

Alternating-Time Stream Logic for Multi-Agent Systems

Sascha Klüppelholz, Christel Baier*

Technische Universität Dresden, Institut für Theoretische Informatik, Germany
{klueppel,baier}@tcs.inf.tu-dresden.de

Abstract. Constraint automata have been introduced to provide a compositional, operational semantics for the exogenous coordination language Reo, but they can also serve interface specification for components and an operational model for other coordination languages. Constraint automata have been used as basis for equivalence checking and model checking temporal logical properties. The main contribution of this paper is to reason about the local view and interaction and cooperation facilities of individual components or coalitions of components by means of a multi-player semantics for constraint automata. We introduce a temporal logic framework that combines classical features of alternating-time logic (*ATL*) for concurrent games with special operators to specify the observable data flow at the I/O-ports of components. Since constraint automata support any kind of synchronous and asynchronous peer-to-peer communication, the resulting game structure is non-standard and requires a series of nontrivial adaptations of the *ATL* model checking algorithm.

1 Introduction

In the last decade several models and specification languages for formal reasoning about the middle-ware layer of software have been developed. Such coordination models consist of ad-hoc libraries of functions providing higher-level inter-process communication support in parallel and especially distributed applications. They aim at a clean separation between individual software components and their interactions within their overall software organization. Our approach is inspired by the coordination language Reo [2], which provides the *glue-code* to coordinate components in an exogenous manner. In this paper we use constraint automata, which have been introduced as an operational semantics for Reo [6]. Constraint automata provide a specification formalism for both, the glue-code (e.g. given as a (Reo) network, or another (channel-based) coordination mechanism) and the behavioral interfaces of components, and can serve to formalize the overall behavior of the composite system. Constraint automata capture any kind of synchronous and asynchronous peer-to-peer communication including data-dependencies of I/O-operations. The syntax of constraint automata is similar to ordinary labeled transition systems and related models, such as timed port automata [15], I/O-automata [20], and interface automata [10]. The differences are mainly based on the fact that constraint automata support any kind of channel-based communication. An extensive discussion on the differences and similarities can be found in [6].

* The authors are supported by the DFG-NWO project SYANCO and the EU project CREDO.

The purpose of this paper is to provide a multi-agent semantics for constraint automata and an alternating-time temporal logic to specify and verify the components considered as individual players of a multi-agent game. The connected components are the individual players and the network sets up the rules how those players interact with each other. The glue-code might be seen as a complex set of social laws [13, 24] the players have to stick to. Constraint automata, interpreted as multi-player game structures, are a special type of concurrent games. The specific challenges of an alternating time approach are caused by the very special mixture of asynchrony and synchrony, mutual dependencies of I/O-operations and data-dependencies. In each state, several concurrent I/O-operations can be enabled, but only some of them might be available once a player refuses some synchronization or declares conditions on the data values accepted on his input ports or on his pending write operations. Furthermore, constraint automata can contain some internal nondeterminism, which yields a rather complex and nonstandard concurrent game structure. We are not aware of any other paper that treats alternating-time aspects for such concurrent games, where the enabledness and also the effect of a concurrent I/O-operation highly depends on the choices of the other players. Our approach allows us to check whether or not some coalition of agents has a strategy to achieve a common goal, no matter how the opponents behave, or which internal nondeterministic choices were made. In contrast to standard concurrent games, see e.g. [1, 9], in our approach a coalition's strategy may select sets of I/O-operations or even refuse any I/O-operations.

For specifying and analyzing the local views and interaction possibilities of (coalitions of) agents, we introduce an alternating-time logic, called alternating-time stream logic (*ASL*). The logic *ASL* is a *CTL*-like branching-time logic which combines the features of standard *ATL* [1] with the operators of *BTSL* [18]. The logic *BTSL* has been introduced as a temporal logic for reasoning about (Reo) networks. Beside the standard modalities of *CTL* [8], *BTSL* supports the specification of the observable data flow at the I/O-ports of channels and components by means of regular expressions. The focus of *ATL* is to ask for the existence (and absence) of a coalition's strategy to achieve (avoid respectively) a specific temporal goal once the behavior for each of the components is specified.

For a simple example, we regard a ticket vending machine, which consists of a number of components (e.g. *I/O-device*, *clock*, *destination*, *price*, *payment*, and *printer*). The exact behavior of the components might be specified in terms of constraint automata. *ASL* can be used to formalize the property stating that the user (possibly together with some other component like the *clock*) can find a way to trick the other players and get a ticket without paying. A dual *ASL* property would state that no matter what strategy the opponents use, the coalition of opponents will not have a chance to avoid that sending the *cancel* signal always resets all components to their initial configuration.

As a first step we assume *perfect recall* on the systems history and *perfect information* on the global state of the system. This interpretation of constraint automata as a multi-player game is consistent with the standard semantics of *ATL* and adequate if the strategies are viewed as a central control that is aware of all activities in the system.

Our approach differs from other *ATL*-like approaches for concurrent multi-player games in various aspects. First, our nonstandard game structure (see explanations above) requires a revised notion of strategies for (coalitions of) components. Second, since com-

ponents may refuse any further interaction from some moment on, the concept of finite runs and fairness plays a crucial role in the logic *ASL*. To reason about liveness properties we need an adaption of the standard notion of strong (process) fairness. Our notion of fairness is not a requirement for strategies, but formalizes the ability of certain strategies of a component C to enforce infinite data flow at the I/O-ports of C . Third, *ASL* provides special operators to reason about the observable data flow at the I/O-ports of the components and the nodes of the given network. To the best of our knowledge, such operators have not yet been investigated in the context of alternating-time game models.

Organization. Section 2 gives a brief introduction to constraint automata. In section 3 we provide the multi-player semantics for constraint automata and introduce the notion of a strategy and its runs. Section 4 introduces the temporal logic *ASL* and presents corresponding model checking algorithms. Section 5 introduces fairness assumptions to *ASL* model checking, before section 6 concludes the paper. An extended technical report including the proofs and other technical material is available on the web [19].

2 Constraint Automata (CA)

This section summarizes the main concepts of CA. We slightly depart from the syntax of CA as introduced in [6] and deal with transitions $q \xrightarrow{c} p$, where c is a *concurrent I/O-operation*, i.e., c consists of a (possibly empty) node-set $N \subseteq \mathcal{N}$ together with data items for each $A \in N$ that are written or received at node A . In the moment where c is executed there is no data flow at the nodes $A \in \mathcal{N} \setminus N$.

Concurrent I/O-operations and I/O-streams. Let \mathcal{N} be a finite, nonempty set of nodes. We define a concurrent I/O-operation as a function $c : \mathcal{N} \rightarrow \text{Data} \cup \{\perp\}$, where the symbol \perp means “undefined”. We write $\text{Nodes}(c)$ for the set of nodes $A \in \mathcal{N}$ such that $c(A) \in \text{Data}$, where Data is the data domain. For technical reasons, we also allow the *empty* concurrent I/O-operation c_\emptyset with $\text{Nodes}(c_\emptyset) = \emptyset$. It represents any internal step of some component or a non-observable step, where data flow appears at some hidden (invisible) nodes only. We refer to CIO as the set of all concurrent I/O-operations (including c_\emptyset). As we suppose \mathcal{N} and Data to be finite, the set CIO of concurrent I/O-operations is finite as well. When reasoning about the data flow in a Reo network we will also need a special symbol \surd that indicates that data flow has stopped. CIO_\surd stands for $\text{CIO} \cup \{\surd\}$.

Definition 1 (Constraint automata [6]). A constraint automaton (CA) is a tuple

$$\mathcal{A} = \langle Q, \mathcal{N}, \longrightarrow, Q_0, AP, L \rangle,$$

where Q is a finite and nonempty set of states, \mathcal{N} a finite set of nodes, \longrightarrow is a subset of $Q \times \text{CIO} \times Q$ called the transition relation of \mathcal{A} , $Q_0 \subseteq Q$ a nonempty set of initial states, AP a finite set of atomic propositions, and $L : Q \rightarrow 2^{AP}$ a labeling function. We write $q \xrightarrow{c} p$ instead of $(q, c, p) \in \longrightarrow$. Furthermore, we define the set of all I/O-operations enabled in q as $\text{CIO}(q) \stackrel{\text{def}}{=} \{c \in \text{CIO} : q \xrightarrow{c} p \text{ for some } p \in Q\}$.


Intuitively, the nodes correspond to the I/O-ports of the components. For the pictures of CAs we shall use symbolic representations of the transition relation by combining transitions with the same starting and target state. For this purpose, we use I/O-constraints, i.e., propositional formulas in positive normal form that stand for sets of concurrent I/O-operations. The I/O-constraints may impose conditions on the nodes that may or may not be involved and on the data items written on or read from them.

I/O-constraints (IOC). The abstract syntax of I/O-constraints is given by the grammar:

$$ioc ::= tt \mid ff \mid A \mid \neg A \mid (d_{A_1}, \dots, d_{A_k}) \in D \mid ioc_1 \wedge ioc_2 \mid ioc_1 \vee ioc_2$$

where $A \in \mathcal{N}$, A_1, \dots, A_k are pairwise distinct nodes in \mathcal{N} and $D \subseteq Data^k$. The meaning of an I/O-constraint ioc is a subset $CIO(ioc)$ of CIO defined in the obvious way. We often use simplified notations for the IOCs of the form $(d_{A_1}, \dots, d_{A_k}) \in D$. E.g., the notation $d_A = d_B$ is a shorthand for $(d_A, d_B) \in \{(d_1, d_2) \in Data^2 : d_1 = d_2\}$, while $A \wedge (d_B \in P)$ stands for the set $\{c \in CIO : \{A, B\} \subseteq Nodes(c) \wedge c(B) \in P\}$.

Example 1 (CA). The following two CAs realize possible implementations for the *destination* component with node set $\mathcal{N}_D = \{E, I, K, O, R\}$ and *price* component with node set $\mathcal{N}_P = \{F, J, M, T, V, W\}$ of the ticket vending machine. Both components are allowed to operate if and only if some data flow occurs on their synchronization ports E and F respectively. In the picture below we use a parameterized representation for states.



The *destination* component simultaneously reads some destination id (variable i) on its input port I and writes the destination string (variable $dest_i$) to the *I/O-device* using port R and its output port O. If the destination number given is too large, i.e., it exceeds a certain maximum μ , the *I/O-device* gets a message that the selected destination is unknown. The *price* component receives two integer values at its input ports T and V for the destination (variable $dest_i$) and ticket type (variable $type_j$) and sends the corresponding price (variable $p_{i,j}$) first to the *I/O-device* using port W and in a second step to the *payment* component using port M. Both automata accept a *cancel* signal at any state and reset to their initial configuration.

Terminal states. A state q is called *terminal* if data flow may stop in state q . This is the case if all enabled concurrent I/O-operations require some activity of a component connected to a sink or source node. Formally, state q is said to be terminal if for all concurrent I/O-operations c that are enabled in state q , the node-set $Nodes(c)$ is non-empty. Stated differently, state q is terminal iff $c_\emptyset \notin CIO(q)$. Note that data flow does not need to stop in terminal states. Instead data flow continues if there is an enabled concurrent I/O-operation c where the involved components agree on interacting with each other by means of performing the write and read operation specified by c . For each

non-terminal node q , an invisible transition is enabled, i.e., we have $c_\emptyset \in \text{CIO}(q)$. This I/O-operation does not require any interaction with the components that are connected to the sink and source nodes and will fire, unless another transition is taken.

Executions, completeness, paths, I/O-streams. An *execution* in \mathcal{A} is a finite or infinite sequence built by instances of consecutive transitions: $\eta = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots$

where $q_0, q_1, \dots \in Q$, $c_1, c_2, \dots \in \text{CIO}$, and $q_i \xrightarrow{c_{i+1}} q_{i+1}$ for all $i \geq 0$.

To reason about “maximal” behaviors of CAs we introduce the notions of complete executions and paths. An execution is said to be *complete* if it is either infinite or it is finite and ends in a terminal state. A *path* of \mathcal{A} is either an infinite execution or arises from a finite complete execution by adding a special transition symbol \surd to denote termination. More precisely, the finite paths have the form $\pi = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n \xrightarrow{\surd} q_n$ where q_n is terminal. In the sequel, we shall use the symbol η for executions and the symbol π to range over paths. We write $\text{Paths}(q)$ to denote the set of all paths starting in q and $\text{Exec}_{\text{fin}}(q)$ for the set of all finite executions starting in q . The length $|\pi|$ of a path π is the total number of transitions taken in π (including the pseudo-transition with label \surd). Thus, the length of an infinite path is ∞ , while the length of a finite path π as above is $n + 1$. Let $\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots$ be a path and $0 \leq n < |\pi|$. Then $\pi \downarrow n$ denotes the prefix of path π with length n , i.e., $\pi \downarrow n \stackrel{\text{def}}{=} q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n$ is an execution, while for $n = |\pi|$ we have that $\pi \downarrow n = \pi$ is still a path. The *I/O-stream* $\text{ios}(\eta)$ of a finite execution η is the word over CIO that is obtained by taking the projection to the labels of the transitions. That is, if $\eta = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n$ then $\text{ios}(\eta) \stackrel{\text{def}}{=} c_1 \dots c_n$. Similarly, the associated I/O-stream for a finite path $\pi = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n \xrightarrow{\surd} q_n$ is defined by $\text{ios}(\pi) \stackrel{\text{def}}{=} c_1 \dots c_n \surd$. Let $\text{IOS} = \text{CIO}^* \cup \text{CIO}^* \surd$ denote the set of all I/O-streams.

