

## Deciding Confluence for a Simple Class of Relational Transducer Networks

Peer-reviewed author version

AMELOOT, Tom & VAN DEN BUSSCHE, Jan (2015) Deciding Confluence for a Simple Class of Relational Transducer Networks. In: THEORY OF COMPUTING SYSTEMS, 57 (4), p. 1038-1111.

DOI: 10.1007/s00224-015-9624-6

Handle: <http://hdl.handle.net/1942/20633>

# Deciding Eventual Consistency for a Simple Class of Relational Transducer Networks

Tom J. Ameloot and Jan Van den Bussche

## Abstract

Networks of relational transducers can serve as a formal model for declarative networking, focusing on distributed database querying applications. In declarative networking, a crucial property is eventual consistency, meaning that the final output does not depend on the message delays and reorderings caused by the network. Here, we show that eventual consistency is decidable when the transducers satisfy some syntactic restrictions, some of which have also been considered in earlier work on automated verification of relational transducers. This simple class of transducer networks computes exactly all distributed queries expressible by unions of conjunctive queries with negation.

## 1 Introduction

Declarative networking [16] is an approach by which distributed computations and networking protocols, as occurring in cloud computing, are modeled and programmed using formalisms based on Datalog. Recently, declarative networking formalisms are enjoying attention from the database theory community, so that now a number of models and languages are available with a formally defined semantics and initial investigations on their expressive power [5, 18, 13, 1, 6].

A major hurdle in using declarative methods for cloud computing is the nondeterminism inherent to such systems. This nondeterminism is typically due to the asynchronous communication between the compute nodes in a cluster or network. Accordingly, one of the challenges is to design distributed programs so that the same outputs can eventually be produced on the same inputs, no matter how messages between nodes have been delayed or received in different orders. When a program has this property, we say it is *eventually consistent* [22, 14, 15, 4]. Of course, eventual consistency is undecidable in general, and there is much recent interest in finding ways to guarantee it [4, 1].

In the present paper, we view eventual consistency as a confluence notion. On any fixed input, let  $J$  be the union of all outputs that can be produced during any possible execution of the distributed program. Then in our definition of eventual consistency, we require that for any two different outputs  $J_1 \subseteq J$  and  $J_2 \subseteq J$  resulting from two (partial) executions on the same input, the same output  $J$  can be produced in an extension of either partial execution. So, intuitively, the prior execution of the program will not prevent outputs from being produced if those outputs can be produced with another execution (on the same input).

In this paper, we consider clusters of compute nodes modeled as *relational transducers*, an established formal model for data-centric agents [3, 21, 11, 10, 12]. In particular, we consider relational transducers where the rules used by the nodes to send messages, to update their state relations, and to produce output, are unions of conjunctive queries with negation. This setting yields a clear model of declarative networking, given the affinity between conjunctive queries and Datalog. We thus believe our results also apply to other declarative networking formalisms, although in this paper we have not yet worked out these applications.

Our first main result is the identification of a number of syntactic restrictions on the rules used in the transducers, *not* so that eventual consistency always holds, but so that checking it becomes decidable. Informally, the restrictions comprise the following.

- The cluster must be recursion-free: the different rules among all local programs cannot be mutually recursive through positive subgoals. Recursive dependencies through negative subgoals are still allowed.
- The local programs must be inflationary: deletions from state relations are forbidden.
- The rules are message-positive: negation on message relations is forbidden.
- The state-update rules must satisfy a known restriction which we call “message-boundedness”. This restriction is already established in the verification of relational transducers: it was first identified under the name “input-boundedness” by Spielmann [21] and was investigated further by Deutsch et al. [11, 12].
- Finally, the message-sending rules must be “static” in the sense that they cannot depend on state relations; they can still depend on input relations and on received messages.

The last two restrictions are the most fundamental; in fact, even if just the last restriction is dropped and all the others are kept in place, the problem is already back to undecidable. The first three restrictions can probably be slightly relaxed without losing decidability, and indeed we just see our work as a step in the right direction. Consistency is not an easy problem to analyze.

The second result of our paper is an analysis of the expressive power of clusters of relational transducers satisfying our above five restrictions; let us call such clusters “simple”. Specifically, we show that simple clusters can compute *exactly* all distributed queries expressible by unions of conjunctive queries with negation, or equivalently, the existential fragment of first-order logic, without any further restrictions. So, this result shows that simple clusters form indeed a rather weak computational model, but not as weak as to be totally useless.

**Related work** The work most closely related to ours is that by Deutsch et al. on verification of communicating data-driven Web services [12]. The main differences between our works are the following. (i) In their setting, message buffers are ordered queues; in our setting, message buffers are unordered multisets. Unordered buffers model the asynchronous communication typical in cloud computing [15] where messages can be delivered out of order. (ii) In their

setting, to obtain decidability, message buffers are bounded and lossy; in our setting, they are unbounded and not lossy. (iii) In their setting, transducers are less severely restricted than in our setting. (iv) In their setting, clusters of transducers are verified for properties expressed in (first-order) linear temporal logic;<sup>1</sup> in our setting, we are really focusing on the property of eventual consistency. It is actually not obvious whether eventual consistency (in the way we define it formally) is a linear-time temporal property, and if it is, whether it is expressible in first-order linear temporal logic.

Also, this paper is a follow-up on our previous paper [6]. In our previous paper, we did not consider the problem of deciding eventual consistency; we simply assumed eventual consistency and were focusing on expressiveness issues. Moreover, while the distributed computing model used in our previous paper is also based on relational transducers, there are differences in the models. In the previous model, we were focusing on standard queries to databases, computed in a distributed fashion by distributing the database in an arbitrary way over the nodes of the network. In the present model, we directly consider distributed queries, i.e., the input to the query is a distributed database, and different distributions of the same dataset may yield different answers to the query. Furthermore, in the previous model, transducer programs are considered to be network-independent, and nodes communicate in an epidemic manner by spreading messages to their neighbors, who read them one at a time; in the present model, the network is given, different nodes can run different programs, and nodes can directly address their messages to specified nodes. The perspective taken in our previous paper is equally interesting but different; we have simply chosen here to focus on the present perspective because it is the one mostly assumed by other authors in the area of declarative networking.

This paper extends our conference paper [7] by detailing all proofs, and by fully characterizing the computational complexity of the decision problem.

**Organization** We start in Section 2 by giving preliminaries about common database constructs, relational transducers, and their networks. Section 3 formalizes consistency for networks, along with syntactic restrictions leading to so-called “simple” networks; related (un)decidability results are also presented. Section 4 shows that consistency of a simple network with multiple nodes can be reduced to consistency of a simple *single-node* network. Next, Section 5 establishes a small model property for simple single-node networks. This is used in Section 6 to give a procedure for deciding whether a simple single-node network is inconsistent, along with a NEXPTIME-completeness result for the complexity.

The expressiveness of simple networks, not necessarily single-node, is analyzed in Section 7. We conclude in Section 8.

## 2 Preliminaries

### 2.1 Database Concepts

We first recall some basic notions from database theory [2]. A *database schema* is a finite set  $\mathcal{D}$  of pairs  $(R, k)$  where  $R$  is a *relation name* and  $k \in \mathbb{N}$  is the

---

<sup>1</sup>Deutsch et al. can also verify branching-time temporal properties, but only when transducer states are propositional.

associated *arity* of  $R$ . A relation name is allowed to occur only once in a database schema. We often write a pair  $(R, k) \in \mathcal{D}$  as  $R^{(k)}$ . An arity of zero is also called *nullary*.

We assume some infinite universe **dom** of atomic data values. A *fact*  $\mathbf{f}$  is a pair  $(R, \bar{a})$ , often denoted as  $R(\bar{a})$ , where  $R$  is a relation name – also called *predicate* – and  $\bar{a}$  is a tuple of values over **dom**. A *database instance*  $I$  over a database schema  $\mathcal{D}$  is a finite set of facts such that for each  $R(a_1, \dots, a_k) \in I$  we have  $R^{(k)} \in \mathcal{D}$ . Let  $Z$  be a subset of relation names in  $\mathcal{D}$ . We write  $I|_Z$  to denote the restriction of  $I$  to the facts whose predicate is a relation name in  $Z$ . For a function  $h : \mathbf{dom} \rightarrow \mathbf{dom}$  we define  $h(I) = \{R(h(a_1), \dots, h(a_k)) \mid R(a_1, \dots, a_k) \in I\}$ . The *active domain* of  $I$ , denoted  $\text{adom}(I) \subseteq \mathbf{dom}$ , is the set of atomic data values that occur in  $I$ . We also use this notation for facts.

A *query*  $Q$  over input database schema  $\mathcal{D}$  and output database schema  $\mathcal{D}'$  is a partial function mapping database instances over  $\mathcal{D}$  to database instances over  $\mathcal{D}'$ . A special but common kind of query are those where the output database schema contains just one relation. A query  $Q$  is called *generic* if for all input instances  $I$  and all permutations  $h$  of **dom**, the query  $Q$  is also defined on the isomorphic instance  $h(I)$  and  $Q(h(I)) = h(Q(I))$ . We recall that a generic query  $Q$  is *domain-preserving*, in the sense that  $\text{adom}(Q(I)) \subseteq \text{adom}(I)$  for all input instances  $I$ . We use the word “query” in this text to mean generic query.

## 2.2 Multisets

A *multiset*  $m$  over a universe  $\mathcal{U}$  is a function that maps each element  $e$  of  $\mathcal{U}$  to a natural number  $m(e)$  that represents the number of times that  $e$  occurs in  $m$ . The set operators  $\cap$ ,  $\cup$ , and  $\setminus$  can be defined for multisets in a natural way. For two multisets  $m_1$  and  $m_2$ , we write  $m_1 \sqsubseteq m_2$  to denote that  $m_1(e) \leq m_2(e)$  for each  $e \in \mathcal{U}$ . For a multiset  $m$ , we write  $\text{set}(m)$  to denote the collapse of  $m$  to a set in which we put only the elements of  $\mathcal{U}$  with multiplicity at least 1. Lastly, when  $m$  is given by a more complicated expression, we will write  $\text{num}(e, m)$  to denote the count of  $e$  in  $m$ .

## 2.3 Unions of Conjunctive Queries

We now recall the query language *unions of conjunctive queries with (safe) negation*, abbreviated UCQ<sup>−</sup>. This language is equivalent to the existential fragment of first-order logic [2]. It will be convenient to use a slightly unconventional formalization of conjunctive queries.

We assume an infinite universe **var** of variables. We will use typewriter font for variables. An *atom* is of the form  $R(\mathbf{u}_1, \dots, \mathbf{u}_k)$  where  $\mathbf{u}_i \in \mathbf{var}$  for each  $i \in \{1, \dots, k\}$ . A *literal* is an atom, or an atom with “ $\neg$ ” prepended; these literals are respectively called *positive* and *negative*.

A *conjunctive query* (or simply *rule*)  $\varphi$  is a four-tuple  $(\text{head}^\varphi, \text{pos}^\varphi, \text{neg}^\varphi, \text{non}^\varphi)$  where  $\text{head}^\varphi$  is an atom, and  $\text{pos}^\varphi$  and  $\text{neg}^\varphi$  are sets of atoms, and  $\text{non}^\varphi$  is a set of nonequalities of the form  $(\mathbf{u} \neq \mathbf{v})$  with  $\mathbf{u}, \mathbf{v} \in \mathbf{var}$ . Note that  $\text{neg}^\varphi$  is a set of atoms, and not negative literals. We call  $\text{head}^\varphi$ ,  $\text{pos}^\varphi$ , and  $\text{neg}^\varphi$  respectively the “head atom”, the “positive body atoms”, and the “negative body atoms”. Let  $\text{var}(\varphi)$  denote all variables that occur in  $\varphi$ . Let  $\text{free}(\varphi)$  denote all free variables of  $\varphi$  (occurring in the head), and let us abbreviate  $\text{bound}(\varphi) = \text{var}(\varphi) \setminus \text{free}(\varphi)$ . Bound variables can be thought of as being existentially quantified. As a safety

restriction, we require that all variables of  $head^\varphi$ ,  $neg^\varphi$  and  $non^\varphi$  occur in  $pos^\varphi$ . Note, nonequalities can be simulated by applying negation to an equality relation  $=$  that would have to be provided in every context where the rule is used, but for technical convenience we will immediately consider  $\neq$  to be a primitive in our language.

A rule  $\varphi$  may be written in the conventional syntax. For example, if  $head^\varphi = T(u, v)$ ,  $pos^\varphi = \{R(u, v)\}$ ,  $neg^\varphi = \{S(v)\}$ , and  $non^\varphi = \{u \neq v\}$ , then we may write  $\varphi$  as

$$T(u, v) \leftarrow R(u, v), \neg S(v), u \neq v.$$

The ordering of atoms and nonequalities in the body is immaterial. We will often refer to the literals of the body more directly, by prepending the symbol “ $\neg$ ” to the negative body atoms. For the previous example, the body literals are  $R(u, v)$  and  $\neg S(v)$ .

A rule  $\varphi$  is said to be *over* a database schema  $\mathcal{D}$  if for each atom  $R(u_1, \dots, u_k) \in \{head^\varphi\} \cup pos^\varphi \cup neg^\varphi$  we have  $R^{(k)} \in \mathcal{D}$ . A *valuation* for  $\varphi$  is a total function  $V : var(\varphi) \rightarrow \mathbf{dom}$ . The *application* of  $V$  to an atom  $R(u_1, \dots, u_k)$  of  $\varphi$ , denoted  $V(R(u_1, \dots, u_k))$ , results in the *fact*  $R(a_1, \dots, a_k)$  with  $a_i = V(u_i)$  for each  $i \in \{1, \dots, k\}$ . We will also use this notation for applying  $V$  to a set of atoms, which results in a set of facts. Let  $I$  be a database instance over  $\mathcal{D}$ . The valuation  $V$  is said to be *satisfying for  $\varphi$  on  $I$*  if  $V(pos^\varphi) \subseteq I$ ,  $V(neg^\varphi) \cap I = \emptyset$ , and  $V(u) \neq V(v)$  for each  $(u \neq v) \in non^\varphi$ . In that case,  $\varphi$  is said to *derive* the fact  $V(head^\varphi)$ . The *result of  $\varphi$  applied to  $I$* , denoted  $\varphi(I)$ , is defined as the set of facts derived by all possible satisfying valuations for  $\varphi$  on  $I$ . Note that rules can only define generic queries.

A *union of conjunctive queries* is a finite set  $\Phi$  of conjunctive queries that all have the same predicate and arity for the head atom. The resulting language is denoted as  $UCQ^\neg$ , and  $\Phi$  will also be called a *UCQ $^\neg$ -program*. Let  $I$  be a database instance. The *result of  $\Phi$  applied to  $I$* , denoted  $\Phi(I)$ , is defined as  $\bigcup_{\varphi \in \Phi} \varphi(I)$ . If  $\Phi = \emptyset$  then always  $\Phi(I) = \emptyset$ .

## 2.4 Distributed Databases and Queries

We now formalize how input data is distributed across a network and define a notion of queries over this data. A *network*  $\mathcal{N}$  is a finite, nonempty set of *nodes*, which are values in  $\mathbf{dom}$ . A *distributed database schema*  $\mathcal{E}$  is a pair  $(\mathcal{N}, \eta)$  where  $\mathcal{N}$  is a network, and  $\eta$  is a function that maps each  $x \in \mathcal{N}$  to an ordinary database schema. A *distributed database instance*  $H$  over schema  $\mathcal{E}$  is a function that assigns to each node  $x \in \mathcal{N}$  an ordinary database instance over the local schema  $\eta(x)$ .

Let  $\mathcal{F}$  be another distributed database schema over the *same* network as  $\mathcal{E}$ . A *distributed query*  $\mathcal{Q}$  over input schema  $\mathcal{E}$  and output schema  $\mathcal{F}$  is a function that maps instances over  $\mathcal{E}$  to instances over  $\mathcal{F}$ .

## 2.5 Transducers

The computation on a single node of a network is formalized by means of relational transducers [3, 21, 12, 11, 10, 6, 23]. First, a *transducer schema*  $\Upsilon$  is a tuple  $(\Upsilon_{in}, \Upsilon_{out}, \Upsilon_{msg}, \Upsilon_{mem}, \Upsilon_{sys})$  of database schemas, called respectively “input”, “output”, “message”, “memory”, and “system”. A relation name can

occur in at most one database schema of  $\Upsilon$ . We fix  $\Upsilon_{\text{sys}}$  to always contain two unary relations `Id` and `All`. A *transducer state* for  $\Upsilon$  is a database instance over  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}} \cup \Upsilon_{\text{sys}}$ .

An *relational transducer*  $\Pi$  over  $\Upsilon$  is a collection of queries, where each query has the input schema  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{msg}} \cup \Upsilon_{\text{mem}} \cup \Upsilon_{\text{sys}}$ :

- for each  $R^{(k)} \in \Upsilon_{\text{out}}$  there is a query  $\mathcal{Q}_{\text{out}}^R$  having output schema  $\{R^{(k)}\}$ ;
- for each  $R^{(k)} \in \Upsilon_{\text{mem}}$  there are queries  $\mathcal{Q}_{\text{ins}}^R$  and  $\mathcal{Q}_{\text{del}}^R$  both having output schema  $\{R^{(k)}\}$ ;
- for each  $R^{(k)} \in \Upsilon_{\text{msg}}$  there is a query  $\mathcal{Q}_{\text{snd}}^R$  having output schema  $\{R^{(k+1)}\}$ ;

These queries will form the internal mechanism that a node uses to update its local storage and to send messages. The reason for the incremented arity in the message queries is that the extra component will be used to indicate the addressee, as will be explained in the next section.

Let  $\Pi$  be a transducer over schema  $\Upsilon$ . A *local transition* of  $\Pi$  is a 4-tuple  $(I, I_{\text{rcv}}, J, J_{\text{snd}})$ , also denoted as  $I, I_{\text{rcv}} \rightarrow J, J_{\text{snd}}$ , where  $I$  and  $J$  are transducer states for  $\Upsilon$ ,  $I_{\text{rcv}}$  is an instance over  $\Upsilon_{\text{msg}}$  and  $J_{\text{snd}}$  is an instance over  $\Upsilon_{\text{msg}}$  but where each fact has one extra component, such that (denoting  $I' = I \cup I_{\text{rcv}}$ ):

$$\begin{aligned} J|_{\Upsilon_{\text{in}}, \Upsilon_{\text{sys}}} &= I|_{\Upsilon_{\text{in}}, \Upsilon_{\text{sys}}}; \\ J|_{\Upsilon_{\text{out}}} &= I|_{\Upsilon_{\text{out}}} \cup \bigcup_{R^{(k)} \in \Upsilon_{\text{out}}} \mathcal{Q}_{\text{out}}^R(I'); \\ J|_{\Upsilon_{\text{mem}}} &= \bigcup_{R^{(k)} \in \Upsilon_{\text{mem}}} (I|_R \cup R^+(I')) \setminus R^-(I') \\ J_{\text{snd}} &= \bigcup_{R^{(k)} \in \Upsilon_{\text{msg}}} \mathcal{Q}_{\text{snd}}^R(I'), \end{aligned}$$

where, following the presentation in [23],

$$\begin{aligned} R^+(I') &= \mathcal{Q}_{\text{ins}}^R(I') \setminus \mathcal{Q}_{\text{del}}^R(I'); \text{ and,} \\ R^-(I') &= \mathcal{Q}_{\text{del}}^R(I') \setminus \mathcal{Q}_{\text{ins}}^R(I'). \end{aligned}$$

Intuitively, on the receipt of message facts  $I_{\text{rcv}}$ , a local transition updates the old transducer state  $I$  to new transducer state  $J$  and sends the facts in  $J_{\text{snd}}$ . When compared to  $I$ , in  $J$  potentially more output facts are produced; and the update semantics for each memory relation  $R$  adds the facts produced by *insertion* query  $\mathcal{Q}_{\text{ins}}^R$ , removes the facts produced by *deletion* query  $\mathcal{Q}_{\text{del}}^R$ , and there is no-op semantics in case a fact is both added and removed at the same time [21]. Output facts can not be removed. Note that local transitions are deterministic in the following sense: if  $I, I_{\text{rcv}} \rightarrow J, J_{\text{snd}}$  and  $I, I_{\text{rcv}} \rightarrow J', J'_{\text{snd}}$  then  $J = J'$  and  $J_{\text{snd}} = J'_{\text{snd}}$ .

For the current paper, we immediately restrict attention to transducers whose queries are specified with UCQ<sup>+</sup>. This results in a rule-based formalism to express the computations, following the idea behind declarative networking [16].

## 2.6 Derivation Trees

We want to formally describe *how* a fact is derived by a transducer, i.e., we want to make visible what rules and valuations are used. To explain a fact, in some cases it suffices to give a so-called *derivation pair*  $(\varphi, V)$ , consisting of a rule  $\varphi$  and a satisfying valuation. In other cases, we want to explain all facts that are recursively needed by the satisfying valuation, i.e., the facts  $V(\text{pos}^\varphi)$ . For this purpose, we use *derivation trees*, and this is formalized below.

Let  $\Pi$  be a transducer over a schema  $\Upsilon$ . A (full) *derivation tree*  $\mathcal{T}$  of  $\Pi$  is a tuple  $(\text{nodes}^\mathcal{T}, \text{edges}^\mathcal{T}, \text{rule}^\mathcal{T}, \text{val}^\mathcal{T}, \text{lit}^\mathcal{T})$  where

- $\text{nodes}^\mathcal{T}$  and  $\text{edges}^\mathcal{T}$  are respectively the nodes and parent-child edges that together form a tree;
- $\text{rule}^\mathcal{T}$  is a function that maps each internal node  $x \in \text{nodes}^\mathcal{T}$  to a rule  $\text{rule}^\mathcal{T}(x)$  of  $\Pi$ ;
- $\text{val}^\mathcal{T}$  is a function that maps each internal node  $x \in \text{nodes}^\mathcal{T}$  to a valuation  $\text{val}^\mathcal{T}(x)$  for  $\text{rule}^\mathcal{T}(x)$  such that the nonequalities are satisfied; and,
- $\text{lit}^\mathcal{T}$  is a function that maps each non-root node  $x \in \text{nodes}^\mathcal{T}$  to a literal  $\text{lit}^\mathcal{T}(x)$  in the body of  $\text{rule}^\mathcal{T}(y)$  where  $y$  is the parent of  $x$ ,

subject to the additional constraints:

- for each internal node  $x \in \text{nodes}^\mathcal{T}$ , for each literal  $\mathbf{l}$  in the body of rule  $\text{rule}^\mathcal{T}(x)$ , there is precisely one child  $y$  of  $x$  such that  $\text{lit}^\mathcal{T}(y) = \mathbf{l}$ ;
- for each non-root node  $x \in \text{nodes}^\mathcal{T}$ , if  $\text{lit}^\mathcal{T}(x)$  is a database literal, or if  $\text{lit}^\mathcal{T}(x)$  is negative, then  $x$  must be a leaf; and,
- for all non-root internal nodes  $x \in \text{nodes}^\mathcal{T}$ , having a parent  $y$ , applying valuation  $\text{val}^\mathcal{T}(x)$  to the head of rule  $\text{rule}^\mathcal{T}(x)$  results in the same fact as applying the parent valuation  $\text{val}^\mathcal{T}(y)$  to the (positive) atom inside literal  $\text{lit}^\mathcal{T}(x)$ .

For each internal node  $x$  of  $\mathcal{T}$ , we write  $\text{fact}^\mathcal{T}(x)$  to denote the fact  $\text{val}^\mathcal{T}(x)(\mathbf{a})$ , where  $\mathbf{a}$  is the head of  $\text{rule}^\mathcal{T}(x)$ . For a leaf node  $y$  with parent  $x$ , we write  $\text{fact}^\mathcal{T}(y)$  to denote the fact  $\text{val}^\mathcal{T}(x)(\mathbf{a})$ , where  $\mathbf{a}$  is the atom inside the literal  $\text{lit}^\mathcal{T}(y)$ . We write  $\text{int}^\mathcal{T}$  to denote the set of internal nodes of  $\mathcal{T}$ .

From  $\text{nodes}^\mathcal{T}$  and  $\text{edges}^\mathcal{T}$  we can always uniquely identify the root node of  $\mathcal{T}$ , which we denote as  $\text{root}^\mathcal{T}$ . Let  $\mathbf{f}$  be a fact over a relation  $R^{(k)} \in \Upsilon_{\text{out}} \cup \Upsilon_{\text{msg}} \cup \Upsilon_{\text{mem}}$ . A derivation tree  $\mathcal{T}$  is said to be *for* fact  $\mathbf{f}$  if applying valuation  $\text{val}^\mathcal{T}(\text{root}^\mathcal{T})$  to the head of rule  $\text{rule}^\mathcal{T}(\text{root}^\mathcal{T})$  results in the fact  $\mathbf{f}$ .

### 2.6.1 Schedulings

To relate derivation trees to runs, we use the concept of schedulings. Formally, a *scheduling* for a derivation tree  $\mathcal{T}$  is a function  $\kappa$  that assigns to each internal node  $x$  of  $\mathcal{T}$  a nonzero natural number  $\kappa(x)$ , subject to the constraint that nodes always get strictly lower numbers than their ancestors. Intuitively,  $\kappa(x)$  represents the transition number of a run in which the rule  $\text{rule}^\mathcal{T}(x)$  should fire under valuation  $\text{val}^\mathcal{T}(x)$ .



