groom_I.dont problem
   bw3 do_you_bride ms1
11  ms1 bride_I.do bw4
   ms1 bride_I.dont problem
13  bw4 poll_objections bw5
   bw5 objection problem
15  bw5 pronounce_marriage ms2
   ms2 kiss finished
17 .marking {outside}
   .end
```

../source_code/cases/bride.out

```
1 .model groom
2 .inputs bride_enters do_you_groom do_you_bride bride_I.do bride_I.dont poll_objections
   objection pronounce_marriage
3 .outputs groom_I.do groom_I.dont kiss
4 .internal hesitate
5 .state graph
6   gw1 bride_enters gw2
   gw2 do_you_groom ms1
8   ms1 hesitate ms1
   ms1 groom_I.do gw3
10  ms1 groom_I.dont problem
   gw3 do_you_bride gw4
12  gw4 bride_I.do gw5
   gw4 bride_I.dont problem
14  gw5 poll_objections gw6
   gw6 objection problem
16  gw6 pronounce_marriage ms2
   ms2 kiss finished
18 .marking {gw1}
   .end
```

../source_code/cases/groom.out

```
1 .model minister
2 .inputs bride_enters groom_I.do groom_I.dont bride_I.do bride_I.dont objection
3 .outputs do_you_groom do_you_bride poll_objections pronounce_marriage
4 .internal
5 .state graph
6   mw1 bride_enters ms1
7   ms1 do_you_groom mw2
   mw2 groom_I.do ms2
9   mw2 groom_I.dont problem
   ms2 do_you_bride mw3
11  mw3 bride_I.do ms3
   mw3 bride_I.dont problem
13  ms3 poll_objections mw4
   mw4 objection problem
15  mw4 pronounce_marriage finished
17 .marking {mw1}
   .end
```

../source_code/cases/minister.out

B.6 Passing the Ball

```
1 .model A1
2 .inputs a5
3 .outputs a1
4 .internal
5 .state graph
6     ball a1 noball
7     noball a5 ball
8 .marking {ball}
9 .end
```

../source_code/cases/A1.out

```
1 .model A2
2 .inputs a1
3 .outputs a2
4 .internal
5 .state graph
6     noball a1 ball
7     ball a2 noball
8 .marking {noball}
9 .end
```

../source_code/cases/A2.out

```
1 .model A3
2 .inputs a2
3 .outputs a3
4 .internal
5 .state graph
6     noball a2 ball
7     ball a3 noball
8 .marking {noball}
9 .end
```

../source_code/cases/A3.out

```
1 .model A4
2 .inputs a3
3 .outputs a4
4 .internal
5 .state graph
6     noball a3 ball
7     ball a4 noball
8 .marking {noball}
9 .end
```

../source_code/cases/A4.out

```
1 .model A5
2 .inputs a4
3 .outputs a5
4 .internal
5 .state graph
6     noball a4 ball
7     ball a5 noball
8 .marking {noball}
9 .end
```

../source_code/cases/A5.out