3 Constraint Automata as Multi-Player Games

In this section we introduce a game-theoretical interpretation for CA. The players are the individual components using (a)synchronous peer-to-peer communication. Each of the players has control over his I/O-behavior at its interface nodes. A player might refuse some or even any synchronization operation with other players. As in ordinary *ATL*, players might build arbitrary coalitions to achieve a certain common goal including a specific temporal behavior. A coalition of players induces a set of controllable nodes $N \subseteq \mathcal{N}$, the union of all controllable coalition nodes, for which the players might try to develop a common strategy to achieve their objective(s). Intuitively, an N -strategy takes the history of the system formalized by a finite execution as input, (i.e., we suppose here perfect recall) and declare the conditions under which the N -agents (members of the coalition) are willing to cooperate with each other and their opponents. For instance, an N -strategy might offer to write data value 0 at a source node $A \in \mathcal{N}$, but refuse to write data value 1. The general notion of N -strategies also permits to couple such constraints for the offered I/O-operations at the N -nodes with conditions on the

IOCs at the nodes in $\mathcal{N} \setminus \mathbf{N}$. Furthermore, an \mathbf{N} -strategy might suggest the \mathbf{N} -agents to refuse any participation in concurrent I/O-operations. The special symbol *stop* will be used for this purpose.

Definition 2 (Strategy). Let \mathcal{A} be a CA as before and let \mathbf{N} be a node-set such that $\mathbf{N} \subseteq \mathcal{N}$. An \mathbf{N} -strategy is a function

$$\mathfrak{S} : \text{Exec}_{\text{fin}}(\mathcal{A}) \rightarrow 2^{\text{CIO} \cup \{\text{stop}\}},$$

assigning to any finite execution η a set $\mathfrak{S}(\eta)$ consisting of I/O-operations $c \in \text{CIO}$ or the special symbol *stop* such that if $c \in \text{CIO}$ and $\text{Nodes}(c) \cap \mathbf{N} = \emptyset$ then $c \in \mathfrak{S}(\eta)$.

The intuitive meaning of the condition required for an \mathbf{N} -strategy asserts that the \mathbf{N} -nodes are not in the position to refuse an I/O-operation c where none of the \mathbf{N} -nodes is involved. In particular, invisible I/O-operations (i.e., concurrent I/O-operations with the empty node-set) cannot be ruled out by an \mathbf{N} -strategy. A possible refinement for the notion of a strategy would be to allow components to restrict their write operations only and not to cut down any input provided at their boundary nodes. Given an \mathbf{N} -strategy \mathfrak{S} , the \mathfrak{S} -paths are those paths in \mathcal{A} , where each of the I/O-operations performed is accepted at any time by the \mathbf{N} -nodes and their strategy \mathfrak{S} .

Notation 3 (\mathfrak{S} -executions, \mathfrak{S} -completeness, \mathfrak{S} -paths) Let \mathfrak{S} be an \mathbf{N} -strategy and $\eta = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots$ a finite or infinite execution in \mathcal{A} . Then, η is called a \mathfrak{S} -*execution* if for any position $i \in \mathbb{N}$ with $i < |\eta|$ we have $c_{i+1} \in \mathfrak{S}(\eta \downarrow i)$. A finite \mathfrak{S} -execution η of length n is called \mathfrak{S} -*complete* if the last state q_n of η is terminal and at least one of the following two conditions holds:

- (i) $\text{stop} \in \mathfrak{S}(\eta)$ or
- (ii) there is no $c \in \text{CIO}(q_n) \cap \mathfrak{S}(\eta \downarrow n)$ such that $\text{Nodes}(c) \subseteq \mathbf{N}$

The first condition indicates that refusing any data flow on the \mathbf{N} -nodes is a potential behavior under strategy \mathfrak{S} , while the second indicates the possibility for the opponents to do the same on their part (i.e. refusing any synchronization on the $\mathcal{N} \setminus \mathbf{N}$ nodes). Furthermore, each infinite \mathfrak{S} -execution is said to be \mathfrak{S} -complete. A \mathfrak{S} -*path* denotes any infinite \mathfrak{S} -execution or any finite path $\pi = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n \xrightarrow{\vee} q_n$, where $\pi \downarrow n$ is a \mathfrak{S} -complete \mathfrak{S} -execution. We write $\text{Paths}(q, \mathfrak{S})$ to denote all \mathfrak{S} -paths starting in q . Similarly, $\text{Exec}_{\text{fin}}(q, \mathfrak{S})$ denotes the set of all finite \mathfrak{S} -executions from q .

Notation 4 (Memoryless, finite-memory strategies) An \mathbf{N} -strategy \mathfrak{S} is called *memoryless* if $\mathfrak{S}(\eta) = \mathfrak{S}(\eta')$ for all finite executions η and η' that end in the same state. Memoryless strategies can be seen as functions $\mathfrak{S} : \mathcal{Q} \rightarrow 2^{\text{CIO} \cup \{\text{stop}\}}$. Obviously, memoryless strategies are special instances of *finite-memory* strategies, i.e., strategies that make their decisions on the basis of a finite automaton rather than the full history.

4 Alternating-Time Stream Logic (ASL)

To reason about the components from a game-theoretic point of view, we introduce *alternating-time stream logic (ASL)* which is inspired by alternating-time temporal logic (*ATL*) [1]. *ASL* extends *BTSL* [18] to state the possibility for components to cooperate in such way that a certain temporal property or property on the observable data flow holds. *ASL* is a branching time logic with state and path formulas. The state formula fragment is as in *ATL*, but adapted to the CA framework where the alternating-time quantifiers range over the strategies of certain node-sets. Intuitively, these node-sets stand for the interface nodes of one or more components. The existential quantifier \mathbb{E}_N is used to indicate that the components with sink and source nodes in N have a strategy ensuring a certain condition, no matter how the other components connected to the nodes in $N \setminus N$ behave. The universal quantifier \mathbb{A}_N is dual and serves to state that the components providing the write and read actions at the N -nodes cannot avoid that a certain condition holds. The syntax of the *ASL* path formulas is the same as in *BTSL* and uses the standard until- and release operator, but replaces the standard next modality \bigcirc with special operators $\langle\langle\alpha\rangle\rangle$ and $\llbracket\alpha\rrbracket$ to impose conditions on the I/O-streams of finite executions. In path formulas of the type $\langle\langle\alpha\rangle\rangle\Phi$ or $\llbracket\alpha\rrbracket\Phi$, the formula Φ is a state formula while α is a regular expression that stands for a regular language over the alphabet CIO_{\surd} . This type of formulas is inspired by propositional dynamic logic [12], extended temporal logic [23], and timed scheduled data stream logic [3].

4.1 Syntax and Standard Semantics of ASL

In the sequel, we assume a fixed, non-empty and finite node-set \mathcal{N} . Furthermore, let AP be non-empty and finite set of atomic propositions, which can be viewed as conditions on the states of the automaton. In case of the CA modeling a FIFO-channel an atomic proposition might state that all buffer cells are empty or that the first buffer cell contains a value d in some set $P \subseteq \text{Data}$.

Regular I/O-stream expressions. The abstract syntax of regular I/O-stream expressions, briefly called stream expressions, is given by the following grammar:

$$\alpha ::= ioc \mid \surd \mid \alpha^* \mid \alpha_1; \alpha_2 \mid \alpha_1 \cup \alpha_2$$

where ioc ranges over all IOCs. Any stream expression represents a regular set of I/O-streams. The formal definition of the regular languages $IOS(\alpha) \subseteq IOS$ is defined by structural induction. $IOS(ioc)$ is the set consisting of the I/O-streams of length 1 given by ioc , i.e., $IOS(ioc) \stackrel{\text{def}}{=} CIO(ioc)$. Similarly, $IOS(\surd)$ is the singleton set consisting of the I/O-stream \surd . Union (\cup) has its standard meaning: $IOS(\alpha_1 \cup \alpha_2) \stackrel{\text{def}}{=} IOS(\alpha_1) \cup IOS(\alpha_2)$, while Kleene star ($*$) and concatenation ($;$) have to ensure that the special termination symbol \surd can only appear at the end of an I/O-stream:

$$IOS(\alpha^*) \stackrel{\text{def}}{=} \{\varepsilon\} \cup \bigcup_{n \geq 1} \{\sigma_1 \dots \sigma_n : \sigma_i \in IOS(\alpha) \cap \text{CIO}^*, i = 1, \dots, n-1, \sigma_n \in IOS(\alpha)\}$$

$$IOS(\alpha_1; \alpha_2) \stackrel{\text{def}}{=} \{\sigma_1 \surd : \sigma_1 \surd \in IOS(\alpha_1)\} \cup \{\sigma_1 \sigma_2 : \sigma_1 \in IOS(\alpha_1) \cap \text{CIO}^*, \sigma_2 \in IOS(\alpha_2)\}$$

Syntax of ASL. State-formulas (denoted by capital greek letters Φ, Ψ) and path-formulas (denoted by small greek letters φ, ψ) of ASL are built by the following grammar:

$$\begin{aligned} \Phi &::= \text{true} \mid a \mid \Phi_1 \wedge \Phi_2 \mid \neg\Phi \mid \exists\varphi \mid \mathbb{E}_N\varphi \\ \varphi &::= \langle\langle\alpha\rangle\rangle\Phi \mid \llbracket\alpha\rrbracket\Phi \mid \Phi_1 \cup \Phi_2 \mid \Phi_1 \text{R} \Phi_2 \end{aligned}$$

where $N \subseteq \mathcal{N}$, $a \in AP$ and α is a regular I/O-stream expression. The quantifier \exists in the syntax of ASL state formulas is the standard existential path quantifier of CTL and ranges over all paths, while the operator \mathbb{E}_N corresponds an existential quantification over all N-strategies. The dual operator $\mathbb{A}_N\varphi$ stating that no strategy for the nodes in N can avoid φ to hold is defined by:

$$\begin{aligned} \mathbb{A}_N\langle\langle\alpha\rangle\rangle\Phi &\stackrel{\text{def}}{=} \neg\mathbb{E}_N\llbracket\alpha\rrbracket\neg\Phi & \mathbb{A}_N(\Phi_1 \cup \Phi_2) &\stackrel{\text{def}}{=} \neg\mathbb{E}_N(\neg\Phi_1 \text{R} \neg\Phi_2) \\ \mathbb{A}_N\llbracket\alpha\rrbracket\Phi &\stackrel{\text{def}}{=} \neg\mathbb{E}_N\langle\langle\alpha\rangle\rangle\neg\Phi & \mathbb{A}_N(\Phi_1 \text{R} \Phi_2) &\stackrel{\text{def}}{=} \neg\mathbb{E}_N(\neg\Phi_1 \cup \neg\Phi_2) \end{aligned}$$

In an analogous way, the universal CTL-path quantifier \forall can be derived by duality from \exists . (Alternatively, $\forall\varphi$ can be defined by $\mathbb{E}_\emptyset\varphi$.) Other boolean connectives, like disjunction or implication, are obtained in the obvious way. In the following we shortly write $\mathbb{E}_A\varphi$ for $\mathbb{E}_{\{A\}}\varphi$ and $\mathbb{A}_A\varphi$ for $\mathbb{A}_{\{A\}}\varphi$.

ASL path formulas are interpreted over paths in a CA. The modalities \cup and R denote the ordinary until-operator and release-operator, respectively. The eventually and always operator are obtained in the usual way by $\diamond\Phi \stackrel{\text{def}}{=} (\text{true} \cup \Phi)$ and $\square\Phi \stackrel{\text{def}}{=} (\text{false} \text{R} \Phi)$. The intended meaning of $\langle\langle\alpha\rangle\rangle\Phi$ is that it holds for a path π iff π has a finite prefix generating an α -stream and Φ holds for the state reached afterwards. $\llbracket\alpha\rrbracket\Phi$ is the dual operator of $\langle\langle\alpha\rangle\rangle\Phi$ and holds for a path π iff for all finite prefixes of π generating an α -stream, formula Φ holds for the last state of the prefix. The standard *next* operator is derived from the path formula $\bigcirc\Phi \stackrel{\text{def}}{=} \langle\langle tt \rangle\rangle\Phi$, which asserts the occurrence for some (non-observable) data flow. Recall that $IOS(tt) = CIO(tt) = CIO$. Thus, $\bigcirc\Phi$ holds for all paths where the underlying execution has at least one transition and Φ holds afterwards. The presence of some observable data flow can be expressed by $\langle\langle A_1 \vee \dots \vee A_n \rangle\rangle\text{true}$, where $\mathcal{N} = \{A_1, \dots, A_n\}$. The path formula $\llbracket tt^*; \sqrt{\ } \rrbracket\text{false}$ is characteristic for the infinite paths, while $\langle\langle tt^*; \sqrt{\ } \rangle\rangle\text{true}$ holds exactly for the finite paths. The terminal states are characterized by the state formula $\exists\langle\langle \sqrt{\ } \rangle\rangle\text{true}$, while $\forall\langle\langle \sqrt{\ } \rangle\rangle\text{true}$ is satisfied in exactly those states where no concurrent I/O-operation is enabled. ASL state formulas are the same as in BTSL except for the \mathbb{E}_N -operator (and its dual).