The *canonical scheduling* of  $\mathcal{T}$ , denoted  $\kappa^{\mathcal{T}}$ , is the (unique) scheduling for which there is at least one internal node  $x$  such that  $\kappa^{\mathcal{T}}(x) = 1$ , and for all parent-child edges  $(x, y)$  we have  $\kappa^{\mathcal{T}}(x) = \kappa^{\mathcal{T}}(y) + 1$ . Intuitively, the canonical scheduling executes the derivations of  $\mathcal{T}$  as tightly as possible at the beginning of a run.

## 2.7 Transducer Networks

We now formalize a network of compute nodes. A *transducer network*  $\mathcal{N}$  is a triple  $(\mathcal{N}, \Upsilon, \Pi)$  where  $\mathcal{N}$  is a network,  $\Upsilon$  is a function that maps each node  $x \in \mathcal{N}$  to a transducer schema, and  $\Pi$  is a function that maps each node  $x \in \mathcal{N}$  to a transducer over the schema  $\Upsilon(x)$ . For technical convenience, we assume that all transducer schemas use the same message relations. This is not really a restriction because the transducers are not obliged to use all message relations. We make no further assumptions about how names for input, output and memory relations might be shared by several nodes.

### 2.7.1 Distributed Schemas

Naturally, we can define the distributed input database schema  $in^{\mathcal{N}}$  for  $\mathcal{N}$  that maps each node  $x$  to the input schema of  $\Upsilon(x)$ . The distributed schemas  $out^{\mathcal{N}}$  and  $mem^{\mathcal{N}}$  can be defined similarly.

### 2.7.2 Operational Semantics

Any distributed database instance over  $in^{\mathcal{N}}$  can be given as input to  $\mathcal{N}$ . Let  $H$  be such an instance. Let  $\Upsilon_{\text{msg}}$  denote the shared message schema of  $\mathcal{N}$ . A *configuration of  $\mathcal{N}$  on  $H$*  is a pair  $\rho = (s, b)$  of functions  $s$  and  $b$  where for each  $x \in \mathcal{N}$ ,

- letting  $\mathcal{D}_1 = \Upsilon(x)_{\text{in}}$  and  $\mathcal{D}_2 = \Upsilon(x)_{\text{sys}}$ , function  $s$  maps  $x$  to a transducer state  $s(x)$  for  $\Upsilon(x)$  such that  $s(x)|_{\mathcal{D}_1} = H(x)$  and  $s(x)|_{\mathcal{D}_2} = \{\text{Id}(x)\} \cup \{\text{All}(y) \mid y \in \mathcal{N}\}$ ; and,
- $b$  maps  $x$  to a finite multiset of facts over the shared message schema of  $\mathcal{N}$ .

We call  $s$  the *state function* and  $b$  the *buffer function*. Intuitively, the instance  $H$  is used to initialize each node, and for each  $x \in \mathcal{N}$ , the system relations  $\text{Id}$  and  $\text{All}$  provide the local transducer  $\Pi(x)$  the identity of the node  $x$  it is running on and the identities of the other nodes. Next, the buffer function maps each  $x \in \mathcal{N}$  to the multiset of messages that have been sent to  $x$  but that have not yet been delivered to  $x$ . A multiset allows us to represent duplicates of the same message (sent at different times).

The *start configuration of  $\mathcal{N}$  on  $H$* , denoted  $start(\mathcal{N}, H)$ , is the unique configuration  $\rho = (s, b)$  where for each  $x \in \mathcal{N}$ , letting  $\mathcal{D} = \Upsilon(x)_{\text{out}} \cup \Upsilon(x)_{\text{mem}}$ , we have  $s(x)|_{\mathcal{D}} = \emptyset$  and  $b(x) = \emptyset$ .

We now describe the actual computation of the transducer network. A *global transition* of  $\mathcal{N}$  on input  $H$  is a 4-tuple  $(\rho_1, x, m, \rho_2)$ , also denoted as  $\rho_1 \xrightarrow{x, m} \rho_2$ , where  $x \in \mathcal{N}$ , and  $\rho_1 = (s_1, b_1)$  and  $\rho_2 = (s_2, b_2)$  are configurations of  $\mathcal{N}$  on  $H$  such that

- $m \sqsubseteq b_1(x)$  and there exists a  $J_{\text{snd}}$  such that

$$s_1(x), \text{set}(m) \rightarrow s_2(x), J_{\text{snd}}$$

is a local transition of transducer  $\Pi(x)$ ;

- for each  $y \in \mathcal{N} \setminus \{x\}$  we have  $s_1(y) = s_2(y)$ ;
- for  $y \in \mathcal{N} \setminus \{x\}$  we have  $b_2(y) = b_1(y) \cup J_{\text{snd}}^{\rightarrow y}$  (multiset union) and for  $x$  we have  $b_2(x) = (b_1(x) \setminus m) \cup J_{\text{snd}}^{\rightarrow x}$  (multiset union and difference) where  $J_{\text{snd}}^{\rightarrow z} = \{R(\bar{a}) \mid R(z, \bar{a}) \in J_{\text{snd}}\}$  for each  $z \in \mathcal{N}$ .

We call  $x$  the *active node* and  $m$  the *delivered messages*. Intuitively, in a global transition, we select an arbitrary node  $x$  and allow it to receive some arbitrary submultiset  $m$  from its message buffer. The messages in  $m$  are then delivered at node  $x$  (as a set, i.e., without duplicates) and  $x$  performs a local transition, in which it updates its memory and output relations, and possibly sends some new messages addressed to specific nodes (possibly itself). The first component of each message fact in  $J_{\text{snd}}$  is regarded as the addressee, and this component is projected away during the transfer of the message to the buffer of that addressee. Messages having an addressee outside the network are lost. If  $m = \emptyset$ , we call this global transition a *heartbeat* transition and otherwise we call it a *delivery* transition. A heartbeat transition corresponds to the real life situation in which a node does a computation step when a local timer goes off and no messages have been received from the network.

A *run*  $\mathcal{R}$  of a transducer network  $\mathcal{N}$  on distributed input database instance  $H$  is a *finite* sequence of global transitions  $\rho_i \xrightarrow{x_i, m_i} \rho_{i+1}$  for  $i = 1, 2, 3, \dots, n$ , with  $n \in \mathbb{N}$ , where  $\rho_1 = \text{start}(\mathcal{N}, H)$ , and the  $i^{\text{th}}$  transition with  $i \geq 2$  operates on the resulting configuration of the previous transition  $i - 1$ . We write  $\text{last}(\mathcal{R})$  to denote the last configuration reached by  $\mathcal{R}$ .

Note, when a node changes its output or memory relations during one global transition, then these changes are visible to that node only starting from the next global transition in which that node is active. Also, several facts can be delivered together during a transition, regardless of whether they were sent during different earlier transitions or during the same earlier transition.

We have not defined global transitions that are concurrent, i.e., global transitions in which multiple nodes simultaneously receive messages from their own message buffer and do a local transition. This can be simulated by multiple sequential global transitions: let the nodes become active in some arbitrary order, and each active node just reads its own message buffer. Because local transitions are deterministic, the nodes will update their state and send messages in the same way as they would during a concurrent transition.

### 2.7.3 Example

Here we give an example transducer network.

**Example 2.1.** Let  $\mathcal{N} = \{x, y\}$  be a network of two nodes. We define a transducer network  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$ . There are no memory relations in this example.

First, define  $\Upsilon(x)_{\text{in}} = \{A^{(1)}\}$ ,  $\Upsilon(x)_{\text{out}} = \{T^{(1)}\}$ ,  $\Upsilon(x)_{\text{msg}} = \{A_{\text{msg}}^{(1)}, B_{\text{msg}}^{(2)}\}$ , and  $\Upsilon(x)_{\text{mem}} = \emptyset$ . Transducer  $\Pi(x)$  is given as

$$A_{\text{msg}}(\mathbf{y}, \mathbf{u}) \leftarrow A(\mathbf{u}), \text{All}(\mathbf{y}), \neg \text{Id}(\mathbf{y}).$$

$$T(\mathbf{u}) \leftarrow B_{\text{msg}}(\mathbf{x}, \mathbf{u}), \text{Id}(\mathbf{x}).$$

Next, define  $\Upsilon(y)_{\text{in}} = \{B^{(2)}\}$ ,  $\Upsilon(y)_{\text{out}} = \{T^{(1)}\}$ ,  $\Upsilon(y)_{\text{msg}} = \Upsilon(x)_{\text{msg}}$  (shared messages), and  $\Upsilon(y)_{\text{mem}} = \emptyset$ . Transducer  $\Pi(y)$  is given as

$$B_{\text{msg}}(\mathbf{y}, \mathbf{u}, \mathbf{v}) \leftarrow B(\mathbf{u}, \mathbf{v}), \text{All}(\mathbf{y}), \neg \text{Id}(\mathbf{y}).$$

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}).$$

On any input distributed database instance  $H$  for  $\mathcal{N}$ , node  $x$  sends its local  $A$ -facts as  $A_{\text{msg}}$ -facts to  $y$ . Similarly,  $y$  sends its local  $B$ -facts as  $B_{\text{msg}}$ -facts to  $x$ . For a received  $B_{\text{msg}}$ -fact, node  $x$  outputs the second component in relation  $T$  if the first component is its identifier. Node  $y$  simply outputs all received  $A_{\text{msg}}$ -facts.  $\square$

## 2.8 Encoding

We specify how a transducer network can be given as input to a decision procedure. Let  $\mathcal{N}$  be a transducer network. The encoding is a sequence of transducers (and their schemas), one for each node of  $\mathcal{N}$ . For each node, (i) the transducer schema is represented by a sequence of (relname, type)-pairs, where relname is a relation name and the type indicates whether the relation is input, output, etc; and, (ii) the transducer itself is given by a sequence of rules that are written in full, like in Example 2.1.<sup>2</sup> We assume that the transducer schema only mentions relations effectively used by the rules. To represent the relation names and variables, binary numbers must be used, so that the number of bits is logarithmic in the total number of relations and variables respectively. Moreover, some small fixed alphabet of auxiliary characters needs to be used, to represent the type of relations in the transducer schema, and to separate the different components (schemas, transducers, rules, etc).

We write  $|\mathcal{N}|$  to denote the size of the encoding of  $\mathcal{N}$ .

## 3 Consistency

Let  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  be a transducer network. Let  $H$  be an input distributed database instance for  $\mathcal{N}$ . By the asynchronous nature of message delivery, different runs of  $\mathcal{N}$  on  $H$  can deliver messages in different orders. So, if a transducer at some node  $x \in \mathcal{N}$  applies negation too quickly, without having seen some crucial messages, we could accidentally produce a wrong output. Worse, output facts can never be retracted once they are produced. By contrast, transducer networks where such problems are not possible are called consistent.

Formally, we call  $\mathcal{N}$  *consistent on  $H$*  if for any two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  of  $\mathcal{N}$  on  $H$ , for every node  $x \in \mathcal{N}$ , for every output fact  $\mathbf{f}$  available at  $x$  in the last configuration of  $\mathcal{R}_1$ , there exists an extension  $\mathcal{R}_2'$  of  $\mathcal{R}_2$  such that  $\mathbf{f}$  is available at  $x$  in the last configuration of  $\mathcal{R}_2'$ . To rephrase, if during one run some node can produce an output, then for any run there exists an extension in which that

<sup>2</sup>The components of the body atoms have to be specified in full, because we need to describe which variables are used, and how they are potentially shared between atoms.

fact can be produced on that node too. Naturally, we call  $\mathcal{N}$  *consistent* if  $\mathcal{N}$  is consistent on all input distributed database instances. If  $\mathcal{N}$  is not consistent, we say that  $\mathcal{N}$  is *inconsistent*. Our definition of consistency is a formalization of the notion of “eventual consistency” [4, 15], but see also Section 8 for a discussion.

The transducer network given in Example 2.1 is consistent. Indeed, say, node  $x$  outputs a fact  $T(a)$  during a run. This means that  $x$  has received  $B_{\text{msg}}(x, a)$ , which was sent by node  $y$  based on an input fact  $B(x, a)$ . On the same input distributed database instance, consider now any run where  $x$  has not yet output  $T(a)$ . We can extend this run as follows. We do a global transition with active node  $y$ , so that  $y$  sends its input  $B$ -facts as  $B_{\text{msg}}$ -facts to  $x$ . One of these messages is  $B_{\text{msg}}(x, a)$ . Then, in a following global transition, we deliver  $B_{\text{msg}}(x, a)$  to  $x$ , and  $x$  again outputs  $T(a)$ . Similarly, we can argue that if the node  $y$  outputs a  $T$ -fact in one run, then any other run on the same input can be extended so that  $y$  outputs again this fact. Therefore the transducer network is consistent.

By contrast, consider the following example of a transducer network that is inconsistent.

**Example 3.1.** Let  $\mathcal{N} = \{x, y\}$  be a network. We define a transducer network  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  as follows. In this example, we do no deletions on memory relations, and we will only explicitly specify the insertions.

First, define  $\Upsilon(x)_{\text{in}} = \{A^{(1)}, B^{(1)}\}$ ,  $\Upsilon(x)_{\text{out}} = \emptyset$ ,  $\Upsilon(x)_{\text{msg}} = \{A_{\text{msg}}^{(1)}, B_{\text{msg}}^{(1)}\}$ , and  $\Upsilon(x)_{\text{mem}} = \emptyset$ . The node  $x$  sends its local  $A$ - and  $B$ -facts to the other node  $y$ . Transducer  $\Pi(x)$  is given as

$$\begin{aligned} A_{\text{msg}}(\mathbf{y}, \mathbf{u}) &\leftarrow A(\mathbf{u}), \text{All}(\mathbf{y}), \neg \text{Id}(\mathbf{y}). \\ B_{\text{msg}}(\mathbf{y}, \mathbf{u}) &\leftarrow B(\mathbf{u}), \text{All}(\mathbf{y}), \neg \text{Id}(\mathbf{y}). \end{aligned}$$

Next, define  $\Upsilon(y)_{\text{in}} = \emptyset$ ,  $\Upsilon(y)_{\text{out}} = \{T^{(1)}\}$ ,  $\Upsilon(y)_{\text{msg}} = \Upsilon(x)_{\text{msg}}$  (shared messages), and  $\Upsilon(y)_{\text{mem}} = \{B^{(1)}\}$ . Transducer  $\Pi(y)$  is given as:

$$\begin{aligned} B(\mathbf{u}) &\leftarrow B_{\text{msg}}(\mathbf{u}). \\ T(\mathbf{u}) &\leftarrow A_{\text{msg}}(\mathbf{u}), \neg B(\mathbf{u}). \end{aligned}$$

Now we show why  $\mathcal{N}$  is inconsistent. Let  $H$  be the following instance over  $\text{in}^{\mathcal{N}}$ :  $H(x) = \{A(1), B(1)\}$  and  $H(y) = \emptyset$ . There are two quite different runs possible, that we describe next. Suppose that both runs start with a global transition with active node  $x$ . This causes  $x$  to send both  $A_{\text{msg}}(1)$  and  $B_{\text{msg}}(1)$  to  $y$ . For the first run, in the second transition we deliver only  $A_{\text{msg}}(1)$  to  $y$ , which causes  $y$  to output  $T(1)$ . For the second run, in the second transition we deliver only  $B_{\text{msg}}(1)$  to  $y$ , which causes  $y$  to only create the memory fact  $B(1)$ . Now, the output fact  $T(1)$  can not be created in any extension of the second run because each time we deliver  $A_{\text{msg}}(1)$  to  $y$ , the presence of  $B(1)$  prevents  $T(1)$  from being created. These two runs show that  $\mathcal{N}$  is not consistent.  $\square$

### 3.1 Decision Problem

Since output facts can not be retracted once they are produced, it seems useful to know if a transducer network could be inconsistent. Formally, we have the

following *inconsistency decision problem*: given a transducer network  $\mathcal{N}$ , decide if  $\mathcal{N}$  is inconsistent (for some input). One can expect this problem to be undecidable in general. For this reason, we consider possible syntactical restrictions on transducer networks in Section 3.2, and Section 3.3 investigates their effect on decidability.

### 3.2 Syntactical Restrictions

We introduce several syntactical restrictions on individual transducers and on transducer networks as a whole.

Let  $\Pi$  be a transducer over a schema  $\Upsilon$ . For an individual rule  $\varphi$  of  $\Pi$ , we consider the following possible restrictions:

- We say that  $\varphi$  is *message-positive* if there are no message atoms in  $neg^\varphi$ . Note, this seems to be a natural constraint in our model because message delivery is asynchronous.
- We say that  $\varphi$  is *static* if  $pos^\varphi$  and  $neg^\varphi$  do not contain output or memory atoms.
- We say that  $\varphi$  is *message-bounded* if  $bound(\varphi) \subseteq A$  and  $bound(\varphi) \cap B = \emptyset$ , where  $A$  and  $B$  are respectively the set of variables of  $\varphi$  occurring in positive message atoms, and the set of variables of  $\varphi$  occurring in output or memory atoms. In words: every bound variable occurs in a positive message atom, and does not occur in output or memory atoms (positive or negative). This is an application of the more general notion of “input-boundedness” [21, 12, 11].<sup>3</sup>

We consider the following restrictions for transducer  $\Pi$ :

- We say that  $\Pi$  is *recursion-free* if there are no cycles in the *positive dependency graph* of  $\Pi$ , which is the graph having as vertices the relations of  $\Upsilon_{out} \cup \Upsilon_{msg} \cup \Upsilon_{mem}$  and there is an edge from relation  $R$  to relation  $S$  if  $S$  occurs positively in a rule for  $R$  in  $\Pi$ .
- We say that  $\Pi$  is *inflationary* if there are no rules for the deletion queries of memory relations. This means that  $\Pi$  can not delete memory facts once they are produced.

We call  $\Pi$  *simple* (for lack of a better name) if

- $\Pi$  is recursion-free and inflationary;
- all send rules are message-positive and static;<sup>4</sup> and,
- all insertion rules for output and memory relations are message-positive and message-bounded.

<sup>3</sup>We have replaced the term “input-boundedness” by “message-boundedness” because the word “input” has a different meaning in our text, namely, as the input that a transducer is locally initialized with.

<sup>4</sup>The restrictions considered by Deutsch et al. [11] for “input-rules”, which are closely related to our send rules, are a bit less restrictive. Roughly speaking, they still allow the use of nullary output and memory facts. It seems plausible that our results can be similarly extended.

Because input facts are never changed, note that static send rules always produce the same result on receipt of the same messages, independently of what output or memory facts might have been derived. Also, if  $\Pi$  is inflationary, memory and output relations basically behave in the same way. However, we preserve the difference between these two kinds of relations to retain the connection to the unrestricted transducer model and because memory relations are useful as a separate construct, namely, as relations used for computation but that don't belong to the final result.

Let  $\mathcal{N}$  be a transducer network. We present a restriction that we can impose on  $\mathcal{N}$  as a whole. Note that messages are the only way to introduce a dependency between different nodes of  $\mathcal{N}$ . Now, we say that  $\mathcal{N}$  is *globally recursion-free* if there are no cycles in the *positive message dependency graph* of  $\mathcal{N}$ , which is the graph having as vertices the (shared) message relations of  $\mathcal{N}$  and there is an edge from relation  $R$  to relation  $S$  if  $S$  occurs positively in a rule for  $R$  in some transducer of  $\mathcal{N}$ .

We call  $\mathcal{N}$  *simple* if

- all transducers of  $\mathcal{N}$  are simple; and,
- $\mathcal{N}$  is globally recursion-free.

The Examples 2.1 and 3.1 are simple transducer networks.

### 3.3 Results on Decidability

One of the difficulties of the inconsistency decision problem is that we need to verify a property of an infinite state system. Intuitively, there are infinitely many inputs and even for a fixed input there are infinitely many configurations because there is no bound on the size of the message buffer. As the following two propositions show, inconsistency for transducer networks is undecidable, even under several restrictions:

**Proposition 3.2.** Inconsistency is undecidable for transducer networks that are simple, except that send rules do not have to be static.

*Proof.* Inspired by the proof technique of Deutsch et al. [12], we reduce the the finite implication problem for functional and inclusion dependencies to the inconsistency decision problem [8]. We sketch the proof; the technical details are in Appendix A.1. An instance of the finite implication problem is a triple  $(\mathcal{D}, \Sigma, \sigma)$ , where  $\mathcal{D}$  is a database schema,  $\Sigma$  is a set of functional and inclusion dependencies over  $\mathcal{D}$ , and  $\sigma$  is a functional or inclusion dependency over  $\mathcal{D}$ . We call  $(\mathcal{D}, \Sigma, \sigma)$  *valid* if  $I \models \Sigma$  implies  $I \models \sigma$  for each instance  $I$  over  $\mathcal{D}$ .<sup>5</sup> We have to check validity of  $(\mathcal{D}, \Sigma, \sigma)$ .

For the instance  $(\mathcal{D}, \Sigma, \sigma)$ , we construct a single-node transducer network  $\mathcal{N}$  that is simple except that send rules are not static, and so that  $\mathcal{N}$  is inconsistent iff  $(\mathcal{D}, \Sigma, \sigma)$  is not valid. Let  $\Pi$  denote the single transducer of  $\mathcal{N}$ . We let the input schema of  $\Pi$  contain  $\mathcal{D}$ . Transducer  $\Pi$  sends a special marker message to itself, and when the marker is received,  $\Pi$  checks whether the input over  $\mathcal{D}$  satisfies  $\Sigma$  and  $\sigma$ . For each violated dependency  $\tau \in \Sigma \cup \{\sigma\}$ , transducer  $\Pi$

<sup>5</sup>We write  $I \models \sigma$  to denote that  $\sigma$  holds in  $I$ . We write  $I \models \Sigma$  to denote that  $I \models \sigma$  for each  $\sigma \in \Sigma$ .

sends a  $\text{viol}_\tau()$ -message to itself. Non-static send rules are needed for checking the inclusion dependencies.

Upon receiving  $\text{viol}_\sigma()$ , the transducer can do something inconsistent, by blocking a rule for output relation  $T$  as was done in Example 3.1, so that an incoming  $A_{\text{msg}}(a)$ -fact is ignored when memory fact  $B(a)$  was previously created. But when some  $\text{viol}_\tau()$  message with  $\tau \in \Sigma$  is received, we can repair the inconsistencies. Concretely, we fill a nullary memory relation **repair**, that is tested positively in another output rule for relation  $T$ . This second rule for  $T$  can henceforth output all received  $A_{\text{msg}}$ -facts.

Now, if  $(\mathcal{D}, \Sigma, \sigma)$  is not valid, there is an instance  $I$  over  $\mathcal{D}$  such that  $I \models \Sigma$  and  $I \not\models \sigma$ . Instance  $I$  can be extended to an input  $J$  for  $\mathcal{N}$ , and we make two runs as follows. In the first run, an output  $T(a)$  is produced by first delivering some fact  $A_{\text{msg}}(a)$  and by postponing the marker message (to postpone the dependency checking). In the second run, we do the converse, i.e., we deliver the marker first. Then, dependency  $\sigma$  turns out to be violated, and upon delivery of  $\text{viol}_\sigma()$ , we can block the output. No repairs are possible because only  $\sigma$  is violated.

Conversely, if  $\mathcal{N}$  is not consistent on some input  $J$ , this can only be explained by  $\sigma$  being violated and no dependency of  $\Sigma$ , so that the input of  $\mathcal{N}$  gives rise to an instance  $I$  over  $\mathcal{D}$  for which  $I \models \Sigma$  and  $I \not\models \sigma$ . Hence,  $(\mathcal{D}, \Sigma, \sigma)$  is not valid.  $\square$

**Proposition 3.3.** Inconsistency is undecidable for transducer networks that are simple, except that messages may participate in cycles in the local positive dependency graphs of individual transducers.

*Proof.* Inspired by the proof technique of Deutsch et al. [12], we reduce the Post correspondence problem to the inconsistency decision problem [19]. We sketch the proof; the technical details are in Appendix A.2. An instance of the Post correspondence problem is a pair  $(U, V)$  where  $U = u_1, \dots, u_n$  and  $V = v_1, \dots, v_n$  are two nonempty equal-length sequences of nonempty words over some alphabet with at least two symbols. A *match* for  $U$  and  $V$  is a sequence  $E = e_1, \dots, e_m$  of indices in  $\{1, \dots, n\}$  such that the words  $u_{e_1} \dots u_{e_m}$  and  $v_{e_1} \dots v_{e_m}$  are equal. Sequence  $E$  may contain the same index multiple times. The problem is to check whether a match exists.

For the instance  $(U, V)$ , we construct a single-node transducer network  $\mathcal{N}$  that is simple except that messages can have cyclic dependencies, and so that  $\mathcal{N}$  is inconsistent iff  $(U, V)$  has a match. Let  $\Pi$  denote the single transducer of  $\mathcal{N}$ . First, we provide  $\Pi$  with input relations to encode a word-structure: a binary relation  $R$  represents a chain, and a binary relation  $L$  assigns a label to each element of the chain.

The idea is to use messages to align the words of  $U$  and  $V$  to the input word-structure, to discover a match for  $(U, V)$ . Concretely, we use messages of the form  $\text{align}[i, k, l](a, b)$ , with  $i \in \{1, \dots, n\}$ ,  $k \in \{1, \dots, |u_i|\}$  and  $l \in \{1, \dots, |v_i|\}$ , expressing that we have already successfully aligned a sequence of  $(u_j, v_j)$ -pairs with  $j \in \{1, \dots, n\}$  to the word-structure, where  $(u_i, v_i)$  is the last pair tried, and the alignment of  $u_i$  and  $v_i$  has progressed partially up to respectively symbols  $k$  and  $l$ , arriving at respectively elements  $a$  and  $b$  of the word-structure. After a message  $\text{align}[i, |u_i|, |v_i|](a, b)$  is sent, indicating that  $(u_i, v_i)$  is fully aligned, we have sending rules to align a next pair  $(u_j, v_j)$ , by

sending message  $\text{align}[j, 1, 1](a', b')$ , where  $a'$  and  $b'$  are the successor-elements of respectively  $a$  and  $b$  on the word-structure. Adding unrestricted message recursion adds some notion of “iteration” to the transducer model: because message relations are allowed to participate in cycles, the alignment to the word-structure can repeatedly use the *same* pair  $(u_i, v_i)$ , allowing us to consider all candidate sequences  $E$  like above (but restricted to the input word structure).