For an intuitive example, consider a FIFO-channel with source node A and sink node B. Then the ASL state formulas $\mathbb{E}_A\square\text{empty}$, $\mathbb{E}_A\square(\text{buffer} \neq 0)$, $\mathbb{A}_B\diamond\text{empty}$ and $\mathbb{A}_B\square\text{empty}$ do hold, where $(\text{buffer} \neq 0)$ states that either the buffer is empty or contains a data value different from zero. In case of the ticket vending machine we may ask whether the user (possibly in coalition with other components) controlling three boundary nodes $N = \{C, D, P\}$ (for the *cancel* signal, data items, and payment) has a strategy to get a ticket without paying, i.e. if state formula $\mathbb{E}_{\{C,D,P\}}\langle\langle \neg\text{pay}^* \rangle\rangle\text{ticket_printed}$ holds. A dual ASL property states that all components except the user respect the *cancel* signal and reset to their initial configuration. This can be expressed by $\mathbb{A}_{\mathcal{N}\setminus N}\llbracket tt^*; C \rrbracket\text{initconf}$.

Standard semantics of ASL. Let \mathcal{A} be a CA and π a path in \mathcal{A} . The satisfaction relation \models for ASL state formulas is defined by structural induction as shown below:


$$\begin{aligned}
q &\models \text{true} \\
q &\models a && \text{iff } a \in L(q) \\
q &\models \Phi_1 \wedge \Phi_2 && \text{iff } q \models \Phi_1 \text{ and } q \models \Phi_2 \\
q &\models \neg\Phi && \text{iff } q \not\models \Phi \\
q &\models \exists\varphi && \text{iff there exists } \pi \in \text{Paths}(q) \text{ such that } \pi \models \varphi \\
q &\models \mathbb{E}_N\varphi && \text{iff there is an N-strategy } \mathfrak{S} \text{ such that:} \\
&&& \text{for all } \pi \in \text{Paths}(q, \mathfrak{S}) : \pi \models \varphi
\end{aligned}$$

The satisfaction relation \models for ASL path-formulas and the path π in \mathcal{A} as follows:

$$\begin{aligned}
\pi &\models \langle\langle\alpha\rangle\rangle\Phi && \text{iff there exists } n \in \mathbb{N} \text{ such that } 0 \leq n \leq |\pi| \text{ and} \\
&&& \text{ios}(\pi \downarrow n) \in \text{IOS}(\alpha) \text{ and } q_n \models \Phi \\
\pi &\models \llbracket\alpha\rrbracket\Phi && \text{iff for all } n \in \mathbb{N} \text{ such that } 0 \leq n \leq |\pi| \text{ we have:} \\
&&& \text{ios}(\pi \downarrow n) \in \text{IOS}(\alpha) \text{ implies } q_n \models \Phi \\
\pi &\models \Phi_1 \cup \Phi_2 && \text{iff there exists } n \in \mathbb{N} \text{ such that } 0 \leq n < |\pi| \text{ where} \\
&&& q_n \models \Phi_2 \text{ and } q_i \models \Phi_1 \text{ for } 0 \leq i < n \\
\pi &\models \Phi_1 \text{R} \Phi_2 && \text{iff at least one of the following conditions (i) or (ii) holds:} \\
&&& \text{(i) for all } n \in \mathbb{N} \text{ with } 0 \leq n < |\pi| \text{ we have: } q_n \models \Phi_2 \\
&&& \text{(ii) there exists some } n \in \mathbb{N} \text{ with } 0 \leq n \leq |\pi| \text{ such that:} \\
&&& q_n \models \Phi_1 \text{ and } q_i \models \Phi_2 \text{ for } 0 \leq i \leq n
\end{aligned}$$

Given a state q and a ASL path formula φ , an N-strategy \mathfrak{S} is called *winning* for the tuple $\langle q, \varphi \rangle$ if φ holds for all \mathfrak{S} -paths starting in q . Thus, $q \models \mathbb{E}_N\varphi$ iff there exists a winning N-strategy for $\langle q, \varphi \rangle$. For the derived operator \mathbb{A}_N we get that $q \models \mathbb{A}_N\varphi$ iff for all N-strategies \mathfrak{S} there exists $\pi \in \text{Paths}(q, \mathfrak{S})$ such that $\pi \models \varphi$, i.e. there is no winning strategy for $\langle q, \varphi \rangle$.


Example 2 (ASL state formulas). The CA with node set $\mathcal{N} = \{A, B\}$ depicted below fulfills the following state formula $\mathbb{A}_A \diamond \neg \exists \bigcirc \text{true}$, stating that an agent controlling A only cannot avoid that a terminal state q_t will eventually be reached.



The multi-player game associated with a CA and an ASL path formula is *not determined*. In fact, there are path formulas φ such that neither the N-agents have a winning strategy for φ nor does the opponents (i.e., the $\mathcal{N} \setminus N$ -agents) have a strategy to ensure that φ does not hold. The reason for this is that the internal nondeterminism can yield the possibility to generate paths where φ holds and paths where φ does not hold. In

particular, the *ASL* state formulas $\mathbb{E}_N \varphi$ and $\mathbb{A}_{N \setminus N} \varphi$ are *not* equivalent¹ and $q \models \mathbb{E}_N \varphi$ implies $q \models \mathbb{A}_{N \setminus N} \varphi$ holds for all states $q \in Q$, but not vice versa. A simple example illustrating this fact is the following CA with node-set $\mathcal{N} = \{A, B\}$.

Example 3 (Internal nondeterminism).



Assume that $a \in AP$ is an atomic proposition which holds in q_1 only, i.e. $L(q_1) = \{a\}$ and $L(q_2) = \emptyset$. Since the internal nondeterminism decides whether q_1 or q_2 will be selected as successor state of q_0 when A fires, neither A can enforce nor B can avoid that q_1 will be entered in the next step. Thus, we have $q_0 \models \mathbb{A}_B \bigcirc a$ and $q_0 \not\models \mathbb{E}_A \bigcirc a$.

4.2 ASL Model Checking

The model checking problem for *ASL* asks whether, for a given CA \mathcal{A} and *ASL* state formula Φ , all initial states q_0 of \mathcal{A} satisfy Φ . The main procedure for *ASL* model checking follows the standard approach for *CTL*-like branching-time logics [8] and recursively calculates the satisfaction sets $Sat(\Psi) = \{q \in Q : q \models \Psi\}$ for all sub-formulas Ψ of Φ . The treatment of the *BTSL*-fragment of *ASL* is the same as for *BTSL* (see [18]). The only interesting part is how to calculate $Sat(\mathbb{E}_N \varphi)$ for an *ASL* path formulas φ and node-set $N \subseteq \mathcal{N}$. The essential ingredient for this is the predecessor operator $Pre(P, N)$ which is defined as the set of all states $q \in Q$ such that the N -nodes have a strategy which guarantees to move within one step to a state in P .

Definition 5 (Predecessors). Let $P \subseteq Q$ and $N \subseteq \mathcal{N}$ a node-set. Then, $Pre(P, N)$ denotes the set of all states $q \in Q$ such that the following two conditions hold:

- (i) for all $c \in CIO(q)$ such that $Nodes(c) \cap N = \emptyset$ we have $Post[c](q) \subseteq P$
- (ii) there exists a $c \in CIO(q)$ such that $Nodes(c) \subseteq N$ and $Post[c](q) \subseteq P$

where $Post[c](q) \stackrel{\text{def}}{=} \{p \in Q : q \xrightarrow{c} p\}$.

Condition (i) is needed to ensure that no uncontrollable transition (from the view of the N -agents) leads to a state outside of P , while condition (ii) asserts the existence of at least one concurrent I/O-operation that can be enforced by the N -agents and certainly leads to a state in P . In fact we have $Pre(P, N) = \{q \in Q : q \models \mathbb{E}_N \bigcirc P\}$.

As for standard *CTL* (and *ATL*), the semantics of the until and release operator have a fixed point characterization. The set $Sat(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ is the least fixpoint, while the set $Sat(\mathbb{E}_N(\Phi_1 R \Phi_2))$ is the greatest fixpoint of the following operators $2^Q \rightarrow 2^Q$:

$$\begin{aligned} P &\mapsto Sat(\Phi_2) \cup (Pre(P, N) \cap Sat(\Phi_1)) && \text{(until)} \\ P &\mapsto Sat(\Phi_2) \cap (Pre(P, N) \cup Sat(\Phi_1)) && \text{(release)} \end{aligned}$$

¹ The same observation holds for *ATL** interpreted over concurrent games, but for other reasons.

Hence, in *ASL* with the standard semantics we have the following *expansion laws*:

$$\mathbb{E}_N(\Phi_1 \cup \Phi_2) \equiv \Phi_2 \vee (\Phi_1 \wedge \mathbb{E}_N \circ \mathbb{E}_N(\Phi_1 \cup \Phi_2)) \quad (1)$$

$$\mathbb{E}_N(\Phi_1 \text{R} \Phi_2) \equiv \Phi_2 \wedge (\Phi_1 \vee \mathbb{E}_N \circ \mathbb{E}_N(\Phi_1 \text{R} \Phi_2)) \quad (2)$$

where \equiv denotes equivalence of *ASL* state formulas. On the basis of (1) and (2), we obtain that for winning objectives formalized by *ASL* path formulas φ of the form $(\Phi_1 \cup \Phi_2)$ or $(\Phi_1 \text{R} \Phi_2)$, memoryless strategies are sufficient and the satisfaction set $Sat(\mathbb{E}_N \varphi)$ can be computed by means of the standard procedures to compute least and greatest fixed points of monotonic operators. The algorithms for until and release including the proof of correctness can be found in the technical report [19]. For *ASL* state formulas of the form $\mathbb{E}_N \langle\langle \alpha \rangle\rangle \Phi$ or $\mathbb{E}_N [\alpha] \Phi$, we follow an automata-theoretic approach which resembles the standard automata-based *LTL* model checking procedure and relies on a representation of α by means of a finite automaton \mathcal{Z} and a graph analysis of the product $\mathcal{A} \bowtie \mathcal{Z}$. As α is roughly an ordinary regular expression, we can apply standard methods to generate a deterministic finite automata \mathcal{Z} over the alphabet $\text{CIO}_{\sqrt{}}$ such that the accepted language of \mathcal{Z} agrees with $IOS(\alpha)$.

Let $\mathcal{Z} = (Z, \text{CIO}_{\sqrt{}}, \delta, Z_0, Z_F)$, i.e., Z stands for the state space, z_0 the initial state, Z_F for the set of final (accept) states and $\delta : Z \times \text{CIO}_{\sqrt{}} \rightarrow Z$ for the transition function. In fact, beside the special $\sqrt{}$ -transitions, \mathcal{Z} can be viewed as a CA where the set Z_F plays the role of the labeling function which separates the final states from the non-final states. Due to the special role of the symbol $\sqrt{}$ (which can only appear at the end of a word in $IOS(\alpha)$), we can assume that there are special states $z_{accept} \in Z_F$ and $z_{reject} \in Z \setminus Z_F$ such that each $\sqrt{}$ -transition leads to one of the states z_{accept} or z_{reject} and that the states z_{accept} or z_{reject} cannot be entered via any other symbol. Given \mathcal{A} and \mathcal{Z} , we built the product $\mathcal{A} \bowtie \mathcal{Z}$, similar to the product of finite automata and the join operator for CAs [6], but with a special treatment of the pseudo-transitions with label $\sqrt{}$. In fact, the product construction we use here differs from those used in the *BTSL* model checking procedure [18] since in the context of the \mathbb{E}_N -operator we have to incorporate the possibilities of the N-agents to enforce termination. Formally, we define the CA $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ as follows:

$$\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z} \stackrel{\text{def}}{=} (S, \mathcal{N} \cup \{A_{stop}\}, \longrightarrow, S_0, AP', L').$$