If there is indeed a match for  $(U, V)$  then we can encode the resulting word as an input word-structure for  $\mathcal{N}$ . So, the above alignment process can eventually send a message of the form  $\text{align}[j, |u_j|, |v_j|](a, a)$ , i.e., we can align a sequence of  $(u_i, v_i)$ -pairs fully to the word-structure, where the implied concatenation of  $U$ -words ends at the same element of the word-structure as the implied concatenation of  $V$ -words. Then we do something inconsistent, like Example 3.1.

For the other direction, when  $\mathcal{N}$  is inconsistent on some input, we can attribute that to the sending of a message  $\text{align}[j, |u_j|, |v_j|](a, a)$ , whose derivation history reveals a match for  $(U, V)$  against a valid word-structure contained in the input of  $\mathcal{N}$ .  $\square$

By disallowing the syntactical liberties of the previous two propositions, we obtain decidability:

**Theorem 3.4.** Inconsistency for simple transducer networks is decidable in NEXPTIME; the problem is NEXPTIME-complete.

Theorem 3.4 is proven in Sections 4, 5, and 6.

## 4 Simulation on Single Node

Let  $\mathcal{N}$  be a simple transducer network. We construct a simple *single-node* transducer network  $\mathcal{M}$  that simulates  $\mathcal{N}$ , and so that  $\mathcal{M}$  is consistent iff  $\mathcal{N}$  is consistent. This will be made more precise below. The transformation can be done in PTIME for reasonable encodings of a transducer network, and so  $|\mathcal{M}|$  is polynomial in  $|\mathcal{N}|$  (cf. Section 2.8). The merit of this section lies in reducing the technical complexity for the decidability result (Sections 5 and 6) and the expressivity analysis (Section 7).

First, Section 4.1 gives syntactical simplifications for single-node networks. Next, Section 4.2 formalizes the notion of simulation and formulates the result. The sections thereafter show the result: Sections 4.3 and 4.4 respectively define the transducer schema and transducer of  $\mathcal{M}$ , and Section 4.5 shows that  $\mathcal{M}$  satisfies the desired properties.

### 4.1 Syntactical Simplifications

For a single-node transducer network  $\mathcal{M}$ , we use the following syntactical simplifications. It will be sufficient to view  $\mathcal{M}$  as consisting of only a transducer schema  $\Upsilon$  and a transducer  $\Pi$  over  $\Upsilon$ ; the actual node of  $\mathcal{M}$  is immaterial. The schemas  $\text{in}^{\mathcal{M}}$ ,  $\text{out}^{\mathcal{M}}$  and  $\text{mem}^{\mathcal{M}}$  (Section 2.7.1) are regarded as ordinary (non-distributed) database schemas. Accordingly, an input for  $\mathcal{M}$  is an ordinary database instance  $I$ . A configuration of  $\mathcal{M}$  on  $I$  is a pair  $(s, b)$  where  $s$  is a transducer state of  $\Pi$  and  $b$  is a multiset of facts over  $\Upsilon_{\text{msg}}$ . Because there is



only a single node, sending rules of  $\Pi$  have no explicit addressee variable in the head. Hence, schema  $\Upsilon_{\text{sys}}$  will not be used.

## 4.2 Simulation Concept and Result

To formalize the notion of “simulation”, we introduce some auxiliary notations. Let  $\mathcal{N}$  denote the network of  $\mathcal{N}$ . For a distributed database schema  $\mathcal{E}$  over  $\mathcal{N}$ , we view each node  $x \in \mathcal{N}$  as a namespace containing the relations  $\mathcal{E}(x)$ : we use symbol “ $x.R$ ” to denote relation  $R$  at  $x$ . Let  $\langle \mathcal{E} \rangle$  denote the (ordinary) database schema

$$\{x.R^{(k)} \mid x \in \mathcal{N}, R^{(k)} \in \mathcal{E}(x)\}.$$

For each distributed database instance  $H$  over  $\mathcal{E}$ , let  $\langle H \rangle$  be the following ordinary database instance over  $\langle \mathcal{E} \rangle$ :

$$\{x.R(\bar{a}) \mid x \in \mathcal{N}, R(\bar{a}) \in H(x)\}.$$

Let  $\text{sch}^{\mathcal{N}}$  denote the database schema  $\{x.\text{Id}^{(1)} \mid x \in \mathcal{N}\} \cup \{\text{Node}^{(1)}\}$ . Let  $\text{inst}^{\mathcal{N}}$  be the following instance over  $\text{sch}^{\mathcal{N}}$ :

$$\{x.\text{Id}(x), \text{Node}(x) \mid x \in \mathcal{N}\}.$$

We abbreviate  $\langle \mathcal{E} \rangle^{\mathcal{N}} = \langle \mathcal{E} \rangle \cup \text{sch}^{\mathcal{N}}$  and  $\langle H \rangle^{\mathcal{N}} = \langle H \rangle \cup \text{inst}^{\mathcal{N}}$ . We say that an instance  $I$  over  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  is *well-formed* if  $I$  is isomorphic to an instance  $J$  over  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  for which  $J|_{\text{sch}^{\mathcal{N}}} = \text{inst}^{\mathcal{N}}$ .<sup>6</sup> An instance that is not well-formed is called *ill-formed*.

For a configuration  $\rho = (s, b)$  of  $\mathcal{N}$ , we write  $\text{out}(\rho)$  to denote the following distributed instance  $H'$  over  $\text{out}^{\mathcal{N}}$ : for each  $x \in \mathcal{N}$ , instance  $H'(x)$  consists of all output facts in  $s(x)$ . If  $\mathcal{N}$  is a single-node network, we consider  $\text{out}(\rho)$  to be an ordinary database instance.

Now, we say that a single-node transducer network  $\mathcal{M}$  *simulates*  $\mathcal{N}$  if (i)  $\text{in}^{\mathcal{M}} = \langle \text{in}^{\mathcal{N}} \rangle^{\mathcal{N}}$ ; (ii)  $\text{out}^{\mathcal{M}} = \langle \text{out}^{\mathcal{N}} \rangle$ ; and, (iii) for each input  $H$  for  $\mathcal{N}$ , the following holds:

- for every run  $\mathcal{R}$  of  $\mathcal{N}$  on  $H$ , there is a run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  such that  $\langle \text{out}(\text{last}(\mathcal{R})) \rangle = \text{out}(\text{last}(\mathcal{S}))$ ,
- for every run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$ , there is a run  $\mathcal{R}$  of  $\mathcal{N}$  on  $H$  such that  $\langle \text{out}(\text{last}(\mathcal{R})) \rangle = \text{out}(\text{last}(\mathcal{S}))$ .

We use  $\text{in}^{\mathcal{M}} = \langle \text{in}^{\mathcal{N}} \rangle^{\mathcal{N}}$  instead of  $\text{in}^{\mathcal{M}} = \langle \text{in}^{\mathcal{N}} \rangle$  because  $\mathcal{M}$  needs the identifiers of the nodes to simulate message sending and the nodes' comparisons of their identifier to input values, and because we do not use values from **dom** directly in rules (cf. Section 2.3).

Now we are ready to present the result:

**Proposition 4.1.** For each simple transducer network  $\mathcal{N}$ , there exists a simple single-node transducer network  $\mathcal{M}$  such that (i)  $\mathcal{M}$  simulates  $\mathcal{N}$ , and (ii)  $\mathcal{M}$  is consistent iff  $\mathcal{N}$  is consistent.

Note, the simulation property says nothing about consistency and vice versa. The following subsections define  $\mathcal{M}$  so that the desired properties are satisfied.

<sup>6</sup> $I$  is isomorphic to  $J$  if there is an injective function  $f : \mathbf{dom} \rightarrow \mathbf{dom}$  such that  $f(I) = J$ .

### 4.3 Transducer Schema

We define the single transducer schema  $\Upsilon$  of  $\mathcal{M}$ . Denote  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$ . We write  $\mathcal{D}_{\text{msg}}^{\mathcal{N}}$  to denote the shared message schema of  $\mathcal{N}$ . We define  $\Upsilon$  as follows:

- $\Upsilon_{\text{in}} = \langle \text{in}^{\mathcal{N}} \rangle^{\mathcal{N}}$ ;  $\Upsilon_{\text{out}} = \langle \text{out}^{\mathcal{N}} \rangle$ ;  $\Upsilon_{\text{mem}} = \langle \text{mem}^{\mathcal{N}} \rangle$ ; and,
- $\Upsilon_{\text{msg}}$  consists of (i) the relations  $R_{\rightarrow x}^{(k+1)}$  for which  $x \in \mathcal{N}$  and  $R^{(k)} \in \mathcal{D}_{\text{msg}}^{\mathcal{N}}$ , (ii) a relation  $\text{do}_x^{(0)}$  for each  $x \in \mathcal{N}$ , (iii) relation **error**<sup>(0)</sup>, and (iv) relation **adom**<sup>(1)</sup>.

Relations of the form  $\text{do}_x$  allow us to explicitly simulate a transition of node  $x$ . Next, a relation  $R_{\rightarrow x}$  is used to send  $R$ -facts specifically to node  $x$ . The latter relations have an incremented arity when compared to  $\mathcal{D}_{\text{msg}}^{\mathcal{N}}$ , for the following reason. Each transition of the transducer  $\Pi$  in  $\mathcal{M}$  can simulate multiple nodes simultaneously, and these simulated nodes could send the same message to the same addressee. But the transition of  $\Pi$  can only send a *set* of messages. So, by letting  $\Pi$  additionally put the simulated sender node in each simulated message, we can avoid that these distinct simulated sending events would all be collapsed. Lastly, the relations **error** and **adom** allow  $\Pi$  to be consistent on ill-formed inputs; see below.

### 4.4 Transducer Rules

We now describe the single transducer  $\Pi$  of  $\mathcal{M}$ . Essentially, the  $\text{UCQ}^\top$  queries of  $\Pi$  are unions of modified  $\text{UCQ}^\top$  queries of the original transducers in  $\mathcal{N}$ . Some extra rules deal with ill-formed inputs.

#### 4.4.1 Output and Memory

We do the following for each node  $x \in \mathcal{N}$ . Let  $T^{(k)}$  be an output or memory relation in  $\Upsilon(x)$ . All rules for relation  $T$  in  $\Pi(x)$  are message-positive and message-bounded. An insertion rule  $\varphi$  for relation  $T$  in transducer  $\Pi(x)$  is modified to insertion rule  $\varphi'$  for relation  $x.T$  in  $\Pi$  as follows:

- input, output and memory atoms  $R(\bar{u})$  in  $\varphi$  become  $x.R(\bar{u})$  in  $\varphi'$ , including the head;
- atoms of the form **Id**( $u$ ) and **All**( $u$ ) in  $\varphi$  become respectively  $x.\text{Id}(u)$  and **Node**( $u$ ) in  $\varphi'$ ;
- (positive) message atoms  $R(\bar{u})$  in  $\varphi$  become  $R_{\rightarrow x}(z, \bar{u})$  in  $\varphi'$  where  $z$  is a new variable that is unique per message atom;
- the nonequalities in  $\varphi$  are the nonequalities in  $\varphi'$ ;
- $\varphi'$  additionally contains the positive body atom  $\text{do}_x()$ .

Intuitively, because relation **All** always contains  $\mathcal{N}$  on every node of  $\mathcal{N}$ , it is replaced by the shared relation **Node** in  $\mathcal{M}$ . For a message atom  $R_{\rightarrow x}(z, \bar{u})$ , the new variable  $z$  represents the extra sender-component (cf. Section 4.3). This component is not used elsewhere in the rule and is basically projected away.

The resulting output and memory insertion rules are message-positive and message-bounded. Because  $\Pi(x)$  is simple, there are no deletion rules for memory relations, so we don't have to translate these.

#### 4.4.2 Messages

We do the following for each node  $x \in \mathcal{N}$ . Let  $T^{(k)}$  be a shared message relation of  $\mathcal{N}$ . All rules for relation  $T$  in  $\Pi(x)$  are message-positive and static. To let simulated node  $x$  send messages in  $\mathcal{M}$ , we add to  $\Pi$  all rules  $\varphi'_y$  obtained by combining a sending rule  $\varphi$  for  $T$  in  $\Pi(x)$  and a node  $y \in \mathcal{N}$ . Intuitively, rule  $\varphi'_y$  models the sending of  $T$ -messages by  $x$  to the specific addressee  $y$ . Denote  $\text{head}^\varphi = T(\mathbf{n}_0, \bar{\mathbf{u}})$ , where  $\mathbf{n}_0$  is the addressee variable. Let  $\mathbf{n}_1$  be a new variable. Rule  $\varphi'_y$  is obtained as follows:

- the head  $T(\mathbf{n}_0, \bar{\mathbf{u}})$  of  $\varphi$  becomes the head  $T_{\rightarrow y}(\mathbf{n}_1, \bar{\mathbf{u}})$  in  $\varphi'_y$ ;
- $\varphi'_y$  contains positive body atoms  $y.\text{Id}(\mathbf{n}_0)$  and  $x.\text{Id}(\mathbf{n}_1)$ ;
- input atoms  $R(\bar{\mathbf{u}})$  in  $\varphi$  become  $x.R(\bar{\mathbf{u}})$  in  $\varphi'_y$ ;
- atoms of the form  $\text{Id}(\mathbf{u})$  and  $\text{All}(\mathbf{u})$ , and message atoms, are transformed as in the output and memory rules above;
- the nonequalities of  $\varphi$  are the nonequalities of  $\varphi'_y$ ;
- $\varphi'_y$  additionally contains the positive body atom  $\text{do}_x()$ .

Variable  $\mathbf{n}_0$  is not removed because it might occur on several places in  $\varphi$ , and by adding the atom  $y.\text{Id}(\mathbf{n}_0)$ , we fix the addressee  $y$ . Variable  $\mathbf{n}_1$  represents the sender  $x$  by addition of the body atom  $x.\text{Id}(\mathbf{n}_1)$ , and  $\mathbf{n}_1$  replaces  $\mathbf{n}_0$  in the head.

Denote  $\mathcal{N} = \{x_1, \dots, x_n\}$ . For each  $x \in \mathcal{N}$ , we also add the following rule to  $\Pi$ , to send simulation messages for  $x$ :

$$\text{do}_x() \leftarrow x_1.\text{Id}(\mathbf{u}_1), \dots, x_n.\text{Id}(\mathbf{u}_n).$$

The above rule has the effect that a message  $\text{do}_y()$  for any  $y \in \mathcal{N}$  can only be sent if *all* relations  $z.\text{Id}$  with  $z \in \mathcal{N}$  are nonempty. And because the simulated output, memory, and sending rules are guarded by message atoms of the form  $\text{do}_y()$ , the *entire* simulation requires that these relations  $z.\text{Id}$  are nonempty.

The above message rules of  $\Pi$  are all message-positive and static.

#### 4.4.3 Ill-formed Inputs

We indicate how  $\mathcal{M}$  can be made consistent on ill-formed instances. First, using message-positive and static send rules, it is possible to send a message **error**() if the following constraints are violated: some relation  $x.\text{Id}$  contains two different values; two relations  $x.\text{Id}$  and  $y.\text{Id}$  with  $x \neq y$  share a value; relation **Node** is not the union of all  $x.\text{Id}$  relations.

We also add new output rules that on receipt of **error**() can produce all possible output facts in  $\Upsilon_{\text{out}}$ . Technically, this is done by adding rules to send all values  $a$  from the input active domain as an **adom**( $a$ )-message, and the additional output rules combine these values upon delivery when **error**() is also jointly delivered.

#### 4.4.4 Check Simple

We verify that  $\Pi$  is simple: (i)  $\Pi$  is inflationary by construction; (ii)  $\Pi$  is recursion-free because the transducers of  $\mathcal{N}$  are recursion-free and because there are no cycles in the positive message dependency graph of  $\mathcal{N}$ ; and, (iii) the desired constraints on output, memory and sending rules hold, as remarked above. Moreover, because  $\Pi$  is the only transducer of  $\mathcal{M}$  and  $\Pi$  is recursion-free, there are no cycles in the positive message dependency graph of  $\mathcal{M}$ , and thus  $\mathcal{M}$  is simple.

### 4.5 Simulation and Consistency Equivalence

We now show that (i)  $\mathcal{M}$  simulates  $\mathcal{N}$  and (ii)  $\mathcal{M}$  is consistent iff  $\mathcal{N}$  is consistent. First we need some additional concepts and notations. Let  $\rho = (s, b)$  be a configuration of  $\mathcal{N}$  on input  $H$  and let  $\sigma = (s', b')$  be a configuration of  $\mathcal{M}$  on input  $\langle H \rangle^{\mathcal{N}}$ . We say that  $\sigma$  and  $\rho$  are *output-equivalent* if for each  $x \in \mathcal{N}$  and each output relation  $R$  at  $x$ , we have  $R(\bar{a}) \in s(x)$  iff  $x.R(\bar{a}) \in s'$ . The notions of *input-*, *memory-*, and *system-equivalence* can be similarly defined, where the latter is about relations  $\text{Id}$  and  $\text{All}$ . By definition of  $\langle H \rangle^{\mathcal{N}}$ , configuration  $\sigma$  is always input- and system-equivalent to  $\rho$ .

We say that  $\sigma$  is *message-equivalent* to  $\rho$  if for each  $x \in \mathcal{N}$ , for each fact  $R(\bar{a})$ , the cardinality of  $R(\bar{a})$  in  $b(x)$  equals the number of messages of the form  $R_{\rightarrow x}(z, \bar{a})$  in  $b'$  (each may have a different sender component). Similarly, we say that  $\sigma$  has its *messages included* in  $\rho$  when for each  $x \in \mathcal{N}$  the number of messages of the form  $R_{\rightarrow x}(z, \bar{a})$  in  $b'$  is less than or equal to the cardinality of  $R(\bar{a})$  in  $b(x)$ .

Claims 4.2 and 4.3 show that  $\mathcal{M}$  simulates  $\mathcal{N}$ , but they are phrased slightly more general for later use in the consistency equivalence:

**Claim 4.2.** Every run  $\mathcal{R}$  of  $\mathcal{N}$  on an input  $H$  can be converted to a run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  such that  $\text{last}(\mathcal{S})$  and  $\text{last}(\mathcal{R})$  are output-, memory-, and message-equivalent.

*Proof.* Let  $n$  be the number of transitions in  $\mathcal{R}$ , and let  $x_1, \dots, x_n$  be the active nodes in order. Run  $\mathcal{S}$  will consist of  $n+1$  transitions: for each  $i = 1, \dots, n$ , we deliver  $\text{do}_{x_i}()$  in transition  $i+1$  of  $\mathcal{S}$  (and no other  $\text{do}_y$ -messages). We start  $\mathcal{S}$  by doing one heartbeat transition, so that at least  $\text{do}_{x_1}()$  is sent. This message is delivered in the second transition of  $\mathcal{S}$ , to simulate the behaviour of node  $x_1$ . By input- and system-equivalence of the second configuration of  $\mathcal{S}$  and the first configuration of  $\mathcal{R}$ , the third configuration of  $\mathcal{S}$  and the second configuration of  $\mathcal{R}$  are output-, memory-, and message-equivalent. We can now repeat the same for nodes  $x_2, x_3$ , etc. Moreover, the message-equivalence allows us to deliver  $k$  messages of the form  $R_{\rightarrow x}(z, \bar{a})$  in a transition of  $\mathcal{S}$  when the corresponding transition in  $\mathcal{R}$  would deliver  $k$  instances of (the same) message  $R(\bar{a})$  to an active node  $x$ .  $\square$

**Claim 4.3.** Let  $H$  be an input for  $\mathcal{N}$ . Every run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  can be converted to a run  $\mathcal{R}$  of  $\mathcal{N}$  on  $H$  such that  $\text{last}(\mathcal{R})$  and  $\text{last}(\mathcal{S})$  are output- and memory-equivalent, and  $\text{last}(\mathcal{S})$  has its messages included in  $\text{last}(\mathcal{R})$ .

*Proof.* First, some transitions of  $\mathcal{S}$  might deliver a message of the form  $R_{\rightarrow x}(z, \bar{a})$  without jointly delivering  $\text{do}_x()$ . Because node  $x$  is only simulated when  $\text{do}_x()$

is delivered, message  $R_{\rightarrow x}(z, \bar{a})$  is effectively lost. So, we can refrain from delivering  $R_{\rightarrow x}(z, \bar{a})$  in this case, without compromising future message deliveries. After doing this modification for all deliveries of  $\mathcal{S}$ , we also drop any resulting (or preexisting) heartbeat transitions except the first transition, because they do not simulate nodes.<sup>7</sup> This results in a new run  $\mathcal{S}'$  such that  $\text{last}(\mathcal{S})$  and  $\text{last}(\mathcal{S}')$  have the same output and memory facts, and such that the buffer of  $\text{last}(\mathcal{S})$  is included in the buffer of  $\text{last}(\mathcal{S}')$  when ignoring the  $\text{do}_x$ -messages.

Next, some transitions  $i$  of  $\mathcal{S}'$  might deliver two messages  $\text{do}_x()$  and  $\text{do}_y()$  with  $x \neq y$ . Such a transition  $i$  simulates multiple nodes in parallel. But in  $\mathcal{M}$ , the simulated rules of each node  $x$  are guarded by  $\text{do}_x()$ , and these rules can only access relations of  $x$  itself. Hence, transition  $i$  can be converted to a sequence of transitions in which only one node is simulated at a time (in some arbitrary order), and in which each node receives the same messages that it received in  $i$ . This results in a new run  $\mathcal{S}''$ , where  $\text{last}(\mathcal{S}')$  and  $\text{last}(\mathcal{S}'')$  are exactly the same when ignoring the  $\text{do}_x$ -messages.

Starting from the second transition, run  $\mathcal{S}''$  simulates precisely one node in each transition. In the opposite fashion as in Claim 4.2, we can now convert  $\mathcal{S}''$  to a run  $\mathcal{R}$  of  $\mathcal{N}$  on input  $H$  so that  $\text{last}(\mathcal{S}'')$  and  $\text{last}(\mathcal{R})$  are output-, memory-, and message-equivalent. Note,  $\text{last}(\mathcal{S})$  and  $\text{last}(\mathcal{R})$  are output- and memory-equivalent, and  $\text{last}(\mathcal{S})$  has its messages included in  $\text{last}(\mathcal{R})$ .  $\square$

Now we are ready for the actual consistency equivalence between  $\mathcal{N}$  and  $\mathcal{M}$ , where each direction is shown in a separate claim:

**Claim 4.4.** If  $\mathcal{M}$  is consistent then  $\mathcal{N}$  is consistent.

*Proof.* Let  $H$  be an input for  $\mathcal{N}$ . Let  $\mathcal{R}_1$  and  $\mathcal{R}_2$  be two runs of  $\mathcal{N}$  on  $H$ , where  $\text{last}(\mathcal{R}_1)$  contains an output fact  $R(\bar{a})$  at some node  $x \in \mathcal{N}$ . We have to show that  $\mathcal{R}_2$  can be extended to a run  $\mathcal{R}'_2$  such that  $\text{last}(\mathcal{R}'_2)$  also contains fact  $R(\bar{a})$  at  $x$ . Using Claim 4.2, we can make two runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  such that for each  $i \in \{1, 2\}$ , configurations  $\text{last}(\mathcal{S}_i)$  and  $\text{last}(\mathcal{R}_i)$  are output-, memory-, and message-equivalent. In particular,  $\text{last}(\mathcal{S}_1)$  contains output fact  $x.R(\bar{a})$ . By consistency of  $\mathcal{M}$ , run  $\mathcal{S}_2$  can be extended to a run  $\mathcal{S}'_2$  such that  $\text{last}(\mathcal{S}_2)$  also contains  $x.R(\bar{a})$ . Lastly, extension  $\mathcal{S}'_2$  gives rise to an extension  $\mathcal{R}'_2$  such that  $\text{last}(\mathcal{R}'_2)$  is output- and memory-equivalent to  $\text{last}(\mathcal{S}'_2)$ , and so  $\text{last}(\mathcal{R}'_2)$  contains  $R(\bar{a})$  at  $x$ : the proof is similar to that of Claim 4.3, with the exception that the configurations in  $\mathcal{S}'_2$  have their messages included in the corresponding configurations of  $\mathcal{R}'_2$ . This is sufficient to guarantee that  $\mathcal{R}'_2$  can mimick the behaviour of  $\mathcal{S}'_2$ .  $\square$

**Claim 4.5.** If  $\mathcal{N}$  is consistent then  $\mathcal{M}$  is consistent.

*Proof.* Let  $I$  be an input for  $\mathcal{M}$ . We have to show that  $\mathcal{M}$  is consistent on  $I$ .

First, suppose that  $I$  is ill-formed. If  $I$  does not contain a value for each relation  $x.\text{Id}$  with  $x \in \mathcal{N}$  then no output can ever be produced. Indeed, no message  $\text{do}_x()$  for any  $x \in \mathcal{N}$  can be sent (and delivered), so no inconsistency could arise because the nodes are not simulated. Otherwise, if  $I$  contains a value for each relation  $x.\text{Id}$ , because  $I$  is still ill-formed, it will be possible to send  $\text{error}()$ . Then any run can be extended to produce all possible output facts, so potential inconsistencies can always be corrected.

<sup>7</sup>This does not compromise the supply of  $\text{do}_x$ -messages because they are sent in each transition.

Now suppose that  $I$  is well-formed, which means there is an instance  $J$  isomorphic to  $I$  with  $J|_{sch\mathcal{N}} = inst^{\mathcal{N}}$  (cf. Section 4.3). Because transducer rules of  $\mathcal{M}$  only express generic queries, it is sufficient to show that  $\mathcal{M}$  is consistent on  $J$ . Let  $H$  be the (unique) input for  $\mathcal{N}$  for which  $\langle H \rangle^{\mathcal{N}} = J$ . Let  $\mathcal{S}_1$  and  $\mathcal{S}_2$  be two runs of  $\mathcal{M}$  on  $J$ , where  $last(\mathcal{S}_1)$  contains an output fact  $x.R(\bar{a})$ . We have to show that there is an extension of  $\mathcal{S}_2$  for which the last configuration also contains  $x.R(\bar{a})$ .

First, applying Claim 4.3 to run  $\mathcal{S}_1$ , we can construct a run  $\mathcal{R}_1$  of  $\mathcal{N}$  on input  $H$  such that  $last(\mathcal{S}_1)$  and  $last(\mathcal{R}_1)$  are output- and memory-equivalent. In particular, output fact  $R(\bar{a})$  is at node  $x$  in  $last(\mathcal{R}_1)$ .

Next, suppose we can construct an extension  $\mathcal{S}_2''$  of  $\mathcal{S}_2$  and a run  $\mathcal{R}_2''$  of  $\mathcal{N}$  on input  $H$  such that  $last(\mathcal{S}_2'')$  and  $last(\mathcal{R}_2'')$  are output-, memory-, and message-equivalent. If by chance  $last(\mathcal{S}_2'')$  already contains  $x.R(\bar{a})$  then we are ready. Otherwise, by output-equivalence of  $last(\mathcal{S}_2'')$  and  $last(\mathcal{R}_2'')$ , fact  $R(\bar{a})$  will not be at  $x$  in  $last(\mathcal{R}_2'')$ . But, by consistency of  $\mathcal{N}$ , because  $R(\bar{a})$  can be derived at  $x$  in  $\mathcal{R}_1$  (see above), there is an extension of  $\mathcal{R}_2''$  to derive  $R(\bar{a})$  at  $x$ . By message-equivalence of  $last(\mathcal{S}_2'')$  and  $last(\mathcal{R}_2'')$ , this extension can be simulated at the end of  $\mathcal{S}_2''$  to derive  $x.R(\bar{a})$ , in a similar vein as in the proof of Claim 4.2.

We are left to construct the runs  $\mathcal{S}_2''$  and  $\mathcal{R}_2''$ .

**Message saturation** Because transducer  $\Pi$  of  $\mathcal{M}$  is recursion-free, we can consider the maximum height  $n$  amongst derivation trees of  $\Pi$ , where the height is the largest number of edges on any path from a leaf to the root. Now, we extend  $\mathcal{S}_2$  to a run  $\mathcal{S}_2'$  by doing  $n$  additional transitions: each transition delivers the *entire* message buffer, and thus simulates all nodes in parallel where each node receives its entire (simulated) message buffer.<sup>8</sup> Because the sending rules of  $\Pi$  are message-positive and static, the message buffer of  $\mathcal{M}$  — degenerated to a set — will monotonously grow. Because  $n$  is the maximum height of a derivation tree,  $last(\mathcal{S}_2')$  contains all messages that could possibly be sent on input  $J$ .

**Run of  $\mathcal{N}$**  Applying Claim 4.3 to  $\mathcal{S}_2'$  (not to  $\mathcal{S}_2$ ), we can construct a run  $\mathcal{R}_2'$  of  $\mathcal{N}$  on input  $H$  such that  $last(\mathcal{S}_2')$  and  $last(\mathcal{R}_2')$  are output- and memory-equivalent, and such that the messages of  $last(\mathcal{S}_2')$  are included in  $last(\mathcal{R}_2')$ . We now show that actually all messages in the buffers of  $last(\mathcal{R}_2')$  are simulated in the (single) buffer of  $last(\mathcal{S}_2')$ , except for maybe their precise cardinalities.

Let  $S(\bar{b})$  be a message in the buffer of some node  $y$  in  $last(\mathcal{R}_2')$ . We can extract from  $\mathcal{R}_2'$  a “global” derivation tree  $\mathcal{T}$  to explain how  $S(\bar{b})$  was sent to  $y$ : this is like a normal derivation tree, except that we also say at which node a message was derived. Letting  $\Pi$  be the single transducer of  $\mathcal{M}$ , and letting  $x$  be the node in the root of  $\mathcal{T}$  (i.e.,  $x$  sends  $S(\bar{b})$  to  $y$ ), the natural correspondence between  $\Pi$  and  $\mathcal{N}$  allows us to convert  $\mathcal{T}$  into a derivation tree  $\mathcal{T}'$  of  $\Pi$ , to explain how to send the message  $S_{\rightarrow y}(x, \bar{b})$ . Because sending rules are message-positive and static, this tree  $\mathcal{T}'$  is successfully executed in the last  $n$  transitions of  $\mathcal{S}_2'$ , so that  $S_{\rightarrow y}(x, \bar{b})$  is in the message buffer of  $last(\mathcal{S}_2')$ , as desired.

<sup>8</sup>We assume run  $\mathcal{S}_2$  contains at least one transition, so that all  $do_x$ -messages are available in the buffer of  $last(\mathcal{S}_2)$ .

**Obtain message-equivalence** Consider the extension  $\mathcal{R}_2''$  of  $\mathcal{R}_2'$  that is obtained by letting each node, in some arbitrary order, receive its entire message buffer from configuration  $last(\mathcal{R}_2')$ . Similarly, consider the extension  $\mathcal{S}_2''$  of  $\mathcal{S}_2'$  obtained by letting each simulated node, in the same order as in  $\mathcal{R}_2''$ , receive its entire message buffer as it is simulated by configuration  $last(\mathcal{S}_2')$ .

As we have seen above,  $last(\mathcal{R}_2')$  and  $last(\mathcal{S}_2')$  essentially represent the same messages in the buffer of each node, except that the cardinalities might be different. But since duplicate messages are collapsed upon delivery, the nodes do not observe the difference in cardinalities when the above two extensions are performed. Hence, configurations  $last(\mathcal{R}_2'')$  and  $last(\mathcal{S}_2'')$  are output- and memory-equivalent. But they are also message-equivalent as we now explain. First, for a node  $y \in \mathcal{N}$ , the extensions deliver equivalent message sets to  $y$ . Hence, in both extensions, node  $y$  in turn sends equivalent message sets. And because node  $y$  has its entire message buffer (of configurations  $last(\mathcal{R}_2')$  and  $last(\mathcal{S}_2')$ ) emptied during the delivery, the cardinalities of messages in  $last(\mathcal{R}_2'')$  and  $last(\mathcal{S}_2'')$  are the same.  $\square$

## 5 Small Model Property

Let  $\mathcal{N}$  be a simple single-node transducer network. We establish a small model property: if  $\mathcal{N}$  is inconsistent, then  $\mathcal{N}$  is inconsistent on an input whose active domain size is upper bounded by an expression purely over syntactical properties of  $\mathcal{N}$ . For this result, we use all syntactical restrictions of simple transducer networks.

Let  $\Pi$  and  $\Upsilon$  denote respectively the transducer and its schema in  $\mathcal{N}$ . Like in Section 4, an input for  $\mathcal{N}$  is an instance  $I$  over  $\Upsilon_{\text{in}}$ , and a configuration of  $\mathcal{N}$  is a pair  $(s, b)$  where  $s$  is a transducer state and  $b$  is a multiset of facts over  $\Upsilon_{\text{msg}}$ . Moreover, the sending rules have no explicit addressee variable in their head, and  $\Upsilon_{\text{sys}}$  will not be used in any rule. Such a network can always be obtained by applying the simulation in Section 4.

### 5.1 Syntactical Quantities

Consider the following syntactically defined quantities about  $\mathcal{N}$ :

- the length  $\mathbf{P}$  of the longest path in the positive dependency graph of  $\Pi$  (defined in Section 3.2), where the length of a path is measured as the number of edges on this path;
- the largest number  $\mathbf{B}$  of positive body atoms in any rule of  $\Pi$ ;
- the largest arity  $\mathbf{I}$  among input relations;
- the largest arity  $\mathbf{O}$  among output relations;
- the number  $\mathbf{C}$  of different output and memory facts that can be made with values in  $A$ , where  $A \subseteq \mathbf{dom}$  is an arbitrary set with  $|A| = \mathbf{O}$ .

Now, let  $sizeDom(\mathcal{N})$  abbreviate the expression  $2\mathbf{ICB}^{\mathbf{P}}$ . We have the following small model property:

**Proposition 5.1.** If  $\mathcal{N}$  is inconsistent, then  $\mathcal{N}$  is inconsistent on an instance  $J$  over  $\Upsilon_{\text{in}}$  for which  $|\text{adom}(J)| \leq \text{sizeDom}(\mathcal{N})$ .

The rest of this section is devoted to showing this result.

## 5.2 Proof Outline

Here we sketch the proof of Proposition 5.1. The details are provided by the following subsections. The proof technique is inspired by *pseudoruns* from Deutsch et al. [11], although it was adapted to deal with the inconsistency problem and to deal with message buffers (multisets). Let  $\mathcal{N}$ ,  $\Pi$  and  $\Upsilon$  be like above, and recall the syntactical quantities of  $\mathcal{N}$  from Section 5.1.

First we give some additional terminology and notations. Let  $A \subseteq \mathbf{dom}$ . We call a fact  $\mathbf{g}$  an *A-fact* if the values in  $\mathbf{g}$  are a subset of  $A$ . For a set of facts  $H$ , we write  $H^{[A]}$  to denote the subset of all *A*-facts in  $H$ . Note, nullary facts of  $H$  are always in  $H^{[A]}$ .

Let  $I$  be an input for  $\mathcal{N}$ . Suppose  $\mathcal{N}$  is inconsistent on  $I$ , i.e., there are two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  of  $\mathcal{N}$  on  $I$  such that  $\text{last}(\mathcal{R}_1)$  contains an output fact  $\mathbf{f}$  that is not in  $\text{last}(\mathcal{R}_2)$ , and there is no extension  $\mathcal{R}'_2$  of  $\mathcal{R}_2$  such that  $\text{last}(\mathcal{R}'_2)$  contains  $\mathbf{f}$ . Let  $C \subseteq \mathbf{dom}$  be the set of values in  $\mathbf{f}$ . Note,  $|C| \leq \mathbf{o}$ .

In Section 5.3, for  $i = 1, 2$ , we will select a subset of input facts  $K_i \subseteq I$  that are needed to make all output and memory  $C$ -facts of run  $\mathcal{R}_i$ , with the property  $|K_i| \leq \mathbf{CB}^{\mathbf{P}}$ . This gives the instances  $K_1$  and  $K_2$ . Note,  $C \subseteq \text{adom}(K_1)$  because  $\mathbf{f}$  is created in  $\mathcal{R}_1$ . Define

$$J = I^{[\text{adom}(K_1) \cup \text{adom}(K_2)]}.$$

Note,  $|\text{adom}(J)| \leq 2\mathbf{ICB}^{\mathbf{P}} = \text{sizeDom}(\mathcal{N})$ .

Next, in Section 5.4, for  $i = 1, 2$ , we will construct a run  $\mathcal{S}_i$  on input  $J$  with the following properties:

- $\text{last}(\mathcal{S}_i)$  and  $\text{last}(\mathcal{R}_i)$  contain precisely the same output and memory  $C$ -facts;
- every extension  $\mathcal{S}'_i$  of  $\mathcal{S}_i$  gives rise to an extension  $\mathcal{R}'_i$  of  $\mathcal{R}_i$  such that  $\text{last}(\mathcal{S}'_i)$  and  $\text{last}(\mathcal{R}'_i)$  again contain precisely the same output and memory  $C$ -facts.

This gives the runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  on  $J$ . The focus on output and memory  $C$ -facts is mainly the result of the message-boundedness constraint. Since  $\mathbf{f}$  is an output  $C$ -fact, the first property above tells us that  $\text{last}(\mathcal{S}_1)$  contains  $\mathbf{f}$  and  $\text{last}(\mathcal{S}_2)$  does not. Moreover, if  $\mathcal{S}_2$  can be extended to a run  $\mathcal{S}'_2$  such that  $\text{last}(\mathcal{S}'_2)$  contains  $\mathbf{f}$ , then the second property above would tell us that  $\mathcal{R}_2$  can be extended to a run  $\mathcal{R}'_2$  such that  $\text{last}(\mathcal{R}'_2)$  also contains  $\mathbf{f}$ . But the latter is not possible by assumption on  $\mathcal{R}_2$ . Hence,  $\mathcal{S}'_2$  does not exist, and  $\mathcal{N}$  is inconsistent on the instance  $J$ , whose active domain size is upper bounded by  $\text{sizeDom}(\mathcal{N})$ , as desired.

## 5.3 Input Selection

Consider the symbols defined in Sections 5.1 and 5.2. Let  $\mathcal{R}$  be either  $\mathcal{R}_1$  or  $\mathcal{R}_2$ . In this section, we select an instance  $K \subseteq I$  that is needed to make all output and memory  $C$ -facts of  $\mathcal{R}$ , and such that  $|K| \leq \mathbf{CB}^{\mathbf{P}}$ .



We construct a derivation history of each output and memory  $C$ -fact in  $\mathcal{R}$ : this includes the rules and valuations that derive the  $C$ -facts, and it also includes the derivation histories of messages recursively needed to make those  $C$ -facts.

### 5.3.1 Derivation History

Let  $\mathbf{g}$  be an output or memory  $C$ -fact derived during  $\mathcal{R}$ . By inflationarity of  $\Pi$ , the derivation of  $\mathbf{g}$  happens in some unique transition  $i$ . We choose *one* pair  $(\varphi, V)$  of a rule  $\varphi$  and satisfying valuation  $V$  such that  $\mathbf{g}$  is derived during transition  $i$  by applying  $V$  to  $\varphi$ . Let us call  $(\varphi, V)$  a *derivation pair*. If  $\varphi$  contains a (positive) body message atom  $\mathbf{a}$ , the message  $\mathbf{h} = V(\mathbf{a})$  is required by  $(\varphi, V)$  to derive  $\mathbf{g}$ . Similarly as we did for  $\mathbf{g}$ , we can go to a transition in which  $\mathbf{h}$  was derived and select there also *one* pair  $(\varphi', V')$  to derive  $\mathbf{h}$ . We can again recursively repeat the selection of derivation pairs for any message facts needed by  $(\varphi', V')$ .

Formally, after the selection of derivation pairs, we obtain a function  $hist_{\mathcal{R}}$  that maps each pair  $(i, \mathbf{g})$  to a derivation pair for  $\mathbf{g}$ , where  $\mathbf{g}$  is an output or memory  $C$ -fact or a recursively needed message derived in transition  $i$ . We also have a set  $msg_{\mathcal{R}}$  containing triples  $(k, \mathbf{h}, l)$  to indicate that a valuation in transition  $l$  needs the message  $\mathbf{h}$  to arrive, and that  $\mathbf{h}$  itself is sent in (an earlier) transition  $k$ . These triples indicate the timing of the required messages.

Now, let  $K$  denote the subset of all input facts  $\mathbf{h} \in I$  for which there exists a pair  $(i, \mathbf{g})$  in the domain of  $hist_{\mathcal{R}}$ , denoting  $hist_{\mathcal{R}}(i, \mathbf{g}) = (\varphi, V)$ , such that  $\mathbf{h} \in V(pos^{\varphi})$ . In words:  $K$  contains the (positive) input facts needed by the derivation history of all output and memory  $C$ -facts in  $\mathcal{R}$  (and any needed messages). We now show  $|K| \leq \mathbf{CB}^{\mathbf{P}}$ . First, let us fix one output or memory  $C$ -fact  $\mathbf{g}$ . Any chain of messages recursively needed by  $\mathbf{g}$  has length at most  $\mathbf{P}$  by recursion-freeness of  $\Pi$ . Moreover, in the worst case, each message recursively requires  $\mathbf{B}$  other messages. Therefore, the number of input facts needed by  $\mathbf{g}$  alone is bounded by  $\mathbf{B}^{\mathbf{P}}$ . And since at most  $\mathbf{C}$  different output and memory  $C$ -facts are created in  $\mathcal{R}$ , we overall have that  $|K| \leq \mathbf{CB}^{\mathbf{P}}$ , as desired.

### 5.3.2 Natural Properties

Section 5.3.1 allows much liberty in which  $hist_{\mathcal{R}}$  and  $msg_{\mathcal{R}}$  may be chosen. We now demand that some natural properties hold on  $msg_{\mathcal{R}}$ , upon which the construction in Section 5.4 crucially depends.

First, based on  $msg_{\mathcal{R}}$ , for each transition  $i$  of  $\mathcal{R}$ , we define the message multisets  $\beta_i$ ,  $\gamma_i$ , and  $\mathcal{E}_i$  as follows, with the intuition provided below:

- the multiplicity of a message  $\mathbf{h}$  in  $\beta_i$  is the number of triples  $(k, \mathbf{h}, l) \in msg_{\mathcal{R}}$  for which  $l = i$ ;
- the multiplicity of a message  $\mathbf{h}$  in  $\gamma_i$  is the number of triples  $(k, \mathbf{h}, l) \in msg_{\mathcal{R}}$  for which  $k < i$  and  $i \leq l$ ;
- the multiplicity of a message  $\mathbf{h}$  in  $\mathcal{E}_i$  is the number of triples  $(k, \mathbf{h}, l) \in msg_{\mathcal{R}}$  for which  $k = i$ .

Let  $\rho_1, \dots, \rho_n, \rho_{n+1}$  denote the sequence of configurations of  $\mathcal{R}$ , where  $n$  is the number of transitions. Intuitively,  $\beta_i$  contains the messages needed in transition  $i$ ;  $\gamma_i$  contains the needed messages that are sent before configuration  $\rho_i$  and that

travel through configuration  $\rho_i$  to be delivered in transition  $i$  (when  $l = i$ ) or later (when  $i < l$ ); and,  $\mathcal{E}_i$  contains the needed messages that should be sent in transition  $i$ .

In Appendix B.1, we show that  $hist_{\mathcal{R}}$  and  $msg_{\mathcal{R}}$  can be chosen so that the following properties are satisfied, with the intuition provided below:

1.  $\gamma_i \subseteq b_i^{\mathcal{R}}$  for each transition  $i$  of  $\mathcal{R}$ , where  $\rho_i = (s_i^{\mathcal{R}}, b_i^{\mathcal{R}})$ ;
2.  $\beta_i$  is a set for each transition  $i$  of  $\mathcal{R}$ , i.e., for each  $(k, \mathbf{h}, i)$  and  $(k', \mathbf{h}, i)$  in  $msg_{\mathcal{R}}$ , we have  $k = k'$ ;
3.  $\mathcal{E}_i = \gamma_{i+1} \cap \delta_i^{\mathcal{R}}$ , where  $\delta_i^{\mathcal{R}}$  is the set of messages sent in transition  $i$  of  $\mathcal{R}$ .

Intuitively, property 1 means that all needed messages whose transmission overlaps in time, also jointly occur in the message buffer, with the correct cardinalities. Property 2 means that if multiple derivation pairs in the same transition need the same message, the same origin of this message is used. Lastly, property 3 implies that for each needed message, its origin transition is chosen as late as possible: whenever for some needed message  $\mathbf{h} \in \gamma_{i+1}$  we have the opportunity to explain its origin in transition  $i$  (i.e.,  $\mathbf{h} \in \delta_i^{\mathcal{R}}$ ), we take this opportunity (i.e.,  $\mathbf{h} \in \mathcal{E}_i$ ).

## 5.4 Run Projection

Consider the symbols defined in Section 5.2. Let  $\mathcal{R}$  be either  $\mathcal{R}_1$  or  $\mathcal{R}_2$ . We construct a run  $\mathcal{S}$  on input  $J$  with the following properties:

- $last(\mathcal{S})$  and  $last(\mathcal{R})$  contain the same output and memory  $C$ -facts;
- every extension  $\mathcal{S}'$  of  $\mathcal{S}$  gives rise to an extension  $\mathcal{R}'$  of  $\mathcal{R}$  such that  $last(\mathcal{S}')$  and  $last(\mathcal{R}')$  again contain precisely the same output and memory  $C$ -facts.

To improve the readability of this section, helper claims are placed in Appendix B.2. First, Claim B.4 tells us that the second property above holds when the first property holds *and* when the message buffer of  $last(\mathcal{S})$  is included in the message buffer of  $last(\mathcal{R})$ . Intuitively, this inclusion allows every extension  $\mathcal{S}'$  of  $\mathcal{S}$  to be converted to an extension  $\mathcal{R}'$  of  $\mathcal{R}$  so that the buffer of  $\mathcal{S}'$  *remains* included in the buffer of  $\mathcal{R}'$ , allowing  $\mathcal{R}'$  to make precisely the same message deliveries as  $\mathcal{S}'$ .

We first sketch the main idea in the construction of  $\mathcal{S}$ . For run  $\mathcal{R}$ , let  $hist_{\mathcal{R}}$ ,  $msg_{\mathcal{R}}$ ,  $\beta_i$ ,  $\gamma_i$ , and  $\mathcal{E}_i$  be as defined in Section 5.3. We assume that  $msg_{\mathcal{R}}$  satisfies the properties given in Section 5.3.2. Run  $\mathcal{S}$  will be a projected version of  $\mathcal{R}$ : we do the same number of transitions as  $\mathcal{R}$ , and perform the message deliveries selected by  $msg_{\mathcal{R}}$ , so that the output and memory  $C$ -facts of  $\mathcal{R}$  are faithfully created. One caveat, however, is that some transitions of  $\mathcal{S}$  should sometimes deliver more messages than just those of  $msg_{\mathcal{R}}$  because we want the message buffer of  $\mathcal{S}$  to be included in the corresponding message buffer of  $\mathcal{R}$  (see above).

Let  $n$  be the number of transitions in  $\mathcal{R}$ . For each  $i \in \{1, \dots, n+1\}$ , we denote the  $i^{\text{th}}$  configuration of  $\mathcal{R}$  and  $\mathcal{S}$  respectively as  $\rho_i = (s_i^{\mathcal{R}}, b_i^{\mathcal{R}})$  and  $\sigma_i = (s_i^{\mathcal{S}}, b_i^{\mathcal{S}})$ . We inductively specify the message deliveries of  $\mathcal{S}$  so that the following properties are satisfied for each  $i \in \{1, \dots, n+1\}$ :

1.  $s_i^{\mathcal{S}}$  and  $s_i^{\mathcal{R}}$  have the same output and memory  $C$ -facts;

2. message buffer  $b_i^S$  a submultiset of message buffer  $b_i^R$ ; and,
3.  $\gamma_i$  is a submultiset of the message buffer  $b_i^S$ .

The need for the first two properties was already explained above, and property 3 helps in proving them. For the base case ( $i = 1$ ), properties 1 and 2 are satisfied because  $\rho_1$  and  $\sigma_1$  are start configurations, in which there are no output or memory facts and the message buffers are empty; and, property 3 is satisfied because  $\gamma_1 = \emptyset$ , which follows from  $b_1^R = \emptyset$  and the property  $\gamma_1 \subseteq b_1^R$  of  $msg_{\mathcal{R}}$ . For the induction hypothesis, we assume that the properties are satisfied for  $\rho_i$  and  $\sigma_i$ . For the inductive step, we show that they are satisfied for  $\rho_{i+1}$  and  $\sigma_{i+1}$ . In transition  $i$  of  $\mathcal{S}$ , which transforms  $\sigma_i$  into  $\sigma_{i+1}$ , we deliver the following message *multiset*:

$$m_i^S = (b_i^S \setminus (\gamma_i \setminus \beta_i)) \cap m_i^R,$$

where  $m_i^R$  denotes the message multiset delivered in transition  $i$  of  $\mathcal{R}$ , and where we use multiset difference and intersection. Intuitively, the set  $\beta_i$  of messages needed in transition  $i$ , is delivered, but we have to protect the messages in  $\gamma_i \setminus \beta_i$ , because they are needed *after* transition  $i$ . All remaining facts can be delivered, on condition that they are delivered in  $\mathcal{R}$ .

The following subsections show the properties 1 to 3.