The state space S is $Q \times Z$ and A_{stop} is a new node-name (not contained in \mathcal{N}). This new node is supposed to be controllable. (Thus, for $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ we will ask for $(\mathcal{N} \cup \{A_{stop}\})$ -strategies rather than N-strategies.) The initial states are given by

$$S_0 = \{ \langle q, z_0 \rangle : q \in Q_0 \}.$$

The atomic propositions and labeling function in $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ are given by the set $AP' = \{a_\Phi, accept\}$, where $a_\Phi \in L'(\langle q, z \rangle)$ iff $q \models \Phi$ and $accept \in L'(\langle q, z \rangle)$ iff $z \in Z_F$. The transitions in $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ are obtained by the following synchronization rule for concurrent I/O-operations $c \in \text{CIO}$ (i.e., $c \neq \sqrt{}$), state q in \mathcal{A} , and state $z \in Z \setminus \{z_{accept}, z_{reject}\}$:

$$\frac{q \xrightarrow{c} \mathcal{A} q' \wedge z \xrightarrow{c} \mathcal{Z} z'}{\langle q, z \rangle \xrightarrow{c} \langle q', z' \rangle} \quad (3)$$

where we use the subscript \mathcal{A} for the transition relations in \mathcal{A} . In addition, we have the following rules for each terminal state q in \mathcal{A} and state $z \in Z \setminus \{z_{accept}, z_{reject}\}$ where c_{stop} is a concurrent I/O-operation with $\text{Nodes}(c_{stop}) = \{A_{stop}\}$ and $c_{stop}(A_{stop})$ is an arbitrary element from the data domain $Data$:

$$\frac{\neg \exists c \in \text{CIO}(q) \text{ s.t. } \text{Nodes}(c) \subseteq N \wedge c_\emptyset \notin \text{CIO}(q)}{\langle q, z \rangle \xrightarrow{c_\emptyset} \langle q, \delta(z, \sqrt{\ }) \rangle} \quad (4)$$

$$\frac{\exists c \in \text{CIO}(q) \text{ s.t. } \text{Nodes}(c) \cap N \neq \emptyset \wedge c_\emptyset \notin \text{CIO}(q)}{\langle q, z \rangle \xrightarrow{c_{stop}} \langle q, \delta(z, \sqrt{\ }) \rangle} \quad (5)$$

Rule (4) formalizes the fact that if q is terminal (i.e., $c_\emptyset \notin \text{CIO}(q)$) and there is no $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \subseteq N$ then the opponents of the N -agents may refuse any write or read operation and can therefore enforce data flow to stop. This is modeled in the product by a transition with the label c_\emptyset . Rule (5) stands for the fact that whenever q is a terminal node for which some concurrent I/O-operation c is enabled where the N -nodes are involved then the N -agents might decide not to participate in any further I/O-operation. This is modeled in the product by a transition with the label c_{stop} where the new node A_{stop} is supposed to be controllable. We obtain the following two lemmas for ASL state formulas of the form $\mathbb{E}_N \langle \langle \alpha \rangle \rangle \Phi$ and $\mathbb{E}_N \llbracket \alpha \rrbracket \Phi$.

Lemma 1. Let \mathcal{A} be a CA, $\mathcal{Z} = (Z, \text{CIO}_\sqrt{\ }, \delta, Z_0, Z_F)$ a DFA for α , q in \mathcal{A} , node-sets $N \subseteq \mathcal{N}$ and ASL state formulas Φ . Then, the following statements are equivalent:

- (a) $q \models \mathbb{E}_N \langle \langle \alpha \rangle \rangle \Phi$
- (b) $\langle q, z_0 \rangle \models \mathbb{E}_{N \cup \{A_{stop}\}} \diamond (\mathbf{a}_\Phi \wedge \text{accept})$
- (c) There exists a finite-memory N -strategy \mathfrak{S} for \mathcal{A} that is winning for $\langle q, \langle \langle \alpha \rangle \rangle \Phi$

Lemma 2. Let \mathcal{A} be a CA, $\mathcal{Z} = (Z, \text{CIO}_\sqrt{\ }, \delta, Z_0, Z_F)$ a DFA for α , q in \mathcal{A} , node-sets $N \subseteq \mathcal{N}$ and ASL state formulas Φ . Then, the following statements are equivalent:

- (a) $q \models \mathbb{E}_N \llbracket \alpha \rrbracket \Phi$
- (b) $\langle q, z_0 \rangle \models \mathbb{E}_{N \cup \{A_{stop}\}} \square (\text{accept} \rightarrow \mathbf{a}_\Phi)$
- (c) there exists a finite memory N -strategy \mathfrak{S} which is winning for $\langle q, \llbracket \alpha \rrbracket \Phi$

Thanks to lemmas 1 and 2 the satisfaction sets $\text{Sat}(\mathbb{E}_N \langle \langle \alpha \rangle \rangle \Phi)$ and $\text{Sat}(\mathbb{E}_N \llbracket \alpha \rrbracket \Phi)$ can be computed by means of a reduction to the model checking problem for the \mathbb{E}_N -operator in combination with the eventually- and always-modalities. More precisely, we first have to construct a DFA \mathcal{Z} for α , then build the product $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ and finally apply the algorithm for until and release respectively, to compute the satisfaction sets for $\mathbb{E}_{N \cup \{A_{stop}\}} \diamond (\mathbf{a}_\Phi \wedge \text{accept})$ and $\mathbb{E}_{N \cup \{A_{stop}\}} \square (\text{accept} \rightarrow \mathbf{a}_\Phi)$ in the product. Furthermore the memoryless $(N \cup \{A_{stop}\})$ -strategies for the product yield finite-memory winning N -strategies in \mathcal{A} for the objectives $\langle \langle \alpha \rangle \rangle \Phi$ and $\llbracket \alpha \rrbracket \Phi$, respectively.

Assuming that $\text{Sat}(\Phi)$ has already been computed the time complexity for computing $\text{Sat}(\mathbb{E}_N \langle \langle \alpha \rangle \rangle \Phi)$ and $\text{Sat}(\mathbb{E}_N \llbracket \alpha \rrbracket \Phi)$ is linear in the size of CA \mathcal{A} and the DFA \mathcal{Z} for α (which can be exponential in the length of α). However, when restricting to the *ATL*-fragment of *ASL* which just uses the standard path modalities \cup , \mathbb{R} and \bigcirc , but not $\langle \langle \alpha \rangle \rangle$

or $\llbracket \alpha \rrbracket$, then the worst complexity of the *ASL* model checking algorithm is the same as for standard *ATL*, i.e., linear in the size of \mathcal{A} and the length of the formula.

We conclude this section by a simple observation concerning the case that α is a \surd -free expression (i.e., does not contain a subexpression of the form $\beta; \surd$). In fact, for \surd -free expressions, the “best” strategy for the N-agents to ensure $\llbracket \alpha \rrbracket \Phi$ is to stop the data flow whenever possible. This is formalized in the following lemma.

Lemma 3 (Winning strategies for \surd -free expressions). Let \mathfrak{S}_{stop} be the memoryless N-strategy given by $\mathfrak{S}_{stop}(q) = \{stop\} \cup \{c \in \text{CIO} : \text{Nodes}(c) \cap \mathbb{N} = \emptyset\}$ for all states q . Then, for each \surd -free stream expression α and state q we have:

$$q \models \mathbb{E}_{\mathbb{N}} \llbracket \alpha \rrbracket \Phi \text{ iff } \mathfrak{S}_{stop} \text{ is winning for } \langle q, \llbracket \alpha \rrbracket \Phi \rangle.$$

Thus, if α is \surd -free then the set $\text{Sat}(\mathbb{E}_{\mathbb{N}} \llbracket \alpha \rrbracket \Phi)$ can be computed by considering the sub-automaton \mathcal{A}' of \mathcal{A} that results by the memoryless strategy \mathfrak{S}_{stop} and then computing the satisfaction set for $\text{Sat}_{\mathcal{A}'}(\forall \llbracket \alpha \rrbracket \Phi)$ in \mathcal{A}' . This can be done by means of a *BTSL* model checker [18].

5 ASL with Fairness

The concept of fairness serves to rule out pathological behaviors, where certain liveness properties are violated, although they are supposed to hold [14]. The nondeterminism within our multi-player setting demand for some *ASL* fairness assumptions. To illustrate the need for some fairness assumptions, we reuse the deadlock example (2). One would expect that the *ASL* state formula $\mathbb{E}_{\mathbb{B}} \diamond \neg \exists \bigcirc \text{true}$ would be fulfilled, since the memoryless strategy \mathfrak{S} , which tries to write on B whenever q_0 is reached during an execution should be winning for $\langle q_0, \diamond \neg \exists \bigcirc \text{true} \rangle$. But

$$\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} q_0 \xrightarrow{c_1} \dots \in \text{Paths}(q_0, \mathfrak{S}) \text{ and } \pi \not\models \neg \exists \bigcirc \text{true}.$$

The goal of this section is to introduce some fairness assumptions to exclude such undesirable behaviors from our observations.

Definition 6 ($\langle \mathbb{N}, \mathfrak{S} \rangle$ -fairness). Let $\mathcal{A} = \langle \mathbb{Q}, \mathbb{N}, \longrightarrow, Q_0, AP, L \rangle$ be a CA, $\mathbb{N} \subseteq \mathbb{N}$ a node-set, \mathfrak{S} an N-strategy, and $\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots$ a \mathfrak{S} -path in \mathcal{A} . Then π is called (*strongly*) $\langle \mathbb{N}, \mathfrak{S} \rangle$ -fair if either π is finite or for all $c \in \text{CIO}$ we have:

$$\overset{\infty}{\exists} i \geq 0. c \in \text{CIO}(q_i) \cap \mathfrak{S}(\pi \downarrow i) \text{ and } \text{Nodes}(c) \subseteq \mathbb{N} \text{ implies } \overset{\infty}{\exists} i \geq 0. c_i = c,$$

where $\overset{\infty}{\exists} i$ means “there exists infinitely many i ”. We write $\text{FairPaths}_{\langle \mathbb{N}, \mathfrak{S} \rangle}(q)$ for all $\langle \mathbb{N}, \mathfrak{S} \rangle$ -fair paths starting in q and $\text{FairPaths}_{\langle \mathbb{N}, \mathfrak{S} \rangle}(\mathcal{A})$ for the set of $\langle \mathbb{N}, \mathfrak{S} \rangle$ -fair paths.

In the above example, $\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} q_0 \xrightarrow{c_1} \dots \notin \text{FairPaths}_{\langle \{\mathbb{B}\}, \mathfrak{S} \rangle}(q_0)$ because \mathfrak{S} is willing to write infinitely often on B, but no write operation is ever executed. The semantics of the fair *ASL* path formulas is the same as for *ASL* without fairness (see section 4.1).

The semantics for fair ASL state formulas also corresponds to the one without fairness except for:

$$q \models_{\text{fair}} \mathbb{E}_N \varphi \text{ iff there is an } N\text{-strategy } \mathcal{G} \text{ s.t. for all } \pi \in \text{FairPaths}_{(N, \mathcal{G})}(q) : \pi \models \varphi$$

The underlying model checking algorithms need to be modified and now rely on the bottom up computation of the sets $Sat_{\text{fair}}(\Psi) = \{q \in Q \mid q \models_{\text{fair}} \Psi\}$ for all subformulas Ψ . The computation for $Sat_{\text{fair}}(\mathbb{E}_N(\Phi_1 R \Phi_2))$ does not involve any modification at all, as shown in the following lemma.

Lemma 4 (Release with fairness). Let \mathcal{A} be a CA, $N \subseteq \mathcal{N}$ a node-set, $q \in Q$ a state in \mathcal{A} and Φ_1, Φ_2 ASL state formulas. Then $q \models_{\text{fair}} \mathbb{E}_N(\Phi_1 R \Phi_2)$ iff $q \models \mathbb{E}_N(\Phi_1 R \Phi_2)$.

The computation of $Sat_{\text{fair}}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ relies on an iterative SCC-calculation in subgraphs of \mathcal{A} . The following lemma emerges that the remaining fair computation of $Sat_{\text{fair}}(\mathbb{E}_N \langle\langle \alpha \rangle\rangle \Phi)$ and $Sat_{\text{fair}}(\mathbb{E}_N \llbracket \alpha \rrbracket \Phi)$ can be reduced to eventually and always in the product $\mathcal{A} \bowtie \mathcal{Z}$.

Lemma 5 (Fairness for ASL I/O-stream expression formulas). Let \mathcal{A} be a CA, $N \subseteq \mathcal{N}$ a node-set, α a regular I/O-stream expression, \mathcal{Z} a deterministic CA for α , and let Φ be ASL state formula. Then, the following observation holds for all states $q \in Q$.

- i) $q \models_{\text{fair}} \mathbb{E}_N \langle\langle \alpha \rangle\rangle \Phi$ in \mathcal{A} iff $\langle q, z_0 \rangle \models_{\text{fair}} \mathbb{E}_{N \cup \{\mathcal{A}_{\text{stop}}\}} \diamond (\text{accept} \wedge \mathbf{a}_\Phi)$ in $\mathcal{A} \bowtie \mathcal{Z}$.
- ii) $q \models_{\text{fair}} \mathbb{E}_N \llbracket \alpha \rrbracket \Phi$ iff $\langle q, z_0 \rangle \models_{\text{fair}} \mathbb{E}_{N \cup \{\mathcal{A}_{\text{stop}}\}} \square (\text{accept} \rightarrow \mathbf{a}_\Phi)$ in $\mathcal{A} \bowtie \mathcal{Z}$.

6 Conclusion and Future Work

This paper introduces a framework to verify alternating-time properties for a multi-player games derived from CA. The introduced concurrent game semantics captures any complex behavior caused by synchronous and asynchronous peer-to-peer communication, mutual dependencies of I/O-operations and also data-dependencies. Since this game structure is non-standard it takes numerous nontrivial adaptations of the *ATL* model checking algorithm. In future work we will drop our assumption on *perfect information* and *perfect recall* to switch to a more realistic view for exogenous coordination taking the local view [5, 7, 16, 17, 21, 11, 22] into account. In future work we will consider *observation-based strategies* in case of *incomplete information*.