#### 5.4.1 Property 1

We show that  $s_{i+1}^S$  and  $s_{i+1}^R$  contain the same output and memory  $C$ -facts. First, because  $m_i^S \subseteq m_i^R$ , Claim B.6 tells us that the output and memory  $C$ -facts of  $s_{i+1}^S$  are a subset of those in  $s_{i+1}^R$ . For the other direction, let  $\mathbf{g}$  be an output or memory  $C$ -fact in  $s_{i+1}^R \setminus s_{i+1}^S$ . Because  $\mathbf{g}$  is a  $C$ -fact, the mapping  $hist_{\mathcal{R}}(i, \mathbf{g}) = (\varphi, V)$  is defined, where valuation  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}$  and derives  $\mathbf{g}$ . We show that this is also true during transition  $i$  of  $\mathcal{S}$ , so that  $\mathbf{g} \in s_{i+1}^S$ . We look at the different components in the body of  $\varphi$ :

- Consider the input atoms. Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{in}}$ . We have to show  $\mathbf{h} \in J$ . First, because  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}$ , we have  $\mathbf{h} \in I$ . Moreover, because  $\mathbf{h}$  is an input fact needed in  $hist_{\mathcal{R}}$ , we have  $\mathbf{h} \in K$  (Section 5.3). Hence,  $\mathbf{h} \in I^{[adom(K)]} \subseteq I^{[adom(K_1) \cup adom(K_2)]} = J$ .  
Let  $\mathbf{h} \in V(neg^\varphi)|_{\Upsilon_{in}}$ . We have to show  $\mathbf{h} \notin J$ . This follows from  $\mathbf{h} \notin I$  (because  $V$  is satisfying in  $\mathcal{R}$ ) and  $J \subseteq I$ .
- Consider the message atoms. Recall that  $\varphi$  is message-positive. Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{msg}}$ . We have to show that  $\mathbf{h}$  is delivered in transition  $i$  of  $\mathcal{S}$ , i.e.,  $\mathbf{h} \in set(m_i^S)$ . Because  $\mathbf{h}$  is a message needed in  $hist_{\mathcal{R}}$ , there is a triple  $(k, \mathbf{h}, i) \in msg_{\mathcal{R}}$  for some  $k < i$ . Hence,  $\mathbf{h} \in \beta_i$ . Finally, Claim B.3 applied to  $\gamma_i \subseteq b_i^S$  (induction hypothesis) gives  $\beta_i \subseteq set(m_i^S)$ .
- Consider the output and memory atoms. Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$ . We have to show  $\mathbf{h} \in s_i^S$ . First, because  $V$  is satisfying in  $\mathcal{R}$ , we have  $\mathbf{h} \in s_i^R$ . Moreover, because  $\mathbf{g}$  is a  $C$ -fact, the message-boundedness of  $\varphi$  implies that  $\mathbf{h}$  is a  $C$ -fact. Hence,  $\mathbf{h} \in s_i^S$  by the induction hypothesis. Similarly, for each  $\mathbf{h} \in V(neg^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$  we can show  $\mathbf{h} \notin s_i^S$ .
- The nonequalities of  $\varphi$  are satisfied under  $V$  in  $\mathcal{R}$ , hence in  $\mathcal{S}$  as well.

We conclude that  $V$  is satisfying for  $\varphi$  in transition  $i$  of  $\mathcal{S}$ .

### 5.4.2 Property 2

We show  $b_{i+1}^S \sqsubseteq b_{i+1}^R$ . By the operational semantics,  $b_{i+1}^S = (b_i^S \setminus m_i^S) \cup \delta_i^S$  and  $b_{i+1}^R = (b_i^R \setminus m_i^R) \cup \delta_i^R$ , where  $\delta_i^S$  and  $\delta_i^R$  denote the set of messages sent in transition  $i$  of  $\mathcal{S}$  and  $\mathcal{R}$  respectively. Because  $\delta_i^S \subseteq \delta_i^R$  by Claim B.6, it is sufficient to show  $b_i^S \setminus m_i^S \sqsubseteq b_i^R \setminus m_i^R$ . Let  $\mathbf{g}$  be an arbitrary fact. We show  $\text{num}(\mathbf{g}, b_i^S \setminus m_i^S) \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ .

Because  $m_i^S \sqsubseteq b_i^S$ , we have  $\text{num}(\mathbf{g}, b_i^S \setminus m_i^S) = \text{num}(\mathbf{g}, b_i^S) - \text{num}(\mathbf{g}, m_i^S)$ . Applying the definition of  $m_i^S$  further gives

$$\begin{aligned} \text{num}(\mathbf{g}, b_i^S \setminus m_i^S) &= \text{num}(\mathbf{g}, b_i^S) - \min\{\text{num}(\mathbf{g}, b_i^S \setminus (\gamma_i \setminus \beta_i)), \text{num}(\mathbf{g}, m_i^R)\} \\ &= \max\{e_1, e_2\}, \end{aligned}$$

where

$$\begin{aligned} e_1 &= \text{num}(\mathbf{g}, b_i^S) - \text{num}(\mathbf{g}, b_i^S \setminus (\gamma_i \setminus \beta_i)), \text{ and} \\ e_2 &= \text{num}(\mathbf{g}, b_i^S) - \text{num}(\mathbf{g}, m_i^R). \end{aligned}$$

We show that both  $e_1 \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$  and  $e_2 \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ .

- We show  $e_1 \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ . First, rewriting  $e_1 = \text{num}(\mathbf{g}, b_i^S \setminus (b_i^S \setminus (\gamma_i \setminus \beta_i)))$  and applying  $\gamma_i \setminus \beta_i \sqsubseteq b_i^S$  (follows from induction hypothesis  $\gamma_i \sqsubseteq b_i^S$ ), we obtain  $e_1 = \text{num}(\mathbf{g}, \gamma_i \setminus \beta_i)$ .

Now, since  $\gamma_{i+1} = (\gamma_i \setminus \beta_i) \cup \mathcal{E}_i$  (Claim B.2), we further have  $e_1 = \text{num}(\mathbf{g}, \gamma_{i+1} \setminus \mathcal{E}_i)$ . If we can show  $\text{num}(\mathbf{g}, \gamma_{i+1} \setminus \mathcal{E}_i) = \text{num}(\mathbf{g}, \gamma_{i+1} \setminus \delta_i^R)$  then  $\gamma_{i+1} \sqsubseteq b_{i+1}^R$  (property of  $\text{msg}_{\mathcal{R}}$ ) implies  $e_1 \leq \text{num}(\mathbf{g}, b_{i+1}^R \setminus \delta_i^R) = \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ , as desired.

To show  $\text{num}(\mathbf{g}, \gamma_{i+1} \setminus \mathcal{E}_i) = \text{num}(\mathbf{g}, \gamma_{i+1} \setminus \delta_i^R)$ , it suffices to show that if  $\mathbf{g} \in \gamma_{i+1}$  then  $\text{num}(\mathbf{g}, \delta_i^R) = \text{num}(\mathbf{g}, \mathcal{E}_i)$ . This equality holds, because  $\text{msg}_{\mathcal{R}}$  satisfies  $\mathcal{E}_i = \gamma_{i+1} \cap \delta_i^R$ .

- We show  $e_2 \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ . We have  $b_i^S \sqsubseteq b_i^R$  by the induction hypothesis. Hence,  $e_2 \leq \text{num}(\mathbf{g}, b_i^R) - \text{num}(\mathbf{g}, m_i^R)$ . But since  $m_i^R \sqsubseteq b_i^R$ , we may write  $e_2 \leq \text{num}(\mathbf{g}, b_i^R \setminus m_i^R)$ , as desired.

### 5.4.3 Property 3

We show  $\gamma_{i+1} \sqsubseteq b_{i+1}^S$ . First, Claim B.2 tells us that  $\gamma_{i+1} = (\gamma_i \setminus \beta_i) \cup \mathcal{E}_i$ . It is sufficient to show  $\gamma_i \setminus \beta_i \sqsubseteq b_i^S \setminus m_i^S$  and  $\mathcal{E}_i \subseteq \delta_i^S$  because then  $\gamma_{i+1} \sqsubseteq (b_i^S \setminus m_i^S) \cup \delta_i^S = b_{i+1}^S$ , as desired.

We show that  $\gamma_i \setminus \beta_i \sqsubseteq b_i^S \setminus m_i^S$ . First, from the definition of  $m_i^S$ , we get  $m_i^S \sqsubseteq b_i^S \setminus (\gamma_i \setminus \beta_i)$ . By adding  $\gamma_i \setminus \beta_i$  to both sides of this inclusion, and using  $\gamma_i \setminus \beta_i \sqsubseteq b_i^S$  (by induction hypothesis  $\gamma_i \sqsubseteq b_i^S$ ), we obtain  $m_i^S \cup (\gamma_i \setminus \beta_i) \sqsubseteq b_i^S$ . Hence,  $\gamma_i \setminus \beta_i \sqsubseteq b_i^S \setminus m_i^S$ .

We show that  $\mathcal{E}_i \subseteq \delta_i^S$ . Let  $\mathbf{g} \in \mathcal{E}_i$ . By definition of  $\mathcal{E}_i$ , there is a triple  $(i, \mathbf{g}, l) \in \text{msg}_{\mathcal{R}}$  for some  $l > i$ , i.e.,  $\mathbf{g}$  is a needed message that should be sent in transition  $i$ . By construction of  $\text{hist}_{\mathcal{R}}$ , the mapping  $\text{hist}_{\mathcal{R}}(i, \mathbf{g}) = (\varphi, V)$  is defined, where valuation  $V$  is satisfying for rule  $\varphi$  during transition  $i$  of  $\mathcal{R}$  and derives  $\mathbf{g}$ . We show that  $V$  is also satisfying for  $\varphi$  in transition  $i$  of  $\mathcal{S}$ , which gives  $\mathbf{g} \in \delta_i^S$ . This goes similarly as in property 1, where we showed that the  $C$ -facts of  $s_{i+1}^R$  are in  $s_{i+1}^S$ , except that this time we only have to consider input atoms, message atoms and nonequalities of  $\varphi$  (because sending rules are static).

## 6 Decidability

Note, Proposition 5.1 does not immediately give decidability of inconsistency for simple transducer networks because even on a fixed input instance, we still have an infinite state system since the message buffers have no size limit. In this section we show that inconsistency of simple single-node transducer networks is decidable. In Section 6.1, we give a nondeterministic exponential time (NEXPTIME) decision procedure. In Section 6.2, we give a NEXPTIME lower bound, thus making the problem NEXPTIME-complete. This also makes inconsistency for multi-node networks NEXPTIME-complete: (i) the NEXPTIME upper bound follows from the PTIME reduction to a single-node network (Section 4), and (ii) the NEXPTIME lower bound is because single-node networks are a special case of multi-node networks.

### 6.1 Decision Procedure

In Section 6.1.1 we give the description of the decision procedure. Next, Sections 6.1.2 and 6.1.3 investigate the correctness, and Section 6.1.4 investigates the complexity.

Let  $\mathcal{N}$  be a simple single-node transducer network. Let  $\Pi$  and  $\Upsilon$  respectively denote the transducer and transducer schema of  $\mathcal{N}$ . We use the syntactical simplifications for single-node networks (Section 4.1).

#### 6.1.1 Procedure

We give a nondeterministic procedure for checking whether  $\mathcal{N}$  is inconsistent. We say that the procedure *accepts*  $\mathcal{N}$  if at least one computation branch has found evidence that  $\mathcal{N}$  is inconsistent, in which case that branch executes the *accept*-statement. A branch can also stop early by executing *reject*.

Let  $\mathbf{p}$ ,  $\mathbf{B}$ ,  $\mathbf{c}$ , and  $\text{sizeDom}(\mathcal{N})$  be as defined in Section 5.1. Consider the expression  $\text{runLen} = \mathbf{cB}^{\mathbf{p}} + \mathbf{c}$ . For  $A \subseteq \text{dom}$ , we say that a fact  $\mathbf{f}$  is a  $A$ -fact if  $\text{adom}(\mathbf{f}) \subseteq A$ . The procedure does the following steps, in order:

1. [Input] Guess an input instance  $I$  for  $\mathcal{N}$  with  $|\text{adom}(I)| \leq \text{sizeDom}(\mathcal{N})$ .
2. [Two runs] Guess two runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  of  $\mathcal{N}$  on input  $I$ , such that both runs do at most  $\text{runLen}$  transitions. Concretely, such a run is guessed by first choosing how much transitions are done ( $\leq \text{runLen}$ ), and by choosing for each transition which submultiset of the message buffer should be delivered. For simulating these runs, it is sufficient to store only the last configuration, and not all previous configurations.
3. [Output] Choose an output fact  $\mathbf{f}$  in  $\text{last}(\mathcal{S}_1)$  that is not in  $\text{last}(\mathcal{S}_2)$ . If no such fact can be chosen, then *reject*.
4. [Extension] Denote  $C = \text{adom}(\mathbf{f})$ . We extend  $\mathcal{S}_2$  by doing  $\mathbf{p} + 1$  more transitions, and in each transition we deliver the entire message buffer. If no output or memory  $C$ -fact is created in this extension, then *accept* and else *reject*.

### 6.1.2 Correctness Part 1

Suppose that  $\mathcal{N}$  is inconsistent. We show that the procedure accepts. Helper claims can be found in Appendix C.1.

First, by the small model property (Section 5), there is an input  $I$  for  $\mathcal{N}$  such that  $|dom(I)| \leq sizeDom(\mathcal{N})$  and  $\mathcal{N}$  is inconsistent on input  $I$ . Thus, there are two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  of  $\mathcal{N}$  on input  $I$  such that  $last(\mathcal{R}_1)$  contains an output fact  $\mathbf{f}$  that is not in  $last(\mathcal{R}_2)$ , and there is no extension of  $\mathcal{R}_2$  in which  $\mathbf{f}$  can be output. The procedure can guess an instance  $I'$  that is isomorphic to  $I$ , but for notational simplicity we may assume that simply  $I' = I$ .

Denote  $C = adom(\mathbf{f})$ . By inflationarity of  $\Pi$ , we can always extend  $\mathcal{R}_2$  to a run  $\mathcal{R}'_2$  such that no more output or memory  $C$ -facts can be created in any extension of  $\mathcal{R}'_2$ . By assumption on  $\mathcal{R}_2$ , configuration  $last(\mathcal{R}'_2)$  does not contain  $\mathbf{f}$ . We now convert  $\mathcal{R}_1$  and  $\mathcal{R}'_2$  to runs that the procedure can guess: by Claim C.1, there exists two runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  of  $\mathcal{N}$  on input  $I$  with at most **runLen** transitions such that  $last(\mathcal{S}_1)$  and  $last(\mathcal{S}_2)$  contain exactly the same output and memory  $C$ -facts as respectively  $last(\mathcal{R}_1)$  and  $last(\mathcal{R}'_2)$ . Hence,  $last(\mathcal{S}_1)$  contains  $\mathbf{f}$  and  $last(\mathcal{S}_2)$  does not. So, the procedure can choose  $\mathbf{f}$  as the output fact to focus on.

Next, let  $\mathcal{S}'_2$  denote the extension of  $\mathcal{S}_2$  as performed by the procedure: we do  $\mathbf{P} + 1$  additional transitions, in each of which we deliver the entire message buffer. We show that no more output or memory  $C$ -facts are created in this extension, so that the procedure accepts, as desired. Towards a proof by contradiction, suppose that there is some new transition  $i \in \{1, \dots, \mathbf{P} + 1\}$  that derives an output or memory  $C$ -fact  $\mathbf{g}$ , with the assumption that  $i$  is the *first* such transition. Let  $(\varphi, V)$  be a derivation pair for  $\mathbf{g}$  in transition  $i$ . We show that  $\mathcal{R}'_2$  can be extended to output  $\mathbf{g}$  as well, giving the desired contradiction.

Extend  $\mathcal{R}'_2$  to a run  $\mathcal{R}''_2$  by doing  $\mathbf{P} + 1$  more transitions in each of which we also deliver the entire message buffer. We show that  $V$  is satisfying for  $\varphi$  in the last transition of  $\mathcal{R}''_2$ . We consider the different body components of  $\varphi$ :

- The input literals of  $\varphi$  are satisfied under  $V$  in the last transition of  $\mathcal{R}''_2$  because  $\mathcal{S}'_2$  and  $\mathcal{R}''_2$  have the same input  $I$ .
- Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{msg}}$ . Because  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}'_2$ , message  $\mathbf{h}$  can be sent, and then Claim C.2 can be applied to know that  $\mathbf{h}$  is delivered in the last transition of  $\mathcal{R}''_2$ .
- Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$ . We have to show that  $\mathbf{h}$  is available in the last transition of  $\mathcal{R}''_2$ . First, because  $\mathbf{g}$  is a  $C$ -fact, the message-boundedness of  $\varphi$  implies that  $\mathbf{h}$  is a  $C$ -fact. Because  $\mathbf{g}$  is assumed to be the first output or memory  $C$ -fact to be created in the extension of  $\mathcal{S}_2$ , fact  $\mathbf{h}$  is in  $last(\mathcal{S}_2)$ . Thus  $\mathbf{h}$  is in  $last(\mathcal{R}'_2)$  by construction of  $\mathcal{S}_2$ , so  $\mathbf{h}$  can be read in the last transition of  $\mathcal{R}''_2$ .
- Let  $\mathbf{h} \in V(neg^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$ . We have to show that  $\mathbf{h}$  is not read in the last transition of  $\mathcal{R}''_2$ . Like in the previous case,  $\mathbf{h}$  is a  $C$ -fact. It is sufficient to show that  $\mathbf{h}$  is not in  $last(\mathcal{R}'_2)$  because no output or memory  $C$ -fact can be created in an extension of  $\mathcal{R}'_2$ , including  $\mathcal{R}''_2$ . Now, because  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}'_2$ , the inflationarity of transducer  $\Pi$  implies that  $\mathbf{h}$  is not in  $last(\mathcal{S}_2)$ . Thus  $\mathbf{h}$  is not in  $last(\mathcal{R}'_2)$  by construction of  $\mathcal{S}_2$ .

- Also, the nonequalities of  $\varphi$  are satisfied under  $V$  in  $\mathcal{R}_2''$ .

### 6.1.3 Correctness Part 2

Suppose that the procedure accepts. We show that  $\mathcal{N}$  is inconsistent.

Because the procedure accepts, there is a computation branch that has done the following. The branch has guessed an input instance  $I$  for  $\mathcal{N}$  such that  $|adom(I)| \leq sizeDom(\mathcal{N})$ . Next, the branch has guessed two runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  of  $\mathcal{N}$  on input  $I$ , and has been able to choose an output fact  $\mathbf{f}$  in  $last(\mathcal{S}_1)$  that is not in  $last(\mathcal{S}_2)$ . Denote  $C = adom(\mathbf{f})$ . Lastly, the branch has extended  $\mathcal{S}_2$  to a run  $\mathcal{S}_2'$  by doing  $\mathbf{P} + 1$  additional transitions in which the entire message buffer is delivered each time, and the procedure has observed that no output or memory  $C$ -facts were created in this extension, including  $\mathbf{f}$ .

To show that  $\mathcal{N}$  is inconsistent, it is sufficient to show that no output or memory  $C$ -facts (including  $\mathbf{f}$ ) can be created in any extension of  $\mathcal{S}_2'$ . Towards a proof by contradiction, suppose that an output or memory  $C$ -fact  $\mathbf{g}$  can be created in an extension  $\mathcal{S}_2''$  of  $\mathcal{S}_2'$ . Let us assume that  $\mathbf{g}$  is the *first* such output or memory  $C$ -fact. Let  $\varphi$  and  $V$  be a rule and valuation that are responsible for deriving  $\mathbf{g}$ . We show that  $V$  is satisfying for  $\varphi$  in the last transition of  $\mathcal{S}_2'$  itself, so that  $\mathbf{g}$  would already have been created in  $\mathcal{S}_2'$ , which is the desired contradiction. To show that  $V$  is satisfying in  $\mathcal{S}_2'$ , we proceed similarly as in the first correctness proof above. We note the differences:

- Let  $\mathbf{h} \in V(pos^\varphi)|_{\Gamma_{out} \cup \Gamma_{mem}}$ . We have to show that  $\mathbf{h}$  is available in the last transition of  $\mathcal{S}_2'$ . Like before,  $\mathbf{h}$  is a  $C$ -fact by message-boundedness. Because  $\mathbf{g}$  is assumed to be the first output or memory  $C$ -fact to be created in the extension of  $\mathcal{S}_2'$ , it must be that  $\mathbf{h}$  is in  $last(\mathcal{S}_2')$ . Moreover, because the decision procedure has not observed the creation of an output or memory  $C$ -fact in the transitions of  $\mathcal{S}_2'$  after  $last(\mathcal{S}_2)$ , fact  $\mathbf{h}$  is in  $last(\mathcal{S}_2)$ . Hence,  $\mathbf{h}$  can be read in the last transition of  $\mathcal{S}_2'$ .
- Let  $\mathbf{h} \in V(neg^\varphi)|_{\Gamma_{out} \cup \Gamma_{mem}}$ . We have to show that  $\mathbf{h}$  is not present in the last transition of  $\mathcal{S}_2''$ . Because  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}_2''$ , fact  $\mathbf{h}$  must be absent there. Hence, by inflationarity,  $\mathbf{h}$  is not in  $last(\mathcal{S}_2)$ .

### 6.1.4 Time Complexity

Here we analyze the time complexity of each computation branch of the decision procedure. We sketch how the procedure might be implemented in an imperative programming language where blocks of code can be guarded by a nondeterministic choice, that could either execute the corresponding block or skip it. In this framework, we show that each branch uses at most single-exponential time, making the decision procedure be in NEXPTIME.

**Encoding** We use the encoding of transducer networks from Section 2.8. Let  $|\mathcal{N}|$  denote the input size. Now, consider the syntactical quantities defined in Section 5.1. The quantities  $\mathbf{I}$  and  $\mathbf{O}$  are upper bounded by  $|\mathcal{N}|$  because all input and output relations are used in rules (whose atoms are written in full). The quantities  $\mathbf{B}$  and  $\mathbf{P}$  are also upper bounded by  $|\mathcal{N}|$ . Letting  $n$  denote the number of different transducer relations, again upper bounded by  $|\mathcal{N}|$ , the number  $\mathbf{C}$  is

upper bounded by  $n\mathbf{O}^{\mathbf{O}} = n2^{\mathbf{O} \log \mathbf{O}}$ , which is single-exponential in  $|\mathcal{N}|$ . Hence,  $\text{sizeDom}(\mathcal{N})$  is single-exponential in  $|\mathcal{N}|$ .

Let  $\mathbf{numFc}$  denote the number of different facts that can be created with  $\text{sizeDom}(\mathcal{N})$  unique domain values (across all relations). Note that  $\mathbf{numFc}$  is single-exponential in  $|\mathcal{N}|$ .

**Input** For each input instance  $I'$  for  $\mathcal{N}$  with  $|\text{adom}(I')| \leq \text{sizeDom}(\mathcal{N})$ , the procedure can guess an isomorphic instance  $I$ . Because  $\text{sizeDom}(\mathcal{N})$  is single-exponential in  $|\mathcal{N}|$ , an active domain value of  $I$  can be represented as a number encoded by  $p$  bits, where  $p$  is polynomial in  $|\mathcal{N}|$ . We omit the algorithmic details to guess  $I$ .

**Two runs** Next, the procedure needs to guess two runs  $\mathcal{S}_1$  and  $\mathcal{S}_2$  of  $\mathcal{N}$  on  $I$ , such that each run does at most  $\mathbf{runLen}$  transitions. We describe how to guess one run  $\mathcal{S} \in \{\mathcal{S}_1, \mathcal{S}_2\}$ ; the other run can be guessed similarly after the first one.

To guess  $\mathcal{S}$ , we do a for-loop with  $\mathbf{runLen}$  iterations in which we incrementally modify a configuration, starting with the start configuration. Note that  $\mathbf{runLen}$  is single-exponential in  $|\mathcal{N}|$ . In each iteration, we choose whether or not we do a transition. To do a transition, we select a submultiset  $m$  of the message buffer to deliver. The size of the message buffer is at most  $\mathbf{runLen} \cdot \mathbf{numFc}$ , so this selection can be done in single-exponential time. We are left to show that simulating the subsequent local transition can be done in single-exponential time. Let  $J$  denote the transducer state in the last configuration obtained. Now, for all transducer rules  $\varphi$ , for all valuations  $V$  for  $\varphi$ , if  $V$  is satisfying for  $\varphi$  with respect to  $J \cup \text{set}(m)$  then derive  $\mathbf{g} = V(\text{head}^\varphi)$ . The number of rules is linear in  $|\mathcal{N}|$ . For one rule, the number of variables is also linear in  $|\mathcal{N}|$ . Hence, the number of valuations for one rule, using values in  $\text{adom}(I)$ , is single-exponential in  $|\mathcal{N}|$ . Finally, checking whether a valuation  $V$  is satisfying for a rule  $\varphi$  is done by (i) checking that the nonequalities are satisfied, which can be done in polynomial time; and, (ii) going over all body literals  $\mathbf{l}$  of  $\varphi$ , applying  $V$ , and checking whether  $J \cup \text{set}(m) \models V(\mathbf{l})$ , which can be done in single-exponential time because  $|J \cup \text{set}(m)| \leq \mathbf{numFc}$ .

**Output** The procedure then selects an output fact  $\mathbf{f}$  in  $\text{last}(\mathcal{S}_1)$  that is not in  $\text{last}(\mathcal{S}_2)$ . Because the number of output facts in  $\text{last}(\mathcal{S}_1)$  is at most  $\mathbf{numFc}$ , we can select  $\mathbf{f}$  in single-exponential time. Possibly  $\text{last}(\mathcal{S}_2)$  has at least the output facts of  $\text{last}(\mathcal{S}_1)$ , in which case the procedure does *reject*. Otherwise, we continue.

**Extension** In the last step, the procedure extends  $\mathcal{S}_2$  with  $\mathbf{P} + 1$  transitions, in each of which we deliver the entire message buffer. The message buffer in  $\text{last}(\mathcal{S}_2)$  contains at most  $\mathbf{runLen} \cdot \mathbf{numFc}$  facts, and all the subsequent buffers in the extension contain at most  $\mathbf{numFc}$  facts because the buffer has degenerated to a set. Hence, we can apply the same time complexity analysis for simulating the local transitions as above.

Letting  $C = \text{adom}(\mathbf{f})$ , checking whether a newly derived output or memory fact is a  $C$ -fact can be done in polynomial time. Overall, simulating the additional  $\mathbf{P} + 1$  transitions can be done in single-exponential time.



## 6.2 Complexity Lower Bound

In Section 6.1 we gave a NEXPTIME upper bound on the time complexity for deciding inconsistency for simple single-node transducer networks. In this section, we complement this result by giving a NEXPTIME lower bound, making the decision problem NEXPTIME-complete. Concretely, we show that any problem in NEXPTIME is polynomial time reducible to this decision problem.

Let  $A$  be a problem from NEXPTIME. Formally,  $A$  is a set of words over some alphabet  $\Sigma$ , and there exists a nondeterministic Turing machine  $M$  such that (i) for each word  $w$  over  $\Sigma$ ,  $M$  accepts  $w$  iff  $w \in A$ ; and, (ii) every computation trace of  $M$  on an input  $w$  over  $\Sigma$  eventually halts and uses at most  $O(2^{|w|^k})$  steps, where  $k$  is a constant specific to  $M$  [20].

Fix some word  $w$  over  $\Sigma$ . We construct a simple single-node transducer network  $\mathcal{N}$  for  $w$  such that  $\mathcal{N}$  is inconsistent iff  $M$  accepts  $w$ . We use the syntactical simplifications of single-node networks (Section 4.1).

### 6.2.1 Turing Machine

First, following the conventions in Sipser [20], the Turing machine  $M$  is given as a tuple

$$(Q, \Sigma, \Gamma, \delta, q_0, q_{\text{accept}}, q_{\text{reject}}),$$

where  $Q$  is the set of states,  $\Sigma$  is the alphabet of the language  $A$ ,  $\Gamma$  is the tape-alphabet (satisfying  $\Sigma \subseteq \Gamma$ ),  $\delta$  is the transition function,  $q_0 \in Q$  is the start state,  $q_{\text{accept}} \in Q$  is the accept state, and  $q_{\text{reject}} \in Q$  is the reject state. Function  $\delta$  has the signature  $Q \times \Gamma \rightarrow \mathcal{P}(Q \times \Gamma \times \{L, R\})$ , where L and R indicate whether the tape head moves left or right after performing a transition.

### 6.2.2 Construction

We define the transducer schema  $\Upsilon$  and transducer  $\Pi$  of  $\mathcal{N}$ . The main idea is as follows. We provide  $\Upsilon$  with input relations to encode a computation trace of the Turing machine  $M$  on input  $w$ . By simulating the Turing machine  $M$ , transducer  $\Pi$  checks that the input contains a valid and accepting computation trace. If so,  $\Pi$  sends a special message `accept()` to itself, whose delivery is a trigger for inconsistent behaviour. On a more technical note, the sending rules might sometimes send `accept()` when the trace is actually partially incorrect. To solve this, like in Section 4, we also check explicitly for errors in the input: when an error is detected, a message `error()` is sent, and this acts as a signal to correct any inconsistent behavior.

**Inconsistency** Independently of  $w$  or  $M$ , we add the following relations to  $\Upsilon$ : input relation  $A^{(1)}$ ; memory relation  $B^{(1)}$ ; output relation  $T^{(1)}$ ; and, message relations  $A_{\text{msg}}^{(1)}$ ,  $B_{\text{msg}}^{(1)}$ , `accept`<sup>(0)</sup> and `error`<sup>(0)</sup>. The following rules implement the basic idea of making  $\Pi$  inconsistent when `accept()` is received; we can vary the delivery order of  $A_{\text{msg}}$ -facts and  $B_{\text{msg}}$ -facts. The purpose of relation `error` was explained above.

Relation	Purpose
<b>state</b> <sup>(2)</sup>	configuration state
<b>head</b> <sup>(1+n<sup>k</sup>)</sup>	configuration head position
<b>tape</b> <sup>(1+n<sup>k</sup>+1)</sup>	configuration tape cell contents

Table 1: Computation trace input relations

$$A_{\text{msg}}(\mathbf{u}) \leftarrow A(\mathbf{u}), \text{accept}().$$

$$B_{\text{msg}}(\mathbf{u}) \leftarrow A(\mathbf{u}), \text{accept}().$$

$$B(\mathbf{u}) \leftarrow B_{\text{msg}}(\mathbf{u}).$$

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}), \neg B(\mathbf{u}).$$

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}), \text{error}().$$

**Computation trace** We represent a computation trace of  $M$  on  $w$  with new input relations. Henceforth we write  $n$  to denote the length of  $w$ . We can select a  $k \in \mathbb{N}$  such that for each string  $w'$  over  $\Sigma$ , if  $M$  accepts  $w'$  then  $M$  has an accepting computation trace on  $w'$  with at most  $2^{n^k}$  transitions. Note,  $k$  is considered a constant in the construction of the transducer.

A number  $a$  in the interval  $[0, 2^{n^k}]$  indicates a (zero-based) configuration ordinal in the trace. Moreover, since time usage upper bounds space usage,  $a$  can also be used to indicate an individual tape cell. The number  $a$  has a binary representation with  $n^k$  bits, which is polynomial in  $n$ . Now, Table 1 gives the input relations, with their precise arities, to represent a computation trace. The first component in relations **state**, **head**, and **tape**, is an identifier of a Turing machine configuration. This identifier only serves to join the different aspects of one configuration across all three relations: relation **state** gives the current state symbol; relation **head** gives the head position; and, relation **tape** gives the contents of each tape cell.

**Sending accept** We now provide rules to send **accept**( ). Newly mentioned relations are assumed to be added to  $\Upsilon_{\text{msg}}$ . The idea is as follows: in the relations of Table 1, we look for a path of length at most  $2^{n^k}$  configurations that connects the start configuration to an accepting configuration, and such that each pair of subsequent configurations is allowed by a valid transition of  $M$ .

Suppose we could send a message of the form **reach**<sub>0</sub>( $i, j$ ) to say that configuration  $j$  can be reached from configuration  $i$  by a valid transition of  $M$ . The subscript 0 indicates that the distance between  $i$  and  $j$  is  $2^0 = 1$ . Since the desired path is of length at most  $2^{n^k}$ , the following *recursion-free* rules can consider all such paths:<sup>9</sup>

<sup>9</sup>We use that any length between 0 and  $2^{n^k}$  can be represented by a sum of unique powers of two.

$$\mathbf{reach}_m(\mathbf{i}, \mathbf{j}) \leftarrow \mathbf{reach}_{m-1}(\mathbf{i}, \mathbf{l}), \mathbf{reach}_p(\mathbf{l}, \mathbf{j})$$

for each  $m = 1, \dots, n^k$ , and each  $p = 0, \dots, m-1$ .

Suppose too we could send a message of the form  $\mathbf{start}(i)$  to say that configuration  $i$  satisfies the properties of the start configuration of  $M$  on  $w$ . We send  $\mathbf{accept}()$  with these rules:

$$\mathbf{accept}() \leftarrow \mathbf{start}(\mathbf{i}), \mathbf{reach}_m(\mathbf{i}, \mathbf{j}), \mathbf{state}(\mathbf{j}, \mathbf{q}), q_{\mathbf{accept}}(\mathbf{q})$$

for each  $m = 0, \dots, n^k$ .

Here,  $q_{\mathbf{accept}}^{(1)}$  is an extra input relation containing the symbol of the start state.

Note, the number and size of the above sending rules is polynomial in  $n$ . Appendix C.2 fills in the missing details regarding the messages  $\mathbf{reach}_0$ ,  $\mathbf{start}$ ,  $\mathbf{error}$ , and argues the correctness.

## 7 Expressivity

We investigate the expressivity of simple transducer networks. First we define how a transducer network can compute a distributed query. We consider only *consistent* transducer networks because otherwise the output might vary depending on the run. Let  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  be a consistent transducer network, not necessarily simple. Let  $\mathbf{in}^{\mathcal{N}}$  and  $\mathbf{out}^{\mathcal{N}}$  be the distributed schemas for  $\mathcal{N}$  as defined in Section 2.7. We say that  $\mathcal{N}$  *computes* the following distributed query  $\mathcal{Q}$ , that is over input schema  $\mathbf{in}^{\mathcal{N}}$  and output schema  $\mathbf{out}^{\mathcal{N}}$ :  $\mathcal{Q}$  maps each instance  $H$  over  $\mathbf{in}^{\mathcal{N}}$  to the instance  $\mathcal{Q}(H) = J$  over  $\mathbf{out}^{\mathcal{N}}$  such that  $J(x)$  for each  $x \in \mathcal{N}$  is the set of all output facts that can be produced at  $x$  during any run of  $\mathcal{N}$  on  $H$ . The instance  $\mathcal{Q}(H)$  could be defined even if  $\mathcal{N}$  is inconsistent, but when  $\mathcal{N}$  is consistent, all runs on  $H$  can be extended to obtain  $\mathcal{Q}(H)$ . We call  $\mathcal{Q}(H)$  the *output* of  $\mathcal{N}$  on input  $H$ .

We now define how  $\text{UCQ}^\neg$  can express distributed queries in a more direct way, i.e., without transducer networks. This will provide insight in the expressivity of simple transducer networks. First, for a distributed database schema  $\mathcal{E}$  over a network  $\mathcal{N}$ , and an instance  $H$  over  $\mathcal{E}$ , let  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  and  $\langle H \rangle^{\mathcal{N}}$  be as defined in Section 4.2. Intuitively, a  $\text{UCQ}^\neg$ -program over  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  can directly access all relations of all nodes. To make such a program generic, node identifiers are provided in the relations  $x.\text{Id}$  with  $x \in \mathcal{N}$  and  $\text{Node}$ . Let  $\mathcal{Q}$  be a distributed query over an input schema  $\mathcal{E}$  and an output schema  $\mathcal{F}$ , where both schemas are over the same network  $\mathcal{N}$ . We say that  $\mathcal{Q}$  is *expressible* in  $\text{UCQ}^\neg$  if for each pair  $x \in \mathcal{N}$  and  $R^{(k)} \in \mathcal{F}(x)$  we can give a  $\text{UCQ}^\neg$  program  $\Phi_{x,R}$  over input schema  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  and output schema  $\{R^{(k)}\}$  such that for all instances  $H$  over  $\mathcal{E}$  we have

$$\mathcal{Q}(H)(x)|_R = \Phi_{x,R}(\langle H \rangle^{\mathcal{N}}).$$

Now we can present the expressivity result:

**Theorem 7.1.** Consistent simple transducer networks capture the distributed queries expressible in  $\text{UCQ}^\neg$ .

This result requires showing a lower and upper bound on the expressivity of simple transducer networks. These directions are given in the following subsections. Currently, this result depends on our definition of  $\text{UCQ}^\neg$  as a language with built-in nonequalities (or equivalently by having a built-in equality relation). In particular, for showing the upper bound, we do a nontrivial simulation of runs of transducer networks with  $\text{UCQ}^\neg$ , and there we depend on the availability of nonequalities. It remains open whether the result really needs this feature.

## 7.1 Lower Bound

Let  $\mathcal{Q}$  be a distributed query over input distributed schema  $\mathcal{E}$  and output distributed schema  $\mathcal{F}$ , and that is expressible in  $\text{UCQ}^\neg$ . Let  $\mathcal{N}$  be the network of  $\mathcal{E}$  and  $\mathcal{F}$ . Over  $\mathcal{N}$ , we define a simple transducer network  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  to compute  $\mathcal{Q}$ . We assume  $\mathcal{E}(x)$  and  $\mathcal{F}(x)$  have disjoint relation names for each  $x \in \mathcal{N}$ ; that  $\mathcal{E}(x)$  and  $\mathcal{F}(x)$  do not contain  $\text{Id}$  or  $\text{All}$ ; and, that any relations we add to  $\mathcal{N}$  do not yet occur in  $\mathcal{E}$  and  $\mathcal{F}$ . Any conflicts can always be resolved with appropriate renamings.

### 7.1.1 Transducer Schemas

First, we give the shared message relations of  $\mathcal{N}$ , where relation names containing “ $\neg$ ” indicate the absence of a fact:

- the relations  $x.R^{(k)}$  and  $x.R_\neg^{(k)}$  for each  $x \in \mathcal{N}$  and  $R^{(k)} \in \mathcal{E}(x)$ , to broadcast local inputs;
- the relations  $x.\text{Id}^{(1)}$  and  $x.\text{Id}_\neg^{(1)}$  for each  $x \in \mathcal{N}$ , to broadcast identifiers;
- the relations  $x.T^{(k)}$  for each  $x \in \mathcal{N}$  and  $T^{(k)} \in \mathcal{F}(x)$ , to compute local outputs; and,
- the relation  $\text{Adom}^{(1)}$ , to share active domain values.

For each  $x \in \mathcal{N}$ , we define  $\Upsilon(x)_{\text{in}} = \mathcal{E}(x)$ ;  $\Upsilon(x)_{\text{out}} = \mathcal{F}(x)$ ;  $\Upsilon(x)_{\text{mem}} = \emptyset$ ; and,  $\Upsilon(x)_{\text{msg}}$  is the set of message relations from above.

### 7.1.2 Transducer Rules

Let  $x \in \mathcal{N}$ . We incrementally specify the rules of  $\Pi(x)$ . First, to send the active domain of the input, for each  $R^{(k)} \in \Upsilon(x)_{\text{in}} \cup \{\text{Id}^{(1)}\}$  and each  $i \in \{1, \dots, k\}$ , we add the following rule:

$$\text{Adom}(\mathbf{n}, \mathbf{u}_i) \leftarrow \text{All}(\mathbf{n}), R(\mathbf{u}_1, \dots, \mathbf{u}_i, \dots, \mathbf{u}_k).$$

Also, for each  $R^{(k)} \in \Upsilon(x)_{\text{in}} \cup \{\text{Id}^{(1)}\}$ , we add the following rules to send the presence or absence of local facts at  $x$ :

$$x.R(\mathbf{n}, \mathbf{u}_1, \dots, \mathbf{u}_k) \leftarrow \text{All}(\mathbf{n}), R(\mathbf{u}_1, \dots, \mathbf{u}_k).$$

$$x.R_\neg(\mathbf{n}, \mathbf{u}_1, \dots, \mathbf{u}_k) \leftarrow \text{All}(\mathbf{n}), \text{Adom}(\mathbf{u}_1), \dots, \text{Adom}(\mathbf{u}_k), \neg R(\mathbf{u}_1, \dots, \mathbf{u}_k).$$

Now we let  $\Pi(x)$  produce output. Let  $T^{(k)} \in \Upsilon(x)_{\text{out}}$ . To satisfy the message-boundedness restriction for the output rules, we add sending rules for message relation  $x.T^{(k)}$  and copy any received  $x.T$ -messages to output relation  $T$ . Because  $\mathcal{Q}$  is expressible in  $\text{UCQ}^-$ , there is a  $\text{UCQ}^-$  program  $\Phi$  over  $\langle \mathcal{E} \rangle^{\mathcal{N}}$  that expresses the  $T$ -facts at  $x$ . For each  $\varphi \in \Phi$ , we transform  $\varphi$  into a sending rule  $\varphi'$  for relation  $x.T^{(k)}$ , as follows:

- the head  $T(u_1, \dots, u_k)$  of  $\varphi$  becomes the head  $x.T(n, u_1, \dots, u_k)$  of  $\varphi'$ , where  $n$  is a new variable;
- the positive body atoms of  $\varphi'$  are (i)  $\text{Id}(n)$ , with  $n$  as defined previously; (ii) the atoms  $\text{All}(m)$  for which  $\text{Node}(m) \in \text{pos}^\varphi$ ; (iii) the atoms  $y.R(v_1, \dots, v_l) \in \text{pos}^\varphi$ , which are now messages; (iv) the (positive) message atoms  $y.R_-(v_1, \dots, v_l)$  for which  $y.R(v_1, \dots, v_l) \in \text{neg}^\varphi$ ;
- the negative body atoms of  $\varphi'$  are the atoms  $\text{All}(m)$  for which  $\text{Node}(m) \in \text{neg}^\varphi$ ; and,
- the nonequalities of  $\varphi'$  are those of  $\varphi$ .

The positive body atom  $\text{Id}(n)$  has the effect that  $x.T$ -messages are sent only to  $x$ . Now, the final output for  $T^{(k)}$  is created by adding this rule:

$$T(u_1, \dots, u_k) \leftarrow x.T(u_1, \dots, u_k).$$

This completes the specification of  $\Pi(x)$ . Note that transducer  $\Pi(x)$  is simple: all message rules are message-positive and static; all output rules are message-positive and message-bounded;  $\Pi(x)$  is inflationary (there are no memory relations); and,  $\Pi(x)$  is recursion-free.

Following the above instructions, we can build the transducer at each node of  $\mathcal{N}$ . There are also no cycles through message relations in  $\mathcal{N}$ . Hence,  $\mathcal{N}$  is simple.

### 7.1.3 Example

The following example illustrates the construction of the transducer network.

**Example 7.2.** Let  $\mathcal{N} = \{x, y\}$ . Consider the following distributed schemas  $\mathcal{E}$  and  $\mathcal{F}$ , that are over  $\mathcal{N}$ :  $\mathcal{E}(x) = \{A^{(2)}\}$ ,  $\mathcal{E}(y) = \{B^{(1)}\}$ ,  $\mathcal{F}(x) = \{S^{(1)}\}$  and  $\mathcal{F}(y) = \{T^{(1)}\}$ . Consider the following distributed query  $\mathcal{Q}$  with input schema  $\mathcal{E}$  and output schema  $\mathcal{F}$ , expressed in  $\text{UCQ}^-$ :

$$S(u) \leftarrow x.A(u, v), \neg y.B(u), u \neq v.$$

$$T(u) \leftarrow x.A(u, v), x.\text{Id}(u).$$

Each rule corresponds to one of the output relations.

We construct a transducer network  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  to compute  $\mathcal{Q}$ . To save space, we will not literally follow the general construction from above, but instead restrict attention to the relations and rules that affect the output. Also, the sending rules for  $\text{Adom}$  are clear, so we do not explicitly give them.

First, the shared message relations of  $\mathcal{N}$  are:  $x.A^{(2)}$ ,  $x.\text{Id}^{(1)}$ ,  $y.B^{(1)}$  and  $\text{Adom}^{(1)}$ . For node  $x$ , we define  $\Upsilon(x)_{\text{in}} = \{A^{(2)}\}$ ,  $\Upsilon(x)_{\text{out}} = \{B^{(1)}\}$ , and  $\Upsilon(x)_{\text{mem}} = \emptyset$ . Transducer  $\Pi(x)$  contains the rules:

$$\begin{aligned}
x.A(\mathbf{n}, \mathbf{u}, \mathbf{v}) &\leftarrow \text{All}(\mathbf{n}), A(\mathbf{u}, \mathbf{v}). \\
x.\text{Id}(\mathbf{n}, \mathbf{u}) &\leftarrow \text{All}(\mathbf{n}), \text{Id}(\mathbf{u}). \\
x.S(\mathbf{n}, \mathbf{u}) &\leftarrow \text{Id}(\mathbf{n}), x.A(\mathbf{u}, \mathbf{v}), y.B_{\neg}(\mathbf{u}), \mathbf{u} \neq \mathbf{v}. \\
S(\mathbf{u}) &\leftarrow x.S(\mathbf{u}).
\end{aligned}$$

For node  $y$ , we define  $\Upsilon(y)_{\text{in}} = \{B^{(1)}\}$ ,  $\Upsilon(y)_{\text{out}} = \{T^{(1)}\}$ , and  $\Upsilon(y)_{\text{mem}} = \emptyset$ . Transducer  $y$  contains the rules:

$$\begin{aligned}
y.B_{\neg}(\mathbf{n}, \mathbf{u}) &\leftarrow \text{All}(\mathbf{n}), \text{Adom}(\mathbf{u}), \neg B(\mathbf{u}). \\
y.T(\mathbf{n}, \mathbf{u}) &\leftarrow x.A(\mathbf{u}, \mathbf{v}), x.\text{Id}(\mathbf{u}). \\
T(\mathbf{u}) &\leftarrow y.T(\mathbf{u}).
\end{aligned}$$

This completes the network  $\mathcal{N}$ . □

## 7.2 Upper Bound

Let  $\mathcal{N} = (\mathcal{N}, \Upsilon, \Pi)$  be a consistent simple transducer network. Let  $\mathcal{Q}$  denote the distributed query computed by  $\mathcal{N}$ . Let  $x \in \mathcal{N}$  and let  $R^{(k)}$  be a local output relation of  $x$ . We have to construct a UCQ<sup>-</sup>-program  $\Phi$  over input schema  $\langle \text{in}^{\mathcal{N}} \rangle^{\mathcal{N}}$  and output schema  $\{R^{(k)}\}$ , such that  $\mathcal{Q}(H)(x)|_R = \Phi(\langle H \rangle^{\mathcal{N}})$  for each input distributed database instance  $H$  over  $\text{in}^{\mathcal{N}}$ .

The basic idea is to describe the computation of  $\mathcal{N}$  with UCQ<sup>-</sup>-program  $\Phi$ , for output relation  $R$  at  $x$ . To make this technically easier, we first convert  $\mathcal{N}$  to a single-node network in Section 7.2.1. Some common notations are introduced in Section 7.2.2, and program  $\Phi$  is described in Section 7.2.3. The correctness is shown in Appendix D.

### 7.2.1 Reduction to Single-node

Consider the concepts from Section 4.2. Using Proposition 4.1, let  $\mathcal{M}$  be the simple single-node transducer network that simulates  $\mathcal{N}$ , and that is consistent because  $\mathcal{N}$  is consistent. By the syntactical simplifications of single-node networks (Section 4.1), the query  $\mathcal{Q}'$  computed by  $\mathcal{M}$  is regarded as an ordinary database query over input schema  $\langle \text{in}^{\mathcal{N}} \rangle^{\mathcal{N}}$  and output schema  $\langle \text{out}^{\mathcal{N}} \rangle$ . If for every input  $H$  for  $\mathcal{N}$  we would know that  $\mathcal{Q}'(\langle H \rangle^{\mathcal{N}}) = \langle \mathcal{Q}(H) \rangle$ , because  $x.R$  is in  $\langle \text{out}^{\mathcal{N}} \rangle$ , it will be sufficient to construct the UCQ<sup>-</sup>-program  $\Phi$  as a description of the computation of  $\mathcal{M}$  for relation  $x.R$ . To keep the notation simpler, we may assume without loss of generality that output relation  $R$  only occurs at  $x$ . So, we will write “ $R$ ” instead of “ $x.R$ ”.

Now we are left to show  $\mathcal{Q}'(\langle H \rangle^{\mathcal{N}}) = \langle \mathcal{Q}(H) \rangle$  for every input  $H$  over  $\text{in}^{\mathcal{N}}$ . Let the output of a configuration  $\rho$ , denoted  $\text{out}(\rho)$ , be as defined in Section 4.2. Abbreviate  $J = \mathcal{Q}'(\langle H \rangle^{\mathcal{N}})$ . We show  $J \subseteq \langle \mathcal{Q}(H) \rangle$ . By consistency of  $\mathcal{M}$ , there is a run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  such that  $\text{out}(\text{last}(\mathcal{S})) = J$ . Next, because  $\mathcal{M}$  simulates  $\mathcal{N}$ , there is a run  $\mathcal{R}$  of  $\mathcal{N}$  on  $H$  such that  $\langle \text{out}(\text{last}(\mathcal{R})) \rangle = \text{out}(\text{last}(\mathcal{S}))$ . So,  $J = \langle \text{out}(\text{last}(\mathcal{R})) \rangle \subseteq \langle \mathcal{Q}(H) \rangle$ . Now we show  $\langle \mathcal{Q}(H) \rangle \subseteq J$ . By consistency of  $\mathcal{N}$ , there exists a run  $\mathcal{R}$  of  $\mathcal{N}$  on  $H$  such that  $\mathcal{Q}(H) = \text{out}(\text{last}(\mathcal{R}))$ . Because  $\mathcal{M}$  simulates  $\mathcal{N}$ , there exists a run  $\mathcal{S}$  of  $\mathcal{M}$  on  $\langle H \rangle^{\mathcal{N}}$  such that  $\text{out}(\text{last}(\mathcal{S})) = \langle \text{out}(\text{last}(\mathcal{R})) \rangle$ . Hence,  $\langle \mathcal{Q}(H) \rangle = \text{out}(\text{last}(\mathcal{S})) \subseteq J$ .

### 7.2.2 Common Concepts and Notations

A *ground literal* is a fact or a fact with “ $\neg$ ” prepended. For a database instance  $I$  and a ground literal  $l$ , we write  $I \models l$  to mean  $l \in I$  if  $l$  is a fact and otherwise we mean  $f \notin I$ , where  $l = \neg f$ . For a derivation tree  $\mathcal{T}$ , for each internal node  $x$ , we write  $body^{\mathcal{T}}(x)$  to denote the set of ground literals obtained by applying  $val^{\mathcal{T}}(x)$  to the body literals of  $rule^{\mathcal{T}}(x)$ .

Two derivation trees  $\mathcal{T}$  and  $\mathcal{S}$  are said to be *structurally equivalent* if (i) the trees  $(nodes^{\mathcal{T}}, edges^{\mathcal{T}})$  and  $(nodes^{\mathcal{S}}, edges^{\mathcal{S}})$  are isomorphic under a node bijection  $b : nodes^{\mathcal{T}} \rightarrow nodes^{\mathcal{S}}$ ; and, (ii) for every edge  $(x, y) \in edges^{\mathcal{T}}$ , we have  $rule^{\mathcal{T}}(x) = rule^{\mathcal{S}}(b(x))$  and  $lit^{\mathcal{T}}(y) = lit^{\mathcal{S}}(b(y))$ . We call  $b$  the *structural bijection*.