Apart from asking for the existence or absence of a winning strategy for a temporal property the question might raise, if there is a way of connecting the components to make this property hold. This directly leads to the controller synthesis problem where if possible a controlling CA is put in parallel with the other components to ensure the intended behavior. One step further we would like to build the Reo network which glues those components the intended way by using the synthesis approach described in [4].

References

1. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
2. F. Arbab. Reo: A channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14(3):329–366, 2004.
3. F. Arbab, C. Baier, F. de Boer, and J.J.M.M. Rutten. Models and temporal logics for timed component connectors. In *Proc. of SEFM*, pages 198–207. IEEE CS Press, 2004.
4. F. Arbab, C. Baier, F. de Boer, J.J.M.M. Rutten, and M. Sirjani. Synthesis of Reo circuits for implementation of component connector automata specifications. In *Proc. of COORDINATION*, volume 3454 of *LNCS*, 2005.
5. S. Azhar, G.L. Peterson, and J.H. Reif. On multiplayer non-cooperative games of incomplete information: Part 1&2. Technical report, Durham, NC, USA, 1991.
6. C. Baier, M. Sirjani, F. Arbab, and J.J.M.M. Rutten. Modeling component connectors in Reo by constraint automata. In *Science of Computer Programming 61*, pages 75–113., 2006.
7. K. Chatterjee, L. Doyen, T.A. Henzinger, and J.F. Raskin. Algorithms for omega-regular games with imperfect information. *CoRR*, abs/0706.2619, 2007.
8. E.M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM TOPLAS*, 8(2):244–263, 1986.
9. L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *Proc. of LICS*, pages 141–154, Jan. 2000.
10. L. de Alfaro and T.A. Henzinger. Interface automata. In *FSE Proc*, pages 109–120. ACM Press, 2001.
11. M. de Wulf, Laurent Doyen, and Jean-François Raskin. A lattice theory for solving games of imperfect information. In *HSCC*, pages 153–168, 2006.
12. M. J. Fischer and R.J. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Science*, 8:194–211, 1979.
13. David Fitoussi and Moshe Tennenholtz. Choosing social laws for multi-agent systems: minimality and simplicity. *Artif. Intell.*, 119(1-2):61–101, 2000.
14. N. Francez. *Fairness*. Springer-Verlag, 1986.
15. R. Grosu and B. Rumpe. Concurrent timed port automata. Technical Report TUM-I9533, Techn. Univ. München, 1995. <http://www4.informatik.tu-muenchen.de/reports/>.
16. W.v.d. Hoek, M. Roberts, and M. Wooldridge. Knowledge and social laws. In *AAMAS*, pages 674–681, 2005.
17. W.v.d. Hoek and M. Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1):125–157, 2003.
18. S. Klüppelholz and C. Baier. Symbolic model checking for channel-based component connectors. In *Proc. of FOCLASA 2006*, volume 175(2) of *ENTCS*, pages 19–37, 2007.
19. S. Klüppelholz and C. Baier. Alternating-Time Stream Logic for Multi-Agent Systems. Technical report, Technical University Dresden, 2008. <http://wwwtcs.inf.tu-dresden.de/~klueppel/ASLKB2008.pdf>.
20. N. Lynch and M.R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, 1989.
21. J.H. Reif. The complexity of two-player games of incomplete information. *J. Comput. Syst. Sci.*, 29(2):274–301, 1984.
22. P.Y. Schobbens. Alternating-time logic with imperfect recall. In *Proc. of LCMAS*, volume 85(2) of *ENTCS*, pages 1–12, 2004.
23. P. Wolper. Specification and synthesis of communicating processes using an extended temporal logic. In *Proc. of POPL*, pages 20–33, 1982.
24. M. Wooldridge. Social laws in alternating time. In *DEON*, page 2, 2004.

A Appendix

Definition 7 (I/O-constraints (IOC)). We define $CIO(tt) \stackrel{\text{def}}{=} \text{CIO}$, $CIO(ff) \stackrel{\text{def}}{=} \emptyset$, and for the literals A and $\neg A$:

$$CIO(A) \stackrel{\text{def}}{=} \{c \in \text{CIO} : A \in \text{Nodes}(c)\} \text{ and } CIO(\neg A) \stackrel{\text{def}}{=} \{c \in \text{CIO} : A \notin \text{Nodes}(c)\}.$$

The I/O-constraints $(d_{A_1}, \dots, d_{A_k}) \in D$ impose conditions for the written and read data items. That is, $CIO((d_{A_1}, \dots, d_{A_k}) \in D)$ agrees with the set

$$\{c \in \text{CIO} : \{A_1, \dots, A_k\} \subseteq \text{Nodes}(c), (c(A_1), \dots, c(A_k)) \in D\}.$$

Conjunction and disjunction have their standard meaning, i.e.,

$$\begin{aligned} CIO(ioc_1 \wedge ioc_2) &\stackrel{\text{def}}{=} CIO(ioc_1) \cap CIO(ioc_2) \\ CIO(ioc_1 \vee ioc_2) &\stackrel{\text{def}}{=} CIO(ioc_1) \cup CIO(ioc_2) \end{aligned}$$

□

Definition 8 (Finite memory strategies). A finite-memory N -strategy is a tuple $\mathfrak{M} = (\text{Modes}, \Delta, \mu, m_0)$, where

- Modes is a finite set (of so-called modes),
- $m_0 \in \text{Modes}$ the starting mode,
- $\mu : Q \times \text{Modes} \rightarrow 2^{\text{CIO} \cup \{\text{stop}\}}$ the decision function, and
- $\Delta : \text{Modes} \times (Q \times \text{CIO} \times Q) \rightarrow \text{Modes}$ the transition function.

The associated N -strategy $\mathfrak{S}_{\mathfrak{M}}$ is given by:

$$\mathfrak{S}_{\mathfrak{M}}(q_0 \xrightarrow{c_1} \dots \xrightarrow{c_i} q_i) = \mu(q_i, \Delta^*(m_0, q_0 \xrightarrow{c_1} \dots \xrightarrow{c_i} q_i))$$

where $\Delta^*(m, q_0 \xrightarrow{c_1} q_1) = \Delta(m, q_0 \xrightarrow{c_1} q_1)$ and

$$\Delta^*(m, q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots \xrightarrow{c_i} q_i) = \Delta^*(\Delta(m, q_0 \xrightarrow{c_1} q_1), q_1 \xrightarrow{c_2} \dots \xrightarrow{c_i} q_i).$$

A memoryless strategy is a finite-memory strategy with a single mode m_0 .

□

Lemma 6 (Correctness of the Pre-operator). $Pre(P, N) = \{q \in Q : q \models \mathbb{E}_N \circ P\}$.

Proof. “ \subseteq ”: Suppose that $q \in Pre(P, N)$. Let \mathfrak{S} be a memoryless N -strategy such that

$$\begin{aligned} \mathfrak{S}(q) &= \{c \in \text{CIO} : \text{Nodes}(c) \cap N = \emptyset\} \cup \\ &\quad \{c \in \text{CIO}(q) : \text{Nodes}(c) \subseteq N \wedge Post[c](q) \subseteq P\} \end{aligned}$$

Then, $\{c \in \mathfrak{S}(q) : \text{Nodes}(c) \subseteq N\} \neq \emptyset$ by condition (2) and $Post[c](q) \subseteq P$ for all $c \in \mathfrak{S}(q) \cap \text{CIO}(q)$ by condition (1). Hence, each \mathfrak{S} -complete execution η from q starts with a transition $q \xrightarrow{c} p$ where $c \in \mathfrak{S}(q) \cap \text{CIO}(q)$. But then $\pi \models \circ P$ for all $\pi \in \text{Paths}(q, \mathfrak{S})$. Thus, \mathfrak{S} yields a witness for $q \models \mathbb{E}_N \circ P$.

“ \supseteq ”: Suppose now that $q \models \mathbb{E}_N \circ P$. We have to check conditions (1) and (2) in definition 5. Let \mathfrak{S} be a memoryless N -strategy winning for $\langle q, \circ P \rangle$.

- (1) Let $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \cap N = \emptyset$ and let $p \in \text{Post}[c](q)$. By the requirements for N-strategies, we have $c \in \mathfrak{S}(q)$. Let π be an arbitrary \mathfrak{S} -path that starts with the transition $q \xrightarrow{c} p$. Since $\pi \models \bigcirc P$ we get $p \in P$.
- (2) If q is terminal then the execution q of length 0 cannot be \mathfrak{S} -complete, since $q \xrightarrow{\vee} q \not\models \bigcirc \text{true}$ and $\bigcirc P$ holds for all \mathfrak{S} -paths from q . Hence, there exists a concurrent I/O-operation $c \in \text{CIO}(q) \cap \mathfrak{S}(q)$ such that $\text{Nodes}(c) \subseteq N$. We now show that $\text{Post}[c](q) \subseteq P$. For each state $p \in \text{Post}[c](q)$ there exists a \mathfrak{S} -path π that starts with the transition $q \xrightarrow{c} p$. As $\pi \models \bigcirc P$ we get $p \in P$.

□

Algorithms (1) and (2) show how the satisfaction sets for $\mathbb{E}_N(\Phi_1 \cup \Phi_2)$ and $\mathbb{E}_N(\Phi_1 \text{R} \Phi_2)$ can be computed together with a memoryless N-strategy \mathfrak{S} that is winning for all states q where $\mathbb{E}_N(\Phi_1 \cup \Phi_2)$ or $\mathbb{E}_N(\Phi_1 \text{R} \Phi_2)$ holds.

Algorithm 1 Algorithm for computing $\text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$;

```

P0 := Sat(Φ2);
i := 0;
repeat
  Pi+1 := Pi ∪ (Sat(Φ1) ∩ Pre(Pi, N));
  for all states p ∈ Pi+1 \ Pi do
    S(p) := { c ∈ CIO : Nodes(c) ∩ N = ∅ ∨ ∅ ≠ Post[c](p) ⊆ Pi };
  end for
  i := i + 1;
until Pi = Pi-1;
for all states p ∈ (Q \ Pi) ∪ Sat(Φ2) do
  S(p) := CIO ∪ {stop};
end for
return Pi;

```

(* P_i = Sat($\mathbb{E}_N(\Phi_1 \cup \Phi_2)$) *)

Lemma 7 (Correctness of algorithms (1) and (2)). Let \mathcal{A} be a CA as before, $N \subseteq \mathcal{N}$ a node-set and let Φ_1 and Φ_2 be ASL state formulae. Then:

- (a) Algorithm 1 correctly returns the set $\text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ and the computed memoryless N-strategy \mathfrak{S} is winning for all states $q \in \text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ and ASL path formula $(\Phi_1 \cup \Phi_2)$.
- (b) Algorithm 2 correctly returns the set $\text{Sat}(\mathbb{E}_N(\Phi_1 \text{R} \Phi_2))$ and the computed memoryless N-strategy \mathfrak{S} is winning for all states $q \in \text{Sat}(\mathbb{E}_N(\Phi_1 \text{R} \Phi_2))$ and ASL path formula $(\Phi_1 \text{R} \Phi_2)$.

Proof. Both algorithms rely on the standard iterative approach to compute least and greatest fixed points of monotonic operators. This yields that the returned sets P_i agree with the satisfaction set $\text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ and $\text{Sat}(\mathbb{E}_N(\Phi_1 \text{R} \Phi_2))$, respectively. It remains to check that the computed strategies are winning.

Algorithm 2 Algorithm for computing $Sat(\mathbb{E}_N(\Phi_1 R \Phi_2))$;

for all states $p \in Q$ **do**
 $\mathfrak{S}(p) := \text{CIO} \cup \{\text{stop}\}$;
end for
 $P_0 := \text{Sat}(\Phi_2)$;
 $i := 0$;
repeat
 $P_{i+1} := P_i \cap (\text{Sat}(\Phi_1) \cup \text{Pre}(P_i, N))$;
for all states $p \in P_{i+1} \setminus \text{Sat}(\Phi_1)$ **do**
 $\mathfrak{S}(p) := \mathfrak{S}(p) \setminus \{c \in \text{CIO}(p) : \text{Post}[c](p) \not\subseteq P_i\}$;
end for
 $i := i + 1$;
until $P_i = P_{i-1}$;
return P_i ; (* $P_i = \text{Sat}(\mathbb{E}_N(\Phi_1 R \Phi_2))$ *)

- (a) Let $P_j = \text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$. We first observe that if state $q \in P_{i+1} \setminus P_i$ then $q \in \text{Pre}(P_i, N)$. By the definition of the Pre-operator and the definition of $\mathfrak{S}(q)$ we obtain that $\text{Post}[c](q) \subseteq P_i$ for all $c \in \mathfrak{S}(q)$ and that $\mathfrak{S}(q) \cap \{c \in \text{CIO}(q) : \text{Nodes}(c) \subseteq N\} \neq \emptyset$. From this, we get by induction on n that for each finite \mathfrak{S} -execution

$$\eta = q_0 \xrightarrow{c_0} q_1 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n$$

such that $q_0 \models \mathbb{E}_N(\Phi_1 \cup \Phi_2)$ and $q_i \not\models \Phi_2$ for $0 \leq i \leq n$ the following two conditions hold:

- η is not \mathfrak{S} -complete, i.e., there is a $c \in \mathfrak{S}(q_n) \cap \text{CIO}(q_n)$ with $\text{Nodes}(c) \subseteq N$.
- There exist indices $j \geq j_0 > j_1 > j_2 > \dots > j_n$ such that $q_i \in P_{j_i}$ for $0 \leq i \leq n$.