### 7.2.3 Building the UCQ<sup>-</sup>-Program

In this section, we construct the required UCQ<sup>-</sup>-program  $\Phi$ . We gradually build up the different parts of this program, and introduce auxiliary definitions and notations along the way. Using the equivalence between UCQ<sup>-</sup> and existential FO with nonequalities, abbreviated  $\exists$ FO, some parts are specified in  $\exists$ FO for technical convenience.

Let  $\Upsilon$  and  $\Pi$  respectively denote the transducer schema and transducer of single-node transducer network  $\mathcal{M}$ .

**General derivation trees** Let  $\mathcal{T}$  be a derivation tree of  $\Pi$ . We define the *active domain* of  $\mathcal{T}$  to be the set of all values assigned by valuations in  $\mathcal{T}$ . We say that  $\mathcal{T}$  is *general* if there is no structurally equivalent derivation tree  $\mathcal{S}$  with a strictly larger active domain. Intuitively, a general derivation tree assigns a different value to each variable of a rule if possible.

**All output strategies** Let  $forest_R$  be a maximal set of general derivation trees of transducer  $\Pi$  for output relation  $R$ , such that no two trees are structurally equivalent, and such that no two trees have an overlap of their active domains. Because  $\Pi$  is recursion-free, there are only a finite number of structurally different trees, and thus  $forest_R$  is finite. Intuitively,  $forest_R$  represents all possible strategies of  $\Pi$  to derive facts over  $R$ , using as much different values as possible. For each subset  $G \subseteq forest_R$ , we write  $adom(G)$  to denote the union of all active domains of trees in  $G$ .

**Canonical runs** Intuitively, for any particular input for  $\Pi$ , we can make a selection  $G \subseteq forest_R$  of all trees that “work” on that input, i.e., for all trees  $\mathcal{T} \in G$  there is a substitution of the values in  $\mathcal{T}$  by values in the input so that the new valuations are true. If we regard values in  $adom(G)$  as variables (as we will do later), this substitution of values looks very much like a valuation. Next, for  $G$ , we can formally define a *canonical run*  $\mathcal{R}^G$ . The idea is that in  $\mathcal{R}^G$  we execute all trees of  $G$  concurrently, with as few transitions as possible, i.e., by using their canonical schedulings. The run  $\mathcal{R}^G$  will do  $n$  transitions, where  $n$  is the largest height of a tree in  $G$ .<sup>10</sup> Hence, the length of  $\mathcal{R}^G$  is bounded by the syntactical properties of  $\Pi$ .

<sup>10</sup>The height of a derivation tree is the largest number of edges on any path from a leaf to the root.

Note, for an internal node of a derivation tree  $\mathcal{T}$ , by message-positivity,  $body^\mathcal{T}(x)|_{\mathbf{r}_{\text{msg}}}$  contains only facts. Now, for each transition  $i \in \{1, \dots, n\}$  of  $\mathcal{R}^G$ , we (want to) deliver the following message set

$$M_i^G = \bigcup_{\mathcal{T} \in G} \bigcup_{\substack{x \in \text{int}^\mathcal{T}, \\ \kappa^\mathcal{T}(x) = i}} body^\mathcal{T}(x)|_{\mathbf{r}_{\text{msg}}}.$$

In words: for each transition  $i$ , set  $M_i^G$  is the union across all trees of  $G$  of the messages needed by rules scheduled at transition  $i$ . We now make an  $\exists$ FO-formula  $sndMsg_G$  to express that these message sets can be sent. For notational simplicity, the symbols of  $adom(G)$  represent variables. For a derivation tree  $\mathcal{T} \in G$ , let  $msg^\mathcal{T} \subseteq \text{int}^\mathcal{T}$  denote the set of internal nodes  $x$  where  $lit^\mathcal{T}(x)$  is over a message relation. Because sending rules are message-positive and static, it suffices to demand that all involved input literals are satisfied (both positive and negative):

$$sndMsg_G := \bigwedge_{\mathcal{T} \in G} \bigwedge_{x \in msg^\mathcal{T}} body^\mathcal{T}(x)|_{\mathbf{r}_{\text{in}}}.$$

This is a quantifier-free formula, where we write sets of literals in the conjunction, with the understanding that such a set is written using some arbitrary ordering on its elements.

**Canonical runs: output succeeds** Let  $G$  be as above. Fix some  $\mathcal{T} \in G$ . In the following, we specify an  $\exists$ FO-formula to express that  $\mathcal{T}$  succeeds in deriving its root fact in  $\mathcal{R}^G$ . Here, a possible “danger”, is that the concurrent execution of  $\mathcal{T}$  with another tree  $\mathcal{S}$  might make certain valuations in  $\mathcal{T}$  become unsatisfying. This could for instance happen when  $\mathcal{S}$  derives a memory fact that  $\mathcal{T}$  later tests for absence. We formalize this below.

The *alpha* nodes of  $\mathcal{T}$ , denoted  $\alpha^\mathcal{T}$ , are all internal nodes  $x$  of  $\mathcal{T}$  for which  $lit^\mathcal{T}(x)$  is a (positive) output or memory literal.<sup>11</sup> Note,  $root^\mathcal{T} \in \alpha^\mathcal{T}$ . The valuations of these alpha nodes have to be satisfiable to make  $\mathcal{T}$  succeed. For each  $x \in \alpha^\mathcal{T}$ , the *beta* nodes of  $x$ , denoted  $\beta^\mathcal{T}(x)$ , are the child-nodes  $y$  of  $x$  for which  $lit^\mathcal{T}(y)$  is a negative output or memory literal. By definition of derivation tree,  $\beta^\mathcal{T}(x)$  contains only leafs. For each  $x \in \alpha^\mathcal{T}$ , a node  $y \in \beta^\mathcal{T}(x)$  is a potential danger: if the fact in the ground literal  $val^\mathcal{T}(x)(lit^\mathcal{T}(y))$ , henceforth referred to as “beta fact”, is accidentally derived before transition  $\kappa^\mathcal{T}(x)$ , then  $val^\mathcal{T}(x)$  is unsatisfying in transition  $\kappa^\mathcal{T}(x)$  (by inflationarity of  $\Pi$ ). The derivation of beta facts could happen when the message deliveries of  $\mathcal{R}^G$  accidentally trigger some rules of  $\Pi$ .

To represent these unwanted derivations, we consider *truncated derivation trees* that are like normal derivation trees, except that message nodes are also leafs. We only consider truncated derivation trees for deriving output and memory facts. We say that a truncated derivation tree  $\mathcal{S}$  can be *aligned* to  $\mathcal{R}^G$  if there is a scheduling  $\lambda : \text{int}^\mathcal{S} \rightarrow \{1, \dots, n\}$  such that for each  $x \in \text{int}^\mathcal{S}$ , message set  $M_{\lambda(x)}^G$  contains  $body^\mathcal{S}(x)|_{\mathbf{r}_{\text{msg}}}$ , i.e., for each valuation in  $\mathcal{S}$ , the necessary messages occur in some well-chosen transitions. Possibly multiple alignments exist for  $\mathcal{S}$ . For an output or memory fact  $\mathbf{f}$ , we write  $align^G(\mathbf{f})$  to denote the

<sup>11</sup>This literal is always positive because  $x$  is an internal node.



set of all pairs  $(\mathcal{S}, \lambda)$  where  $\mathcal{S}$  is a truncated derivation tree for  $\mathbf{f}$  having alignment  $\lambda$  to  $\mathcal{R}^G$ , and such that no two pairs in  $\text{align}^G(\mathbf{f})$  differ only in the values for representing tree-nodes. This set is finite, as we now argue. First, because  $\Pi$  is recursion-free, there are only a finite number of structurally different (truncated) derivation trees for  $\mathbf{f}$ . Second, only a finite number of valuations can be used in the rules of such trees: because these rules are output or memory rules, by message-boundedness, assigned values must either be in  $\mathbf{f}$  or must occur in a message, and  $\mathcal{R}^G$  contains only a finite number of messages.

Now we specify the formula to express that a derivation tree  $\mathcal{T}$  derives its root fact in  $\mathcal{R}^G$ . To obtain a general construction for later use, we take  $\mathcal{T}$  to be a *truncated* derivation tree for an output or memory relation, that has an alignment  $\kappa$  to  $\mathcal{R}^G$ . Note,  $\alpha^{\mathcal{T}} = \text{int}^{\mathcal{T}}$ . The formula is as follows:

$$\text{succeed}_{G, \mathcal{T}, \kappa} := \text{succeed}_{G, \mathcal{T}, \kappa}^{\text{in}} \wedge \text{succeed}_{G, \mathcal{T}, \kappa}^{\text{deny}}$$

with

$$\begin{aligned} \text{succeed}_{G, \mathcal{T}, \kappa}^{\text{in}} &:= \bigwedge_{x \in \alpha^{\mathcal{T}}} \text{body}^{\mathcal{T}}(x)|_{\Upsilon_{\text{in}}}; \text{ and,} \\ \text{succeed}_{G, \mathcal{T}, \kappa}^{\text{deny}} &:= \bigwedge_{x \in \alpha^{\mathcal{T}}} \bigwedge_{\substack{y \in \beta^{\mathcal{T}}(x), \\ \text{let } \mathbf{f} = \text{fact}^{\mathcal{T}}(y)}} \bigwedge_{\substack{(\mathcal{S}, \lambda) \in \text{align}^G(\mathbf{f}), \\ \lambda(\text{root}^{\mathcal{S}}) < \kappa(x)}}} \neg \text{succeed}_{G, \mathcal{S}, \lambda}. \end{aligned}$$

Intuitively, for each  $x \in \alpha^{\mathcal{T}}$ , we express (i) that the input literals in  $\text{body}^{\mathcal{T}}(x)$  are satisfied; and, (ii) we consider all possible truncated derivation trees for beta facts, and their alignments, and demand that these alignments fail to derive the root (beta) fact. The second requirement is expressed with a recursive construction through negation: intuitively, to protect the alpha facts, we must deny the beta facts, which in turn (recursively) requires letting the alpha facts of trees for these beta facts fail, and so on. This recursion ends because each time we pass a truncated derivation tree to the recursive step, the root of this tree is scheduled strictly closer to the beginning of  $\mathcal{R}^G$ . The final formula  $\text{succeed}_{G, \mathcal{T}, \kappa}$  is quantifier-free, with variables in  $\text{adom}(G)$ .

**Combining everything** Let  $G \subseteq \text{forest}_R$  and  $\mathcal{T} \in G$  be as above. We write  $\mathcal{T}^\downarrow$  to denote the truncated version of  $\mathcal{T}$ , by making the nodes that derive messages into leaf nodes. Note, the canonical scheduling  $\kappa^{\mathcal{T}}$ , when restricted to the internal nodes of  $\mathcal{T}^\downarrow$ , is an alignment of  $\mathcal{T}^\downarrow$  to  $\mathcal{R}^G$ . We can combine our previous formulas to express that the messages of  $\mathcal{R}^G$  can be sent and that  $\mathcal{T}^\downarrow$  successfully derives its root fact when its internal nodes are scheduled by  $\kappa^{\mathcal{T}}$ :

$$\text{derive}_{G, \mathcal{T}} := \exists \bar{z} \left( \text{diffVal}_G \wedge \text{sndMsg}_G \wedge \text{succeed}_{G, \mathcal{T}^\downarrow, \kappa^{\mathcal{T}}} \right),$$

where  $\bar{z}$  is an arbitrary ordering of the values in  $\text{adom}(G)$  that do not occur in the root fact of  $\mathcal{T}$ , and where

$$\text{diffVal}_G = \bigwedge_{\substack{a, b \in \text{adom}(G), \\ a \neq b}} (a \neq b).$$

The subformula  $\text{diffVal}_G$  demands that a valuation is injective, which we need in the correctness proof to convert concrete derivation trees to abstract ones

(i.e., to features of formula  $derive_{G,\mathcal{T}}$ ). By the equivalence of  $\exists\text{FO}$  and  $\text{UCQ}^-$ , we may consider  $derive_{G,\mathcal{T}}$  to be a  $\text{UCQ}^-$ -program, having as free variables the tuple  $\bar{x}$  in the root fact of  $\mathcal{T}$ .<sup>12</sup> We can create such a  $\text{UCQ}^-$ -program for every  $G \subseteq forest_R$  and  $\mathcal{T} \in G$ .

Before we can give the final  $\text{UCQ}^-$ -program  $\Phi$ , we need to consider the following. Although  $derive_{G,\mathcal{T}}$  considers alignments of beta facts, an input for  $\Pi$  possibly has not as many different values as  $adom(G)$ . For this reason, we might overlook some alignments that could occur on a real input. For example, an undesirable beta fact might be derivable by a rule  $S(x, x) \leftarrow A_{\text{msg}}(x, x)$  where  $A_{\text{msg}}^{(2)} \in \Upsilon_{\text{msg}}$ . But because  $G$  contains general trees, in run  $\mathcal{R}^G$  we might deliver only (abstract)  $A_{\text{msg}}$ -facts with two different components, preventing an alignment of this rule. To solve this problem, we consider equivalence relations  $E$  on  $adom(G)$ . Assuming a total order on **dom**, we can replace each value  $a \in adom(G)$  by the smallest value in its equivalence class under  $E$ , giving a set of derivation trees  $E(G)$  with a smaller active domain. Using  $E(G)$  instead of  $G$ , and a tree  $\mathcal{T} \in E(G)$ , the variables in  $\text{UCQ}^-$ -program  $derive_{E(G),\mathcal{T}}$  can represent more specific inputs. We write  $Eq(G)$  to denote all equivalence relations of  $adom(G)$  under which the nonequalities of rules in  $G$  are still satisfied.

Now, we define the final program  $\Phi$  as

$$\Phi := \bigcup_{G \subseteq forest_R} \bigcup_{E \in Eq(G)} \bigcup_{\mathcal{T} \in E(G)} derive_{E(G),\mathcal{T}}.$$

The correctness of  $\Phi$  is shown in Appendix D.

## 8 Discussion and Future Work

We have shown that under five restrictions: recursion-freeness; inflationarity; message-positivity; static message sending; and message-boundedness, one obtains decidability in  $\text{NEXPTIME}$  of inconsistency of relational transducer networks implemented by unions of conjunctive queries with negation (and nonequalities). In fact, the problem turns out to be complete for  $\text{NEXPTIME}$ .

As already mentioned in the Introduction, a topic for further work is to investigate whether decidability can be retained while (slightly) relaxing the restrictions of recursion-freeness, inflationarity, and message-positivity. Also, we have only considered concrete transducer networks, i.e., networks with a particular nodeset. It might be interesting to decide if for a given transducer  $\Pi$ , all transducer networks are consistent where  $\Pi$  is replicated on all nodes [6].

Regarding expressivity, the techniques of the upper bound can transform a given consistent simple transducer network to a query-description in  $\text{UCQ}^-$ . When the techniques of the lower bound are applied to this query-description, we obtain a simple transducer network that does not use memory relations anymore, but still expresses the same query as the original network. This can be considered to be some normal form. It might be interesting to describe the smallest size that the normal form could have in relationship to the original network.

There seem to be several reasonable ways to formalize the intuitive notion of eventual consistency. In contrast to our current formalization, a stronger

<sup>12</sup>A variable may occur multiple times in  $\bar{x}$ .

view of eventual consistency [1, 6] is to require that on every input, all infinite “fair” runs produce the same set of output facts. Again, a number of reasonable fairness conditions could be considered here; a rather standard one would be to require that every node performs a transition infinitely often, and that every sent message is eventually delivered. When a transducer network is consistent in this stronger sense, it is also in the confluence sense of this paper, but the other implication is not obvious. Indeed, our notion of eventual consistency only guarantees that outputs can still be produced when messages are delivered in the “right” way. For example, we might have to deliver two messages simultaneously. But this might never happen in some particular fair run. Clearly, the choice of fairness notion plays an important role. Since eventual consistency is indeed meant to be a very weak guarantee [22], it deserves further research to better understand the relation between consistency and fairness requirements.

There also seems to be a pragmatic lesson: although consistency is an interesting property to guarantee for a network, the cost of automatically deciding it might be too high. Indeed, we have to severely restrict the expressiveness of the language and still the resulting decision problem has high intrinsic complexity. For this reason, other approaches might be more viable, such as providing sufficient syntactic guarantees on consistency without unduly limiting the expressive power (e.g. [4, 17]) and without imposing too much distributed coordination (e.g. [6, 23]).

## References

- [1] S. Abiteboul, M. Bienvenu, A. Galland, et al. A rule-based language for Web data management. In *Proceedings 30th ACM Symposium on Principles of Database Systems*, pages 293–304. ACM Press, 2011.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [3] S. Abiteboul, V. Vianu, et al. Relational transducers for electronic commerce. *Journal of Computer and System Sciences*, 61(2):236–269, 2000.
- [4] P. Alvaro, N. Conway, J. Hellerstein, and W.R. Marczak. Consistency analysis in Bloom: A CALM and collected approach. In *Proceedings 5th Biennial Conference on Innovative Data Systems Research*, pages 249–260. [www.cidrdb.org](http://www.cidrdb.org), 2011.
- [5] P. Alvaro, W. Marczak, et al. Dedalus: Datalog in time and space. Technical Report EECS-2009-173, University of California, Berkeley, 2009.
- [6] T.J. Ameloot, F. Neven, and J. Van den Bussche. Relational transducers for declarative networking. In *Proceedings 30th ACM Symposium on Principles of Database Systems*, pages 283–292. ACM Press, 2011.
- [7] T.J. Ameloot and J. Van den Bussche. Deciding eventual consistency for a simple class of relational transducer networks. In *Proceedings of the 15th International Conference on Database Theory*, pages 86–98. ACM Press, 2012.

- [8] A.K. Chandra and M.Y. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM Journal on Computing*, 14(3):671–677, 1985.
- [9] A. Deutsch. Personal communication. 2011.
- [10] A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic verification of data-centric business processes. In *Proceedings 12th International Conference on Database Theory*, 2009.
- [11] A. Deutsch, L. Sui, and V. Vianu. Specification and verification of data-driven Web applications. *Journal of Computer and System Sciences*, 73(3):442–474, 2007.
- [12] A. Deutsch, L. Sui, V. Vianu, and D. Zhou. Verification of communicating data-driven Web services. In *Proceedings 25th ACM Symposium on Principles of Database Systems*, pages 90–99. ACM Press, 2006.
- [13] S. Grumbach and F. Wang. Netlog, a rule-based language for distributed programming. In M. Carro and R. Peña, editors, *Proceedings 12th International Symposium on Practical Aspects of Declarative Languages*, volume 5937 of *Lecture Notes in Computer Science*, pages 88–103, 2010.
- [14] J.M. Hellerstein. Datalog redux: experience and conjecture. Video available (under the title “The Declarative Imperative”) from <http://db.cs.berkeley.edu/jmh/>, 2010. PODS 2010 keynote.
- [15] J.M. Hellerstein. The declarative imperative: experiences and conjectures in distributed logic. *SIGMOD Record*, 39(1):5–19, 2010.
- [16] B.T. Loo et al. Declarative networking. *Communications of the ACM*, 52(11):87–95, 2009.
- [17] W.R. Marczak, P. Alvaro, N. Conway, J.M. Hellerstein, and D. Maier. Confluence analysis for distributed programs: A model-theoretic approach. In P. Barceló and R. Pichler, editors, *Datalog*, volume 7494 of *Lecture Notes in Computer Science*, pages 135–147. Springer, 2012.
- [18] J.A. Navarro and A. Rybalchenko. Operational semantics for declarative networking. In A. Gill and T. Swift, editors, *Proceedings 11th International Symposium on Practical Aspects of Declarative Languages*, volume 5419 of *Lecture Notes in Computer Science*, pages 76–90, 2009.
- [19] E.L. Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52(4):264–268, 1946.
- [20] M. Sipser. *Introduction to the Theory of Computation, Second Edition, International Edition*. Thomson Course Technology, Boston, Massachusetts, USA, 2006.
- [21] M. Spielmann. Verification of relational transducers for electronic commerce. *Journal of Computer and System Sciences*, 66(1):40–65, 2003.
- [22] W. Vogels. Eventual consistency. *Communications of the ACM*, 52(1):40–44, 2009.

- [23] D. Zinn, T.J. Green, and B. Ludaescher. Win-move is coordination-free. In *Proceedings of the 15th International Conference on Database Theory*, pages 99–113. ACM Press, 2012.

## Appendix

### A Undecidability Results

#### A.1 Proof of Proposition 3.2

Inspired by the work of Deutsch et al. [12, 9], we reduce the finite implication problem for functional and inclusion dependencies to the inconsistency decision problem. Section A.1.1 provides notations for dependencies. Next, Section A.1.2 contains the technical description of the reduction. The correctness is shown in Section A.1.3.

##### A.1.1 Dependencies

We introduce notations for dependencies. Let  $\mathcal{D}$  be a database schema, and let  $R^{(k)} \in \mathcal{D}$ . A *functional dependency*  $\sigma$  over  $R$  is a tuple  $(R, \bar{a}, b)$ , where  $\bar{a}$  is a subsequence of  $[1, \dots, k]$  and  $b \in \{1, \dots, k\}$ . This dependency holds for a database instance  $I$  over  $\mathcal{D}$  if for any pair of facts in  $I$ , if they have the same values on components  $\bar{a}$  then they have the same value on component  $b$ .

Let  $R^{(k)}$  and  $S^{(l)}$  be relations in  $\mathcal{D}$ . An *inclusion dependency*  $\sigma$  from  $R$  to  $S$  is a tuple  $(R, \bar{a}, S, \bar{b})$ , where  $\bar{a}$  and  $\bar{b}$  are subsequences of  $[1, \dots, k]$  and  $[1, \dots, l]$  respectively, and  $\bar{a}$  and  $\bar{b}$  have the same length. Denoting  $\bar{a} = [a_1, \dots, a_m]$  and  $\bar{b} = [b_1, \dots, b_m]$ , this dependency holds for a database instance  $I$  over  $\mathcal{D}$  if

$$\{(u_{a_1}, \dots, u_{a_m}) \mid R(u_1, \dots, u_k) \in I\} \subseteq \{(v_{b_1}, \dots, v_{b_m}) \mid S(v_1, \dots, v_l) \in I\}.$$

##### A.1.2 Transducer Network Construction

Let  $(\mathcal{D}, \Sigma, \sigma)$  be an instance of the finite implication problem. We create a single-node transducer network  $\mathcal{N}$  that is simple except that send rules don't have to be static and such that  $\mathcal{N}$  is inconsistent iff  $(\mathcal{D}, \Sigma, \sigma)$  is not valid.

The syntactical simplifications of Section 4.1 are applied.

Abbreviate  $\Sigma' = \Sigma \cup \{\sigma\}$ . Let  $\Upsilon$  be the transducer schema of  $\Pi$ . We define  $\Upsilon_{\text{in}} = \mathcal{D} \cup \{A^{(1)}\}$  where  $A$  is a new relation name not yet occurring in  $\mathcal{D}$ . Relation  $A$  is used to cause inconsistencies. We define  $\Upsilon_{\text{out}} = \{T^{(1)}\}$ . We introduce the message and memory relations of  $\Upsilon$  while we describe the rules of  $\Pi$  below.

We construct  $\Pi$  to be recursion-free; so  $\mathcal{N}$  is also globally recursion-free. Moreover, the output and memory rules will be message-bounded and all rules are message-positive. We only add rules to insert memory facts, making  $\Pi$  inflationary.

**Send input** First,  $\Pi$  sends all input facts to itself. This helps satisfy the message-boundedness restriction. So, for each relation  $R^{(k)} \in \mathcal{D}$ , we have a rule:

$$R_{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k) \leftarrow R(\mathbf{u}_1, \dots, \mathbf{u}_k).$$

**Projecting** To check violations of  $\Sigma'$ , received input messages are projected onto auxiliary memory relations.

Let  $\tau \in \Sigma'$  be a functional dependency. Denote  $\tau = (R, \bar{a}, b)$ . We add a memory relation  $R_\tau^{(l)}$  where  $l$  is the length of  $\bar{a}$  plus 1 (for  $b$ ). On receipt of an  $R_{\text{msg}}$ -fact, we project components  $\bar{a}$  and  $b$  to  $R_\tau$ , with  $\bar{a}$  placed (in order) before  $b$ . This can be done in a message-bounded manner (details omitted).

Let  $\tau \in \Sigma'$  be an inclusion dependency. Denote  $\tau = (R, \bar{a}, S, \bar{b})$ . We add two memory relations  $R_\tau^{(m)}$  and  $S_\tau^{(m)}$ , where  $m$  is the length of  $\bar{a}$  and  $\bar{b}$ . On receipt of an  $R_{\text{msg}}$ - and  $S_{\text{msg}}$ -fact, we project the components  $\bar{a}$  and  $\bar{b}$  (in order) to the relations  $R_\tau$  and  $S_\tau$  respectively. Again, this can be done in a message-bounded manner.