As $P_i \subseteq \text{Sat}(\Phi_1)$ for $i \geq 1$ and $P_0 = \text{Sat}(\Phi_2)$ we obtain that $\pi \models (\Phi_1 \cup \Phi_2)$ for each \mathfrak{S} -path that starts in a state $q_0 \in P_j = \text{Sat}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$.

- (b) Let $P_j = \text{Sat}(\mathbb{E}_N(\Phi_1 R \Phi_2))$. For the states $q \in P_j \cap \text{Sat}(\Phi_1)$ we have $q \models \Phi_1 \wedge \Phi_2$ (as $P_j \subseteq P_{j-1} \subseteq \dots \subseteq P_0 = \text{Sat}(\Phi_2)$) and therefore $\pi \models (\Phi_1 R \Phi_2)$ for all paths π starting in q .

If $q \in P_j \setminus \text{Sat}(\Phi_1)$ then for all $c \in \mathfrak{S}(q)$ we have $\text{Post}[c](q) \subseteq P_j$ (by definition of \mathfrak{S}) and $q \in \text{Pre}(P_j, N)$. The definition of the Pre-operator yields the existence of a concurrent I/O-operation $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \subseteq N$ and $\text{Post}[c](q) \subseteq P_j$. But then $c \in \mathfrak{S}(q)$, and each \mathfrak{S} -execution ending in q is \mathfrak{S} -incomplete.

These two observations yield that each \mathfrak{S} -path $\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots \in \text{Paths}(q_0, \mathfrak{S})$ starting in a state $q_0 \in P_j$ is either infinite and consists of states in $P_j \setminus \text{Sat}(\Phi_1)$ or has a prefix $q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n$ where $q_0, \dots, q_{n-1} \models \Phi_2$ and $q_n \models \Phi_1 \wedge \Phi_2$. In both cases, we have $\pi \models (\Phi_1 R \Phi_2)$. □

For the following observations on the properties of the product please notice that $\delta(z, \surd) \in \{z_{\text{accept}}, z_{\text{reject}}\}$. Hence, with the corresponding two transition rules a state is entered

where \mathcal{Z} is in one of its special states z_{accept} or z_{reject} . For technical reasons we add self-loops with label c_\emptyset for such states:

$$\langle q, z_{accept} \rangle \xrightarrow{c_\emptyset} \langle q, z_{accept} \rangle \quad \langle q, z_{reject} \rangle \xrightarrow{c_\emptyset} \langle q, z_{reject} \rangle$$

These transitions ensure that all paths in the product that eventually enter a state $\langle q, z \rangle$ where $z \in \{z_{accept}, z_{reject}\}$ are infinite and repeat state $\langle q, z \rangle$ forever. For each transition in the product (obtained by one of the above composition rules) we define its projection to \mathcal{A} as follows.

- If $\langle q, z \rangle \xrightarrow{c} \langle q', z' \rangle$ arises by applying rule (3) then its \mathcal{A} -projection is $q \xrightarrow{c} q'$.
- If $\langle q, z \rangle \xrightarrow{c} \langle q, z' \rangle$ (where $z' \in \{z_{accept}, z_{reject}\}$ and $c \in \{c_\emptyset, c_{stop}\}$) is obtained from rule (4) or rule (5) then the \mathcal{A} -projection is $q \xrightarrow{\vee} q$.
- The \mathcal{A} -projection of the pseudo-transition $\langle q, z \rangle \xrightarrow{\vee} \langle q, z \rangle$ that might appear at the end of a finite path in $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ is $q \xrightarrow{\vee} q$.

Given a path $\tilde{\pi}$ in $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ that does not enter a state of the form $\langle q, z_{accept} \rangle$ or $\langle q, z_{reject} \rangle$ then we define the \mathcal{A} -projection $proj_{\mathcal{A}}(\tilde{\pi})$ as the unique path in \mathcal{A} that results by taking the \mathcal{A} -projection of all transitions in $\tilde{\pi}$. For a path $\tilde{\pi}$ that eventually enters a state of the form $\langle q, z \rangle$ with $z \in \{z_{accept}, z_{reject}\}$ we ignore the (infinite) suffix $\langle q, z \rangle \xrightarrow{c_\emptyset} \langle q, z \rangle \xrightarrow{c_\emptyset} \dots$ and define $proj_{\mathcal{A}}(\tilde{\pi})$ as the \mathcal{A} -projection of the prefix of $\tilde{\pi}$ that leads to $\langle q, z \rangle$. Similarly, we define the \mathcal{Z} -projection $proj_{\mathcal{Z}}(\tilde{\pi})$ as an infinite or finite sequence of elements in \mathcal{Z} of the same length as $proj_{\mathcal{A}}(\tilde{\pi})$. Then, if $\tilde{\pi}$ starts in a state $\langle q_0, z_0 \rangle$ (where z_0 is the initial state of \mathcal{Z}) then $proj_{\mathcal{Z}}(\tilde{\pi})$ is the run for the I/O-stream of $proj_{\mathcal{A}}(\tilde{\pi})$ in \mathcal{Z} . The definitions of the projections are extended for executions (i.e., prefixes of paths) in the obvious way. From now on, we omit the subscript \mathbb{N} and Φ and simply write $\mathcal{A} \bowtie \mathcal{Z}$ for the product whenever they are clear from the context.

Lemma 8 (Properties of the product).

- (i) If $\langle q, z \rangle$ is terminal in $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ then (a) q is terminal in \mathcal{A} , (b) there is a concurrent I/O-operation $c \in \text{CIO}(q)$ such that $\emptyset \neq \text{Nodes}(c) \subseteq \mathbb{N}$, and (c) c_{stop} is enabled in $\langle q, z \rangle$.
- (ii) If $\tilde{\pi}$ is a path in $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ then $proj_{\mathcal{A}}(\tilde{\pi})$ is a path in \mathcal{A} .
- (iii) For each path $\tilde{\pi}$ in the product starting in a state $\langle q, z_0 \rangle$ we have: $\tilde{\pi} \models \diamond(\alpha_\Phi \wedge accept)$ iff $proj_{\mathcal{A}}(\tilde{\pi}) \models \llbracket \alpha \rrbracket \Phi$.
- (iv) For each path $\tilde{\pi}$ in the product starting in a state $\langle q, z_0 \rangle$ we have: $\tilde{\pi} \models \square(accept \rightarrow \alpha_\Phi)$ iff $proj_{\mathcal{A}}(\tilde{\pi}) \models \llbracket \alpha \rrbracket \Phi$.
- (v) Let \mathfrak{S} be an \mathbb{N} -strategy for \mathcal{A} and \mathfrak{T} an $(\mathbb{N} \cup \{A_{stop}\})$ -strategy for $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ such that for all finite executions $\tilde{\eta}$ in $\mathcal{A} \bowtie_{\mathbb{N}, \Phi} \mathcal{Z}$ starting in a state $\langle q, z_0 \rangle$ the following conditions hold:
 - a) If $c : \mathbb{N} \rightarrow (Data \cup \{\perp\})$ is a concurrent I/O-operation for node-set \mathbb{N} then $c \in \mathfrak{T}(\tilde{\eta})$ iff $c \in \mathfrak{S}(proj_{\mathcal{A}}(\tilde{\eta}))$
 - b) $c_{stop} \in \mathfrak{T}(\tilde{\eta})$ iff $stop \in \mathfrak{S}(proj_{\mathcal{A}}(\tilde{\eta}))$
 - c) $stop \notin \mathfrak{T}(\tilde{\eta})$

then the \mathcal{A} -projections of the \mathfrak{T} -paths starting in a state $\langle q, z_0 \rangle$ are exactly the \mathfrak{S} -paths starting in q .

Proof. ad (i). Let $\langle q, z \rangle$ be a terminal state in the product.

- State q is terminal in \mathcal{A} . This is due to the fact that each transition $q \xrightarrow{c_\emptyset} p$ in \mathcal{A} can be lifted to a transition $\langle q, z \rangle \xrightarrow{c_\emptyset} \langle p, \delta(z, c_\emptyset) \rangle$ in the product.
- There is some concurrent I/O-operation $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \subseteq N$, as otherwise rule (4) would yield that c_\emptyset is enabled in $\langle q, z \rangle$.
- Furthermore, $c_{stop} \in \text{CIO}(\langle q, z \rangle)$. This can be seen as follows. We have $c_\emptyset \notin \text{CIO}(q)$ (as q is terminal in \mathcal{A}). Suppose by contradiction that $c_{stop} \notin \text{CIO}(\langle q, z \rangle)$. Then, there is no $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \cap N \neq \emptyset$ (premise of rule (5)). But then $\text{Nodes}(c) \cap N = \emptyset$ for all $c \in \text{CIO}(q)$ and (as $c_\emptyset \notin \text{CIO}(q)$) there is no $c \in \text{CIO}(q)$ such that $\text{Nodes}(c) \subseteq N$. But then rule (4) yields $c_\emptyset \in \text{CIO}(\langle q, z \rangle)$. This contradicts the assumption that $\langle q, z \rangle$ is terminal in the product.

ad (ii). By (i) we get that all paths in the product are infinite or end in a state $\langle q, z \rangle$ where q is terminal in \mathcal{A} and c_{stop} is enabled in $\langle q, z \rangle$. The projection of an infinite path $\tilde{\pi}$ in the product that never enters a state in

$$S_\surd \stackrel{\text{def}}{=} \{ \langle q, z_{accept} \rangle, \langle q, z_{reject} \rangle : q \in Q \}$$

is an infinite path in \mathcal{A} , since all their transitions arise by rule (3) (i.e., their labels are concurrent I/O-operations for the original node-set N). The same holds for all finite paths in the product that do not enter S_\surd . They end in a terminal state $\langle q, z \rangle$ of the product. But then q is terminal in \mathcal{A} and the \mathcal{A} -projection is a finite path in \mathcal{A} . Paths in $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ that eventually enter a state in S_\surd are infinite, but they are projected to finite paths in \mathcal{A} .

ad (iii). Let $\tilde{\pi}$ be a path in the product starting in a state $\langle q, z_0 \rangle$ and let $\pi \stackrel{\text{def}}{=} \text{proj}_{\mathcal{A}}(\tilde{\pi})$ be its \mathcal{A} -projection.

- Suppose first that $\tilde{\pi} \models \diamond(\alpha_\Phi \wedge \text{accept})$. Then, $\tilde{\pi}$ has a finite prefix that leads to a state $\langle p, z \rangle$ where $(\alpha_\Phi \wedge \text{accept})$ holds. Hence, $p \models \Phi$ and $z \in Z_F$. Let n be the length of this prefix, z_0, z_1, \dots, z_n be the sequence of states obtained by the projection $\text{proj}_{\mathcal{Z}}(\tilde{\pi} \downarrow n)$ and $c_1 \dots c_n$ its I/O-stream. We may suppose that $n \leq |\pi|$. (Note that paths that eventually enter a state $\langle p, z \rangle \in S_\surd$ stay in this state $\langle p, z \rangle$ forever.) Then, we have:

- $c_1 \dots c_n = \text{ios}(\pi \downarrow n)$
- z_0, z_1, \dots, z_n is the run for $c_1 \dots c_n$ in \mathcal{Z} and $z_n = z \in Z_F$

But then $c_1 \dots c_n$ is accepted by \mathcal{Z} and we get $c_1 \dots c_n \in \text{IOS}(\alpha)$. Furthermore, the last state of $\pi \downarrow n$ is p . Since Φ holds in p , this yields $\pi \models \langle\langle \alpha \rangle\rangle \Phi$.

- Suppose now that $\pi \models \langle\langle \alpha \rangle\rangle \Phi$. Then, there is some prefix $\pi \downarrow n$ of π such that its I/O-stream $\text{ios}(\pi \downarrow n)$ belongs to $\text{IOS}(\alpha)$ and the last state p of $\pi \downarrow n$ belongs to $\text{Sat}(\Phi)$. Let z_0, \dots, z_n be the run for $\text{ios}(\pi \downarrow n)$ in \mathcal{Z} . Then, $z_n \in Z_F$ and state $\langle p, z_n \rangle$ is the last state of $\tilde{\pi} \downarrow n$. As $(\alpha_\Phi \wedge \text{accept})$ holds in $\langle p, z_n \rangle$ we get $\tilde{\pi} \models \diamond(\alpha_\Phi \wedge \text{accept})$.

ad (iv). Let $\tilde{\pi} = \langle q_0, z_0 \rangle \xrightarrow{c_1} \langle q_1, z_1 \rangle \xrightarrow{c_2} \dots$ be a path in the product starting in a state $\langle q_0, z_0 \rangle$ and let $\pi \stackrel{\text{def}}{=} \text{proj}_{\mathcal{A}}(\tilde{\pi})$ be its \mathcal{A} -projection.

- Suppose first that $\tilde{\pi} \models \Box(\text{accept} \rightarrow \alpha_\Phi)$. Let $n \leq |\pi|$ such that $\text{ios}(\pi \downarrow n) \in \text{IOS}(\alpha)$. Then, z_0, \dots, z_n is the run for $\text{ios}(\pi \downarrow n)$ in \mathcal{Z} . Hence, $z_n \in Z_F$ and therefore $\langle q_n, z_n \rangle \models \text{accept}$. As $\Box(\text{accept} \rightarrow \alpha_\Phi)$ holds for $\tilde{\pi}$ and $\langle q_n, z_n \rangle$ is the $(n+1)$ -st state of $\tilde{\pi}$ we have $\langle q_n, z_n \rangle \models \alpha_\Phi$. This yields that q_n satisfies Φ and therefore $\pi \models \llbracket \alpha \rrbracket \Phi$.
- Assume $\pi \not\models \llbracket \alpha \rrbracket \Phi$. Let $n \leq |\tilde{\pi}|$. The goal is to show that the $(n+1)$ -st state $\langle q_n, z_n \rangle$ of $\tilde{\pi}$ satisfies $(\text{accept} \rightarrow \alpha_\Phi)$. This is obvious, in case accept does not hold for $\langle q_n, z_n \rangle$. Assume now that $\langle q_n, z_n \rangle \models \text{accept}$. Then, $z_n \in Z_F$.
 - If $\langle q_n, z_n \rangle \notin S_\surd$ then $n \leq |\pi|$ and z_0, \dots, z_n is the run for $\text{ios}(\pi \downarrow n)$ in \mathcal{Z} . As $z_n \in Z_F$ we get that $\text{ios}(\pi \downarrow n) \in \text{IOS}(\alpha)$ and therefore $q_n \models \Phi$. But then, $\langle q_n, z_n \rangle \models \alpha_\Phi$.
 - If $\langle q_n, z_n \rangle \in S_\surd$ then there is some $m \leq n$ such that $m \leq |\pi|$ and

$$\langle q_m, z_m \rangle = \langle q_{m+1}, z_{m+1} \rangle = \dots = \langle q_n, z_n \rangle$$

Then, z_0, \dots, z_m is the run for $\text{ios}(\pi \downarrow m)$ in \mathcal{Z} . As $z_m = z_n \in Z_F$ we get that $\text{ios}(\pi \downarrow m) \in \text{IOS}(\alpha)$ and therefore $q_m \models \Phi$. But this yields, $\langle q_m, z_m \rangle = \langle q_n, z_n \rangle \models \alpha_\Phi$.

ad (v). We first show that the projection of each \mathfrak{T} -path is a \mathfrak{S} -path:

- Each \mathfrak{T} -execution that enters S_\surd via a transition $\langle p, z \rangle \xrightarrow{c_\emptyset} \langle p, \delta(z, \surd) \rangle$ is projected to a finite path π that ends with the transition $p \xrightarrow{\surd} p$. By the premise of rule (4) for the product, we get that $\text{Nodes}(c) \setminus N \neq \emptyset$ for all $c \in \text{CIO}(p)$. Hence, the prefix $\pi \downarrow n$ of π (where $n = |\pi| - 1$) leads from some initial state $q_0 \in Q_0$ to p and constitutes a \mathfrak{S} -complete execution. Hence, π is a finite \mathfrak{S} -path.
- Each \mathfrak{T} -executions that enters S_\surd via a transition $\langle p, z \rangle \xrightarrow{c_{\text{stop}}} \langle p, \delta(z, \surd) \rangle$ is also projected to a finite path π that ends with the transition $p \xrightarrow{\surd} p$. Again, let $n = |\pi| - 1$. The concurrent I/O-operation c_{stop} belongs to $\mathfrak{T}(\tilde{\eta})$ for the prefix $\tilde{\eta} = \tilde{\pi} \downarrow n$ that leads from the first state $\langle q, z_0 \rangle$ of $\tilde{\pi}$ to $\langle p, z \rangle$. By the second assumption on the relation between \mathfrak{T} and \mathfrak{S} we get that $\text{stop} \in \mathfrak{S}(\eta)$ for the projection $\eta = \text{proj}_{\mathcal{A}}(\tilde{\eta})$. But since p is terminal (by the premise of rule (5) in the product) we get that $\eta = \pi \downarrow n$ is \mathfrak{S} -complete. Therefore, π is a finite \mathfrak{S} -path.
- We now regard a \mathfrak{T} -complete finite execution $\tilde{\eta}$ that does not visit S_\surd and ends in state $\langle p, z \rangle$. Then, $\langle p, z \rangle$ is terminal and there is no $\tilde{c} \in \mathfrak{T}(\tilde{\eta}) \cap \text{CIO}(\langle p, z \rangle)$ such that $\text{Nodes}(\tilde{c}) \subseteq N \cup \{A_{\text{stop}}\}$. Hence, there is no $c \in \mathfrak{S}(\eta) \cap \text{CIO}(p)$ such that $\text{Nodes}(c) \subseteq N$. But then $\eta \stackrel{\text{def}}{=} \text{proj}_{\mathcal{A}}(\tilde{\eta})$ is a \mathfrak{S} -complete execution and therefore the corresponding path π is a \mathfrak{S} -path.
- Given an infinite \mathfrak{T} execution $\tilde{\eta}$ that does not enter S_\surd , its \mathcal{A} -projection is an infinite path in \mathcal{A} and therefore a \mathfrak{S} -path.

This shows that the projections of all \mathfrak{T} -paths are \mathfrak{S} -paths.

We now regard a \mathfrak{S} -path π in \mathcal{A} and show that it is the \mathcal{A} -projection of some \mathfrak{T} -path. This is obvious if π is infinite since then it can be lifted to an infinite \mathfrak{T} -path in the product that does not enter S_{\surd} . Assume now that

$$\pi = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} q_n \xrightarrow{\surd} q_n$$

is finite of length $n+1$. Let $z_0, z_1, \dots, z_n, z_{n+1}$ be the run for the I/O-stream $c_1 \dots c_n \surd$ of π . Then,

$$\tilde{\eta} = \langle q_0, z_0 \rangle \xrightarrow{c_1} \dots \xrightarrow{c_n} \langle q_n, z_n \rangle$$

is a \mathfrak{T} -execution in the product and its projection $\eta \stackrel{\text{def}}{=} \text{proj}_{\mathcal{A}}(\tilde{\eta}) = \pi \downarrow n$ is \mathfrak{S} -complete. Hence, q_n is terminal and at least one of the following two conditions (1) or (2) holds:

- (a) $\text{stop} \in \mathfrak{S}(\eta)$
- (b) there is no concurrent I/O-operation $c \in \mathfrak{S}(\eta) \cap \text{CIO}(q_n)$ such that $\text{Nodes}(c) \subseteq N$.

If $\langle q_n, z_n \rangle$ is non-terminal then c_{\emptyset} is enabled in $\langle q_n, z_n \rangle$ because of rule (4) for the product and we have $\text{Post}[c_{\emptyset}](\langle q_n, z_n \rangle) = \{\langle q_n, \delta(z_n, \surd) \rangle\}$. But then π is the projection of the infinite \mathfrak{T} -path

$$\tilde{\eta} \xrightarrow{c_{\emptyset}} \langle q_n, \delta(z_n, \surd) \rangle \xrightarrow{c_{\emptyset}} \langle q_n, \delta(z_n, \surd) \rangle \xrightarrow{c_{\emptyset}} \dots$$

Let us now assume that $\langle q_n, z_n \rangle$ is terminal, i.e., c_{\emptyset} is not enabled in $\langle q_n, z_n \rangle$. Then, for all $c \in \text{CIO}(q_n)$ we have $\emptyset \neq \text{Nodes}(c) \subseteq N$ (otherwise the premise of rule (4) applies and $\langle q_n, z_n \rangle$ would be non-terminal).

- Suppose that case (a) applies. Then, $c_{\text{stop}} \in \mathfrak{T}(\tilde{\eta})$,

$$\tilde{\eta} \xrightarrow{c_{\text{stop}}} \langle q_n, \delta(z_n, \surd) \rangle \xrightarrow{c_{\emptyset}} \langle q_n, \delta(z_n, \surd) \rangle \xrightarrow{c_{\emptyset}} \dots$$

is an infinite \mathfrak{T} -path and its projection is π .

- Suppose that case (b), but not case (a) applies. That is, $\text{stop} \notin \mathfrak{S}(\eta)$ and there is no concurrent I/O-operation $c \in \mathfrak{S}(\eta) \cap \text{CIO}(q_n)$ such that $\text{Nodes}(c) \subseteq N$. Then, $c_{\text{stop}} \notin \mathfrak{T}(\tilde{\eta})$. Hence, there is no concurrent I/O-operation $\tilde{c} \in \mathfrak{T}(\tilde{\eta}) \cap \text{CIO}(\langle q_n, z_n \rangle)$ such that $\text{Nodes}(\tilde{c}) \subseteq N \cup \{A_{\text{stop}}\}$. But then $\tilde{\eta}$ is \mathfrak{T} -complete and $\tilde{\eta} \xrightarrow{\surd} \langle q_n, z_n \rangle$ is a \mathfrak{T} -path and its projection is π .

□

Lemma 9 (Lemma 1). Let \mathcal{A} be a CA, $\mathcal{Z} = (Z, \text{CIO}_{\surd}, \delta, Z_0, Z_F)$ a DFA for α , q in \mathcal{A} , node-sets $N \subseteq \mathcal{N}$ and ASL state formulae Φ . Then, the following statements are equivalent:

- (a) $q \models \mathbb{E}_N \langle\langle \alpha \rangle\rangle \Phi$
- (b) $\langle q, z_0 \rangle \models \mathbb{E}_{N \cup \{A_{\text{stop}}\}} \diamond (\alpha_{\Phi} \wedge \text{accept})$
- (c) There exists a finite-memory N -strategy \mathfrak{S} for \mathcal{A} that is winning for $\langle q, \langle\langle \alpha \rangle\rangle \Phi \rangle$

Proof. “(a) \implies (b)” : Suppose that $q \models \mathbb{E}_N \langle\langle \alpha \rangle\rangle \Phi$ and that \mathfrak{S} is an N-strategy for \mathcal{A} that is winning for $\langle q, \langle\langle \alpha \rangle\rangle \Phi \rangle$. The goal is to define a corresponding $(N \cup \{A_{stop}\})$ -strategy \mathfrak{T} for $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$. Given a finite execution $\tilde{\eta}$ in the product we take its \mathcal{A} -projection $\eta \stackrel{\text{def}}{=} \text{proj}_{\mathcal{A}}(\tilde{\eta})$ and define

$$\mathfrak{T}(\tilde{\eta}) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{S}(\eta) & : \text{if } stop \notin \mathfrak{S}(\eta) \\ (\mathfrak{S}(\eta) \setminus \{stop\}) \cup \{c_{stop}\} & : \text{otherwise.} \end{cases}$$

Then, \mathfrak{T} and \mathfrak{S} are related as required in part (v) of Lemma 8. Hence, \mathfrak{T} is winning for $\langle\langle q, z_0 \rangle, \diamond(\alpha_{\Phi} \wedge \text{accept}) \rangle\rangle$ by parts (iii) and (v) of lemma 8.

“(b) \implies (c):” Suppose $\langle q, z_0 \rangle \models \mathbb{E}_{N \cup \{A_{stop}\}} \diamond(\alpha_{\Phi} \wedge \text{accept})$. By part (a) of lemma 7 there is a memoryless $(N \cup \{A_{stop}\})$ -strategy \mathfrak{T} for $\mathcal{A} \bowtie_{N, \Phi} \mathcal{Z}$ that is winning for $\langle\langle q, z_0 \rangle, \diamond(\alpha_{\Phi} \wedge \text{accept}) \rangle\rangle$. We now define a finite-memory N-strategy $\mathfrak{M} = (\text{Modes}, \Delta, \mu, m_0)$ for \mathcal{A} as follows. The set of modes agrees with the state-space of \mathcal{Z} , i.e., $\text{Modes} = \mathcal{Z}$. The decision function μ is given by:

$$\mu(q, z) \stackrel{\text{def}}{=} \begin{cases} \mathfrak{T}(\langle q, z \rangle) & : \text{if } c_{stop} \notin \mathfrak{T}(\langle q, z \rangle) \\ (\mathfrak{T}(\langle q, z \rangle) \setminus \{c_{stop}\}) \cup \{stop\} & : \text{otherwise.} \end{cases}$$

The transition relation Δ is defined by $\Delta(z, q \xrightarrow{c} p) \stackrel{\text{def}}{=} \delta(z, c)$.

It remains to show that \mathfrak{M} is winning for $\langle q, \langle\langle \alpha \rangle\rangle \Phi \rangle$. In fact, \mathfrak{T} and \mathfrak{M} are related as in part (v) of lemma 8. Again, applying parts (iii) and (v) of lemma 8 we get that \mathfrak{M} is winning for $\langle q, \langle\langle \alpha \rangle\rangle \Phi \rangle$.

The implication (c) \implies (a) is obvious. □

Lemma 10 (Lemma 2). Let \mathcal{A} be a CA, $\mathcal{Z} = (Z, \text{CIO}_{\surd}, \delta, Z_0, Z_F)$ a DFA for α , q in \mathcal{A} , node-sets $N \subseteq \mathcal{N}$ and ASL state formulae Φ . Then, the following statements are equivalent:

- (a) $q \models \mathbb{E}_N \llbracket \alpha \rrbracket \Phi$
- (b) $\langle q, z_0 \rangle \models \mathbb{E}_{N \cup \{A_{stop}\}} \square(\text{accept} \rightarrow \alpha_{\Phi})$
- (c) there exists a finite memory N-strategy \mathfrak{S} which is winning for $\langle q, \llbracket \alpha \rrbracket \Phi \rangle$

Proof. Using parts (iv) and (v) of lemma 8, the argument is analogous to the proof of lemma 9. □

Lemma 11 (Lemma 3). Let \mathfrak{S}_{stop} be the memoryless N-strategy given by $\mathfrak{S}_{stop}(q) = \{stop\} \cup \{c \in \text{CIO} : \text{Nodes}(c) \cap N = \emptyset\}$ for all states q . Then, for each \surd -free stream expression α and state q we have:

$$q \models \mathbb{E}_N \llbracket \alpha \rrbracket \Phi \text{ iff } \mathfrak{S}_{stop} \text{ is winning for } \langle q, \llbracket \alpha \rrbracket \Phi \rangle.$$

Proof. The implication \implies is obvious by the semantics for the modality \mathbb{E}_N . Suppose now that $q \models \mathbb{E}_N \llbracket \alpha \rrbracket \Phi$. The goal is to show that \mathfrak{S}_{stop} is a winning strategy for $\langle q, \llbracket \alpha \rrbracket \Phi \rangle$. We pick a winning strategy \mathfrak{T} for $\langle q, \llbracket \alpha \rrbracket \Phi \rangle$. That is, $\pi \models \llbracket \alpha \rrbracket \Phi$ for all \mathfrak{T} -paths π that start in state q . By definition of \mathfrak{S}_{stop} we get that for each incomplete \mathfrak{S}_{stop} -execution $\eta = q_0 \xrightarrow{c_1} \dots \xrightarrow{c_i} q_i$ we have:

$$\mathfrak{S}_{stop}(q_i) \cap \text{CIO}(q_i) \subseteq \mathfrak{T}(\eta)$$

Hence, all incomplete executions in $\text{Exec}_{\text{fin}}(q, \mathfrak{S}_{stop})$ are prefixes of \mathfrak{T} -executions. Thus, if $\eta \in \text{Exec}_{\text{fin}}(q, \mathfrak{S}_{stop})$ and η is an incomplete \mathfrak{T} -execution starting in q then we have:

if $\text{ios}(\eta) \in \text{IOS}(\alpha)$ implies $p \models \Phi$ where p is the last state of η .

In particular, this yields $\pi \models \llbracket \alpha \rrbracket \Phi$ for all infinite \mathfrak{S}_{stop} -paths that start in q . As α is \surd -free, none of the I/O-streams in $\text{IOS}(\alpha)$ contains the termination symbol \surd . Hence, for each finite \mathfrak{S}_{stop} -path π we have $\text{ios}(\pi) \notin \text{IOS}(\alpha)$ and therefore $\pi \models \llbracket \alpha \rrbracket \Phi$. \square

For the computation of $\text{Sat}_{\text{fair}}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ we define the following function:

$$\Gamma(N, \Phi, P) \stackrel{\text{def}}{=} \text{Sat} \left(\mathbb{E}_N(\Box(\Phi \wedge \neg P) \wedge \bigvee_{c \in \text{CIO}(N)} (\Box \diamond a_c \wedge \forall \llbracket c \rrbracket P)) \right)$$

where a_c is a new atomic proposition to reason about the enabledness of a concurrent I/O-operation c in a certain state q such that $a_c \in L(q)$ iff $c \in \text{CIO}(q)$. In other words, $\Gamma(N, \Phi, P)$ computes the set of all states $q \in Q$ such that there is a strategy ensuring to stay in $\text{Sat}(\Phi) \setminus P$ forever. But at the same time there exists a concurrent I/O-operation c , consisting of N -controllable nodes only, which is infinitely often enabled and guarantees to move into a P state in the next step once it has been chosen.

Lemma 12 (Until with fairness). Let \mathcal{A} be a CA, $N \subseteq \mathcal{N}$ a node-set and let Φ_1 and Φ_2 be *ASL* state formulae. Assuming N -fairness, algorithm 3 correctly returns the set $\text{Sat}_{\text{fair}}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ and the computed memoryless N -strategy \mathfrak{S} is winning for all states $q \in \text{Sat}_{\text{fair}}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$ and *ASL* path formula $(\Phi_1 \cup \Phi_2)$.

Proof. Let $P := P_j$ be the set of states and \mathfrak{S} be the N -strategies returned by algorithm 3. We will show the following:

$$q \in P \Leftrightarrow q \models_{\text{fair}} \mathbb{E}_N(\Phi_1 \cup \Phi_2).$$

“ \Rightarrow ”: Let $q \in P_{i+1} \setminus P_i$. Then q came into $P_{i+1} \subseteq P$ because

- a) Either $q \in P_0 = \text{Sat}(\Phi_2)$, then $q \models_{\text{fair}} \mathbb{E}_N(\Phi_1 \cup \Phi_2)$.
- b) Or $q \in \text{Sat}(\Phi_1) \cap \text{Pre}(P_i, N)$. Let

$$Q' := \bigcup_{c \in \mathfrak{S}(q)} \text{Post}[c](q)$$

be the set of successor states. But then $q \models \Phi_1$, and $Q' \subseteq P_i$ and we may apply the same arguments for all states $q' \in Q'$.

Algorithm 3 Algorithm for computing $Sat_{fair}(\mathbb{E}_N(\Phi_1 \cup \Phi_2))$;

```

P0 := Sat(Φ2);
i := 0;
repeat
  P := Pi;
  repeat
    Pi+1 := Pi ∪ (Sat(Φ1) ∩ Pre(Pi, N));
    for all states p ∈ Pi+1 \ Pi do
      S(p) := {c ∈ CIO : Nodes(c) ∩ N = ∅ ∨ ∅ ≠ Post[c](p) ⊆ Pi};
    end for
    i := i + 1;
  until Pi = Pi-1;
  Pi+1 := Pi ∪ Γ(N, Φ1, Pi);
  for all states p ∈ Pi+1 \ Pi do
    S(p) := {c ∈ CIO : Nodes(c) ∩ N = ∅ ∨ ∅ ≠ Post[c](p) ⊆ Pi+1};
  end for
  i := i + 1;
until Pi = P;
for all states p ∈ (Q \ Pi) ∪ Sat(Φ2) do
  S(p) := CIO ∪ {stop};
end for
return Pi;

```

(* P_i = Sat_{fair}(E_N(Φ₁ ∪ Φ₂)) *)

c) Or $q \in \Gamma(N, \Phi_1, P_i)$. Let again

$$Q' := \bigcup_{c \in \mathfrak{S}(q)} Post[c](q)$$

be the set of successor states. But then $q \models \Phi_1$, and $Q' \subseteq P_{i+1}$ and we may apply the same arguments for all states $q' \in Q'$.

“ \Leftarrow ”: Let us assume $q \notin P$ (i.e. $q \notin P_i$ for all $i \in \mathbb{N}$). But then,

$$q \notin Sat\left(\mathbb{E}_N\left(\left(\Phi_1 \cup P\right) \vee \left(\Box(\Phi_1 \wedge \neg P) \wedge \bigvee_{c \in CIO(N)} \Box \Diamond(a_c \wedge \forall [[c]] P)\right)\right)\right)$$

Hence, for all N-strategies \mathfrak{S} exists a path $\pi \in Paths(q, \mathfrak{S})$ such that

$$\pi \not\models (\Phi_1 \cup P) \vee ((\Box(\Phi_1 \wedge \neg P) \wedge \bigvee_{c \in CIO(N)} \Box \Diamond(a_c \wedge \forall [[c]] P))).$$

Let $\pi = q_0 \xrightarrow{c_1} q_1 \xrightarrow{c_2} \dots$ be such a path in $Paths(q, \mathfrak{S})$. Then π fulfills at least one of the following two conditions:

- i) $\pi \models \neg(\Phi_1 \cup P) \wedge \Diamond(\neg\Phi_1 \vee P)$. Let q_i be the first state of π such that $q_i \not\models \Phi_1$. In this case $\pi \downarrow i$ can be extended to a N-fair path, which violates $(\Phi_1 \cup P)$.

ii) $\pi \models \neg(\Phi_1 \cup P) \wedge \bigwedge_{c \in \text{CIO}(\mathcal{N})} \diamond \Box (\alpha_c \rightarrow \exists \langle\langle c \rangle\rangle \neg P)$. In this case π itself is a \mathcal{N} -fair path violating $(\Phi_1 \cup P)$.

Consequently for all \mathcal{N} -strategies \mathcal{S} there exists a path $\pi \in \text{FairPaths}_{\langle \mathcal{N}, \mathcal{S} \rangle}(q)$ such that $\pi \not\models (\Phi_1 \cup \Phi_2)$ and $q \not\models_{\text{fair}} \mathbb{E}_{\mathcal{N}}(\Phi_1 \cup \Phi_2)$. \square

Lemma 13 (Lemma 4). Let \mathcal{A} be a CA, $\mathcal{N} \subseteq \mathcal{N}$ a node-set and let Φ_1 and Φ_2 be *ASL* state formulae. Then, the following observation holds $q \models_{\text{fair}} \mathbb{E}_{\mathcal{N}}(\Phi_1 \mathcal{R} \Phi_2)$ iff $q \models \mathbb{E}_{\mathcal{N}}(\Phi_1 \mathcal{R} \Phi_2)$ for all states $q \in Q$.

Proof. “ \Leftarrow ”: This is obvious since $\text{FairPaths}_{\langle \mathcal{N}, \mathcal{S} \rangle}(q) \subseteq \text{Paths}(q, \mathcal{S})$.

“ \Rightarrow ”: Let us assume that $q \models_{\text{fair}} \mathbb{E}_{\mathcal{N}}(\Phi_1 \mathcal{R} \Phi_2)$ but $q \not\models \mathbb{E}_{\mathcal{N}}(\Phi_1 \mathcal{R} \Phi_2)$.

Let \mathcal{S} be an \mathcal{N} -strategy s.t. $\pi \in \text{FairPaths}_{\langle \mathcal{N}, \mathcal{S} \rangle}(q)$ implies that $\pi \models (\Phi_1 \cup \Phi_2)$, and $\pi = q_0 \xrightarrow{c_1} \dots \in \text{Paths}(q, \mathcal{S})$ a path s.t. $\pi \not\models (\Phi_1 \mathcal{R} \Phi_2)$. This path π can not be \mathcal{N} -fair (i.e. $\pi \in \text{FairPaths}_{\langle \mathcal{N}, \mathcal{S} \rangle}(q)$), since this contradicts our assumption.

Let q_i be the first state of π s.t. $q_i \not\models \Phi_2$. Then the execution $\eta = \pi \downarrow i$ can be extended to a \mathcal{N} -fair path $\pi' \in \text{Paths}(q, \mathcal{S})$. Obviously $\pi' \not\models (\Phi_1 \mathcal{R} \Phi_2)$, which again contradicts our assumption. The extension of η is done in the following manner. Whenever there is a $c \in \text{CIO}(q'_j) \cap \mathcal{S}(\eta \dot{\rightarrow} \dots \dot{\rightarrow} q'_j)$ with $\emptyset \neq \text{Nodes}(c) \subseteq \mathcal{N}$ then append (one of) the transition(s) $q'_j \xrightarrow{c} q'_{j+1}$ to the execution η , otherwise chose one of the enabled I/O-operations $c \in \text{CIO}(q'_j)$. Once the execution becomes maximal the corresponding path is finite and therefore \mathcal{N} -fair (see condition (1) of definition 6). Otherwise the resulting path becomes \mathcal{N} -fair, too (see condition (2) of definition 6). \square

Lemma 14 (Lemma 5). Let \mathcal{A} be a CA, $\mathcal{N} \subseteq \mathcal{N}$ a node-set, α a regular I/O-stream expression, \mathcal{Z} a deterministic CA for α , and let Φ be *ASL* state formula. Then, the following observation holds for all states $q \in Q$:

- i) $q \models_{\text{fair}} \mathbb{E}_{\mathcal{N}} \langle\langle \alpha \rangle\rangle \Phi$ in \mathcal{A} iff $\langle q, z_0 \rangle \models_{\text{fair}} \mathbb{E}_{\mathcal{N} \cup \{\mathcal{A}_{\text{stop}}\}} \diamond (\text{accept} \wedge \alpha_{\Phi})$ in $\mathcal{A} \bowtie \mathcal{Z}$.
- ii) $q \models_{\text{fair}} \mathbb{E}_{\mathcal{N}} [\langle\langle \alpha \rangle\rangle] \Phi$ iff $\langle q, z_0 \rangle \models_{\text{fair}} \mathbb{E}_{\mathcal{N} \cup \{\mathcal{A}_{\text{stop}}\}} \Box (\text{accept} \rightarrow \alpha_{\Phi})$ in $\mathcal{A} \bowtie \mathcal{Z}$.

Proof. The proof comes from combining the arguments of lemma 9, 10, 12, and 13. \square