**Checking** The above auxiliary memory relations depend on message delivery, but we don't know when all input facts have been delivered. For this purpose we introduce a special marker message **datadone**<sup>(0)</sup>. We unconditionally send it in every transition, with the rule

$$\text{datadone}() \leftarrow .$$

On receipt of **datadone**(), we create a snapshot of the input facts. We check dependencies only once in this snapshot, by using the memory relation **checkdone**<sup>(0)</sup>, which is filled by the rule

$$\text{checkdone}() \leftarrow \text{datadone}().$$

To actually check dependencies, we proceed as follows. Let  $\tau \in \Sigma'$  be a functional dependency. Denote  $\tau = (R, \bar{a}, b)$ . We send message **viol** <sub>$\tau$</sub> () if  $\tau$  is violated in the snapshot, where  $k = |\bar{a}|$ :

$$\begin{aligned} \text{viol}_\tau() \leftarrow & R_\tau(\mathbf{a}_1, \dots, \mathbf{a}_k, \mathbf{b}), R_\tau(\mathbf{a}_1, \dots, \mathbf{a}_k, \mathbf{b}'), \mathbf{b} \neq \mathbf{b}', \\ & \text{datadone}(), \neg \text{checkdone}(). \end{aligned}$$

Now, let  $\tau \in \Sigma'$  be an inclusion dependency. Denote  $\tau = (R, \bar{a}, S, \bar{b})$ . We send message **viol** <sub>$\tau$</sub> () if  $\tau$  is violated in the snapshot, where  $m = |\bar{a}| = |\bar{b}|$ :

$$\begin{aligned} \text{viol}_\tau() \leftarrow & R_\tau(\mathbf{a}_1, \dots, \mathbf{a}_m), \neg S_\tau(\mathbf{a}_1, \dots, \mathbf{a}_m), \\ & \text{datadone}(), \neg \text{checkdone}(). \end{aligned}$$

**Inconsistent behavior** We cause inconsistent behavior if  $\sigma$  is violated and  $\Sigma$  is not. First, we (unconditionally) send  $A_{\text{msg}}$ -facts, based on the input  $A$ -facts:

$$A_{\text{msg}}(\mathbf{u}) \leftarrow A(\mathbf{u}).$$

Received  $A_{\text{msg}}$ -facts are copied to output relation  $T$  while new memory relation **blocked**<sup>(0)</sup> is empty:

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}), \neg \text{blocked}().$$

Blocking is triggered by the violation of  $\sigma$ :

$$\text{blocked}() \leftarrow \text{viol}_\sigma().$$

So, if  $\sigma$  is violated, inconsistency can be caused by varying the delivery order of  $A_{\text{msg}}$ -facts and  $\text{viol}_\sigma()$ . But we want to remove the inconsistency if any  $\tau \in \Sigma$  turns out to be violated as well, by adding this output rule:

$$T(u) \leftarrow A_{\text{msg}}(u), \text{repair}().$$

Here,  $\text{repair}^{(0)}$  is a new memory relation that becomes enabled when  $\Sigma$  is violated, denoting  $\Sigma = \{\tau_1, \dots, \tau_n\}$ :

$$\begin{aligned} \text{repair}() &\leftarrow \text{viol}_{\tau_1}(). \\ &\vdots \\ \text{repair}() &\leftarrow \text{viol}_{\tau_n}(). \end{aligned}$$

### A.1.3 Correctness

Let  $(\mathcal{D}, \Sigma, \sigma)$  be as above. Let  $\mathcal{N}$  denote the constructed transducer network.

**First direction** Suppose  $(\mathcal{D}, \Sigma, \sigma)$  is not valid. There is an instance  $I$  over  $\mathcal{D}$  such that  $I \models \Sigma$  and  $I \not\models \sigma$ . We give  $\mathcal{N}$  the input  $J = I \cup \{A(a)\}$  and we obtain inconsistency as follows.

In a first run  $\mathcal{R}_1$ , the message  $A_{\text{msg}}(a)$  is sent during the first transition, and in the second transition we deliver only this message, causing the output fact  $T(a)$  to be derived.

In a second run  $\mathcal{R}_2$ , we do not deliver  $A_{\text{msg}}(a)$ . Instead, in  $\mathcal{R}_2$  we send and deliver all input facts of  $I$ , after which we deliver  $\text{datadone}()$ . Now, message  $\text{viol}_\sigma()$  is sent because  $I \not\models \sigma$ . We deliver this message, causing  $\text{blocked}()$  to be derived. This completes the construction of  $\mathcal{R}_2$ . Run  $\mathcal{R}_2$  produces no output because  $A_{\text{msg}}(a)$  is not delivered. Next, no extension of  $\mathcal{R}_2$  can deliver  $\text{viol}_\tau()$  for some  $\tau \in \Sigma$  because  $I \models \Sigma$ . Hence,  $\text{repair}()$  can not be derived. So,  $\text{blocked}()$  prevents  $T(a)$  from being derived whenever  $A_{\text{msg}}(a)$  would be delivered.

**Second direction** For the other direction, suppose that  $\mathcal{N}$  is inconsistent. There is an input  $J$  for  $\mathcal{N}$ , and two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  of  $\mathcal{N}$  on  $J$ , such that  $\mathcal{R}_1$  derives an output fact  $T(a)$  and  $\mathcal{R}_2$  does not, and neither can  $T(a)$  be derived in any extension of  $\mathcal{R}_2$ . We show there is a subset  $I \subseteq J|_{\mathcal{D}}$  such that  $I \models \Sigma$  and  $I \not\models \sigma$ , so that  $(\mathcal{D}, \Sigma, \sigma)$  is not valid.

First, the derivation of  $T(a)$  in  $\mathcal{R}_1$  implies that  $A_{\text{msg}}(a)$  can be sent in  $\mathcal{R}_1$ . Hence,  $A_{\text{msg}}(a)$  can be sent in  $\mathcal{R}_2$  and in extensions thereof. Therefore, what is preventing  $T(a)$  from being derived in extensions of  $\mathcal{R}_2$  is the presence of  $\text{blocked}()$  and the absence of  $\text{repair}()$ . The fact  $\text{blocked}()$  was derived by the delivery of  $\text{viol}_\sigma()$ . This delivery must have happened inside  $\mathcal{R}_2$  because otherwise in some extension of  $\mathcal{R}_2$  we could postpone the delivery of  $\text{viol}_\sigma()$  until after  $A_{\text{msg}}(a)$  was delivered, deriving  $T(a)$ , which is impossible in any extension of  $\mathcal{R}_2$ .

The sending of  $\text{viol}_\sigma()$  implies that  $\text{datadone}()$  was delivered in some transition  $i$  of  $\mathcal{R}_2$ , and at moment the transducer had received a snapshot  $I \subseteq J|_{\mathcal{D}}$  such that  $I \not\models \sigma$ . Also, because  $\text{repair}()$  was not derived in  $\mathcal{R}_2$  and can not be derived in an extension, it must be that no  $\text{viol}_\tau()$ -fact was ever sent for any  $\tau \in \Sigma$ . So, in transition  $i$  of  $\mathcal{R}_2$ , we have  $I \models \Sigma$ .

## A.2 Proof of Proposition 3.3

Let  $(U, V)$  be an instance of the Post correspondence problem. Denote  $U = u_1, \dots, u_n$  and  $V = v_1, \dots, v_n$ . We construct a single-node transducer network  $\mathcal{N}$  that is simple except that local message recursion is allowed, such that  $(U, V)$  has a match iff  $\mathcal{N}$  is *inconsistent*.

### A.2.1 Notations

For a word  $w$  and an index  $k \in \{1, \dots, |w|\}$ , we write  $w[k]$  to denote the symbol of  $w$  at position  $k$ .

### A.2.2 Transducer Network Construction

We now define the single transducer  $\Pi$  of  $\mathcal{N}$  and its transducer schema  $\Upsilon$ . The syntactical simplifications of Section 4.1 are applied.

**Represent words** For each  $i \in \{1, \dots, n\}$ , we add to  $\Upsilon_{\text{in}}$  unary relations  $U_k^i$  and  $V_l^i$  with  $k \in \{1, \dots, |u_i|\}$  and  $l \in \{1, \dots, |v_i|\}$ . Now, the words  $u_i$  and  $v_i$  can be encoded. To illustrate,  $u_i = aba$  is represented by the facts  $\{U_1^i(a), U_2^i(b), U_3^i(a)\}$ .

To represent a word-structure with arbitrary length, we provide  $\Upsilon_{\text{in}}$  with the input relations  $R^{(2)}$ ,  $L^{(2)}$  and  $F^{(1)}$ . Here,  $L$  and  $F$  respectively stand for “label” and “first”. For instance, the word  $abc$  might be represented as the facts  $\{R(1, 2), R(2, 3), L(1, a), L(2, b), L(3, c), F(1)\}$ . The word  $a$  can be represented by  $\{F(1), L(1, a)\}$ .

We send **error**( ) whenever the previous input relations violate the following natural constraints:

- all relations  $U_k^i$  and  $V_l^j$  contain at most one symbol; for each pair  $u_i$  and  $v_j$ , and each  $k \in \{1, \dots, |u_i|\}$  and  $l \in \{1, \dots, |v_j|\}$ , the relations  $U_k^i$  and  $V_l^j$  contain a different symbol iff  $u_i[k] \neq v_j[l]$ ; similarly for pairs of two  $U$ -words or two  $V$ -words;
- relation  $R$  contains only chains; relation  $F$  designates at most one start element; each element on the chain has at most one label.

We omit the details of the rules to check these constraints.

**Alignment** We search a match for  $(U, V)$  by aligning  $(u_i, v_i)$ -pairs against the input word-structure. Let  $i \in \{1, \dots, n\}$ . To align the single pair  $(u_i, v_i)$ , we use the following binary message relations:

- relations **align** $[i, k, k]$  with  $1 \leq k \leq \min(|u_i|, |v_i|)$  to represent simultaneous alignment, one character at a time;
- relations **align** $[i, k, |v_i|]$  with  $|v_i| + 1 \leq k \leq |u_i|$  to continue aligning  $u_i$  when  $v_i$  has reached its end;
- relations **align** $[i, |u_i|, k]$  with  $|u_i| + 1 \leq k \leq |v_i|$  to continue aligning  $v_i$  when  $u_i$  has reached its end.



Next, we have the *start rule*, to start aligning at the beginning of the word-structure:

$$\text{align}[i, 1, 1](\mathbf{a}, \mathbf{a}) \leftarrow F(\mathbf{a}), L(\mathbf{a}, \mathbf{c}), U_1^i(\mathbf{c}), V_1^i(\mathbf{c}).$$

Then we have *simultaneous continuation rules* for each  $k$  satisfying  $1 \leq k \leq \min(|u_i|, |v_i|) - 1$ :

$$\begin{aligned} \text{align}[i, k+1, k+1](\mathbf{a}', \mathbf{b}') &\leftarrow \text{align}[i, k, k](\mathbf{a}, \mathbf{b}), R(\mathbf{a}, \mathbf{a}'), R(\mathbf{b}, \mathbf{b}'), \\ &\quad L(\mathbf{a}', \mathbf{c}_1), L(\mathbf{b}', \mathbf{c}_2), U_{k+1}^i(\mathbf{c}_1), V_{k+1}^i(\mathbf{c}_2). \end{aligned}$$

We have *separate continuation rules* for  $u_i$ , for each  $k$  satisfying  $|v_i| \leq k \leq |u_i| - 1$ :

$$\text{align}[i, k+1, |v_i|](\mathbf{a}', \mathbf{b}) \leftarrow \text{align}[i, k, |v_i|](\mathbf{a}, \mathbf{b}), R(\mathbf{a}, \mathbf{a}'), L(\mathbf{a}', \mathbf{c}), U_{k+1}^i(\mathbf{c}).$$

Similarly, we have *separate continuation rules* for  $v_i$ , for each  $k$  satisfying  $|u_i| \leq k \leq |v_i| - 1$ :

$$\text{align}[i, |u_i|, k+1](\mathbf{a}, \mathbf{b}') \leftarrow \text{align}[i, |u_i|, k](\mathbf{a}, \mathbf{b}), R(\mathbf{b}, \mathbf{b}'), L(\mathbf{b}', \mathbf{c}), V_{k+1}^i(\mathbf{c}).$$

Lastly, once  $u_i$  and  $v_i$  are both fully aligned, for each pair  $(u_j, v_j)$  with  $j \in \{1, \dots, n\}$  we have the *switch rule* from pair  $i$  to pair  $j$  (with possibly  $i = j$ ):

$$\begin{aligned} \text{align}[j, 1, 1](\mathbf{a}', \mathbf{b}') &\leftarrow \text{align}[i, |u_i|, |v_i|](\mathbf{a}, \mathbf{b}), R(\mathbf{a}, \mathbf{a}'), R(\mathbf{b}, \mathbf{b}'), \\ &\quad L(\mathbf{a}', \mathbf{c}_1), L(\mathbf{b}', \mathbf{c}_2), U_1^j(\mathbf{c}_1), V_1^j(\mathbf{c}_2). \end{aligned}$$

**Inconsistent behavior** Inconsistency is obtained in a similar fashion as in Section A.1. We add input relation  $A^{(1)}$  and message relation  $A_{\text{msg}}^{(1)}$ , and a sending rule:

$$A_{\text{msg}}(\mathbf{u}) \leftarrow A(\mathbf{u}).$$

We also have an output relation  $T^{(1)}$  to which received  $A_{\text{msg}}$ -facts are copied while a memory relation **blocked**() is nonempty:

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}), \neg \text{blocked}().$$

Now, whenever we receive a message of the form  $\text{align}[i, |u_i|, |v_i|](a, a)$ , we have been able to successfully align a sequence of  $(u_i, v_i)$ -pairs to the input word-structure, so that the  $U$ - and  $V$ -side end at the same position. This corresponds to a match for  $(U, V)$ . For each  $i \in \{1, \dots, n\}$ , add the memory insertion rule:

$$\text{blocked}() \leftarrow \text{align}[i, |u_i|, |v_i|](\mathbf{a}, \mathbf{a}).$$

Note, these rules are message-bounded. So, inconsistency is obtained by varying the delivery order of  $A_{\text{msg}}$ -facts and such alignment-messages. Inconsistencies are repaired when **error**() is received (together with  $A_{\text{msg}}$ -facts):

$$T(\mathbf{u}) \leftarrow A_{\text{msg}}(\mathbf{u}), \text{error}().$$

### A.2.3 Correctness

Let  $(U, V)$  be an instance of the Post correspondence problem. Let  $\mathcal{N}$  be the constructed transducer network.

**First direction** Suppose  $(U, V)$  has a match  $E = e_1, \dots, e_m$ . Inconsistency of  $\mathcal{N}$  is obtained as follows. Denote  $w = u_{e_1} \dots u_{e_m}$  (or equivalently  $w = v_{e_1} \dots v_{e_m}$ ). We can naturally encode  $(U, V)$  and  $w$  (as the word-structure) over the input relations. This results in an instance  $J$  on which **error**() can not be sent. We give  $I = J \cup \{A(a)\}$  as input to  $\mathcal{N}$ .

In a first run  $\mathcal{R}_1$  on  $I$ , we immediately send and deliver  $A_{\text{msg}}(a)$ , causing  $T(a)$  to be derived. In a second run  $\mathcal{R}_2$ , we do not deliver  $A_{\text{msg}}(a)$ , but, following sequence  $E$ , we send messages to align pairs of  $(U, V)$  to the encoding of  $w$ . Abbreviating  $z = e_m$ , and assuming the chain in the word-structure consists of consecutive natural numbers starting at 1, at some point we send **align** $[z, |u_z|, |v_z|](|w|, |w|)$ . Upon delivering this message in  $\mathcal{R}_2$ , we derive **blocked**(). Because **error**() can not be sent,  $T(a)$  can not be derived in any extension of  $\mathcal{R}_2$ .

**Second direction** Suppose that  $\mathcal{N}$  is inconsistent. We show that  $(U, V)$  has a match. There is an input  $I$  for  $\mathcal{N}$  and two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  such that  $\mathcal{R}_1$  derives an output fact  $T(a)$  that is not derived in  $\mathcal{R}_2$  or any extensions thereof. The presence of  $T(a)$  in  $\mathcal{R}_1$  implies that  $A_{\text{msg}}(a)$  can be delivered in  $\mathcal{R}_1$ . So,  $A_{\text{msg}}(a)$  can also be delivered in extensions of  $\mathcal{R}_2$ . The reason why  $T(a)$  can not be derived in such extensions is the presence of **blocked**() and because **error**() can never be sent. Fact **blocked**() must have been derived in  $\mathcal{R}_2$  itself, by delivering a message of the form **align** $[i, |u_i|, |v_i|](a, a)$ .<sup>13</sup>

By going over the derivation history of **align** $[i, |u_i|, |v_i|](a, a)$  in a forward manner, we obtain a sequence  $E = e_1, \dots, e_m$  of indices in  $\{1, \dots, n\}$  by looking at the used start- or switch-rules. Sequence  $E$  is a match, because the absence of **error**() implies that the alignment of the  $U$ -words “sees” the same word-structure as the alignment of the  $V$ -words. This would not be the case, for instance, when an element of the word-structure could have two labels or when the other natural constraints on the input are violated.

## B Small Model Property

### B.1 Details of Section 5.3

Let  $\mathcal{R}$  be a run of  $\mathcal{N}$  on input  $I$ . We construct  $\text{hist}_{\mathcal{R}}$  and  $\text{msg}_{\mathcal{R}}$  such that the properties 1, 2, and 3 of Section 5.3 are satisfied. Let  $n$  be the number of transitions of  $\mathcal{R}$ . For each  $i \in \{1, \dots, n+1\}$ , we denote the  $i^{\text{th}}$  configuration of  $\mathcal{R}$  as  $\rho_i = (s_i^{\mathcal{R}}, b_i^{\mathcal{R}})$ . For a transition  $i$ , we denote the multiset of delivered messages and the set of sent messages respectively as  $m_i^{\mathcal{R}}$  and  $\delta_i^{\mathcal{R}}$ .

We will perform the construction backwards, starting in the last transition of  $\mathcal{R}$ . Inductively, for each transition  $j = n, n-1, \dots, 1$ , we define  $\text{hist}_{\mathcal{R}}^j$  and  $\text{msg}_{\mathcal{R}}^j$ , where, intuitively,  $\text{hist}_{\mathcal{R}}^j$  and  $\text{msg}_{\mathcal{R}}^j$  say something about the  $C$ -facts and their needed messages for transition  $j$  and later. In the end, we define  $\text{hist}_{\mathcal{R}} = \text{hist}_{\mathcal{R}}^1$  and  $\text{msg}_{\mathcal{R}} = \text{msg}_{\mathcal{R}}^1$ . For each pair of transitions  $j$  and  $i$ ,  $\text{hist}_{\mathcal{R}}^j$  and  $\text{msg}_{\mathcal{R}}^j$  give rise to the (multi)sets  $\gamma_i^j$ ,  $\beta_i^j$ , and  $\mathcal{E}_i^j$ , defined as in Section 5.3.2. By induction on  $j$ , we want the following properties to be satisfied:

<sup>13</sup>If **blocked**() would not be derived in  $\mathcal{R}_2$  itself, we could simply extend  $\mathcal{R}_2$  by delivering  $A_{\text{msg}}(a)$ , upon which  $T(a)$  would be derived.

1.  $\gamma_i^j \subseteq b_i^{\mathcal{R}}$  for each transition index  $i$ ;
2.  $\beta_i^j$  is a set for each transition index  $i$ ;
3.  $\mathcal{E}_i^j = \gamma_{i+1}^j \cap \delta_i^{\mathcal{R}}$  for each transition index  $i$ ; and,
4.  $hist_{\mathcal{R}}^j$  contains only derivation pairs for transitions  $j$  and later.

To allow for a simple base case, we start the inductive construction at  $j = n + 1$  and we define  $hist_{\mathcal{R}}^{n+1} = \emptyset$  (no mappings) and  $msg_{\mathcal{R}}^{n+1} = \emptyset$ . The induction properties are satisfied for the base case. For the induction hypothesis, we assume that  $hist_{\mathcal{R}}^{j+1}$  and  $msg_{\mathcal{R}}^{j+1}$  are defined such that the properties are satisfied.

### B.1.1 Extend derivation history

We define  $hist_{\mathcal{R}}^j$  to be  $hist_{\mathcal{R}}^{j+1}$  extended with an assignment of a derivation pair  $(\varphi, V)$  to each pair  $(j, \mathbf{g})$  where  $\mathbf{g}$  is either (i) an output or memory  $C$ -fact created during transition  $j$  of  $\mathcal{R}$ , or (ii) a needed message such that  $(j, \mathbf{g}, l) \in msg_{\mathcal{R}}^{j+1}$  for some  $l$ . Note,  $hist_{\mathcal{R}}^j$  is a function because there are no derivation pairs for transition  $j$  in  $hist_{\mathcal{R}}^{j+1}$ .

Now we define  $msg_{\mathcal{R}}^j$  as an extension of  $msg_{\mathcal{R}}^{j+1}$ . Let  $\beta$  be the set of all messages positively needed by the selected derivation pairs in  $hist_{\mathcal{R}}^j$  for transition  $j$ . For each  $\mathbf{g} \in \beta$ , we will select an origin transition  $k$  of  $\mathbf{g}$ , and the resulting triple  $(k, \mathbf{g}, j)$  is added to  $msg_{\mathcal{R}}^j$ . There are two cases:

- If there is no triple  $(k_0, \mathbf{g}, l) \in msg_{\mathcal{R}}^{j+1}$  with  $k_0 < j$  then we define  $k$  to be the largest transition index of  $\mathcal{R}$  for which  $k < j$  and  $\mathbf{g} \in \delta_k^{\mathcal{R}}$ ;
- Otherwise, let  $k_0$  be the smallest transition of  $\mathcal{R}$  for which  $(k_0, \mathbf{g}, l) \in msg_{\mathcal{R}}^{j+1}$  and  $k_0 < j$ . Then we can apply Claim B.1 to know  $num(\mathbf{g}, \gamma_{k_0}^{j+1}) < num(\mathbf{g}, b_{k_0}^{\mathcal{R}})$ . So, intuitively, we have some instance of  $\mathbf{g}$  in  $b_{k_0}^{\mathcal{R}}$  that is not yet used in  $msg_{\mathcal{R}}^{j+1}$ . We now define  $k$  as the largest transition index of  $\mathcal{R}$  for which  $k < k_0$  and  $\mathbf{g} \in \delta_k^{\mathcal{R}}$ .

### B.1.2 Show induction properties

We show that the induction properties are satisfied. First,  $hist_{\mathcal{R}}^j$  by construction only contains derivation pairs for transitions  $j$  and later. Now we show the properties for  $msg_{\mathcal{R}}^j$ . Because we have added triples only for facts in  $\beta$  to  $msg_{\mathcal{R}}^j$  with respect to  $msg_{\mathcal{R}}^{j+1}$ , it is sufficient to focus on one  $\mathbf{g} \in \beta$ . Let  $k$  be the transition index such that  $(k, \mathbf{g}, j) \in msg_{\mathcal{R}}^j$ . Let  $i \in \{1, \dots, n\}$  be an arbitrary transition index. We consider each of the properties:

**Inclusion** We have to show  $num(\mathbf{g}, \gamma_i^j) \leq num(\mathbf{g}, b_i^{\mathcal{R}})$ . If  $i \leq k$  then  $num(\mathbf{g}, \gamma_i^j) = 0$ , because index  $k$  by choice is the smallest transition index of  $\mathcal{R}$  for which  $(k, \mathbf{g}, l) \in msg_{\mathcal{R}}^j$  for some  $l$ . If  $j < i$ , then  $num(\mathbf{g}, \gamma_i^j) = num(\mathbf{g}, \gamma_i^{j+1})$  since  $(k, \mathbf{g}, j)$  is only a delivery for transition  $j$ ; thus the property is satisfied by applying the induction hypothesis.

Lastly, we consider the case  $k < i \leq j$ . If there is no triple  $(k_0, \mathbf{g}, l) \in msg_{\mathcal{R}}^{j+1}$  with  $k_0 < j$  then by choice of  $k$  we have  $num(\mathbf{g}, \gamma_i^j) = 1$ . And because  $\mathbf{g}$  is

not sent between  $k$  and  $j$  and yet  $\text{num}(\mathbf{g}, b_j^{\mathcal{R}}) \geq 1$  (since  $\mathbf{g} \in \beta$ ), it must be  $\text{num}(\mathbf{g}, b_i^{\mathcal{R}}) \geq 1$ ; hence,  $\text{num}(\mathbf{g}, \gamma_i^j) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$ .

Now suppose that  $k_0$  exists. We consider the subcases  $k < i \leq k_0$  and  $k_0 < i \leq j$ . If  $k < i \leq k_0$  then  $\text{num}(\mathbf{g}, \gamma_i^j) = 1$ , and since  $\mathbf{g}$  is not sent between  $k$  and  $k_0$  and yet  $\text{num}(\mathbf{g}, b_{k_0}^{\mathcal{R}}) \geq 1$  (Claim B.1), it must be  $\text{num}(\mathbf{g}, b_i^{\mathcal{R}}) \geq 1$ ; hence,  $\text{num}(\mathbf{g}, \gamma_i^j) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$ . If  $k_0 < i \leq j$ , we have  $\text{num}(\mathbf{g}, \gamma_i^j) = \text{num}(\mathbf{g}, \gamma_i^{j+1}) + 1$  because  $(k, \mathbf{g}, j) \in \text{msg}_{\mathcal{R}}^j$  is new (and  $k < k_0$ ) and  $\text{num}(\mathbf{g}, \gamma_i^{j+1}) < \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$  (Claim B.1); hence,  $\text{num}(\mathbf{g}, \gamma_i^j) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$ .

**Set** We have to show  $\text{num}(\mathbf{g}, \beta_i^j) \leq 1$ . If  $i < j$  then  $\text{num}(\mathbf{g}, \beta_i^j) = 0$  and if  $j < i$  then  $\text{num}(\mathbf{g}, \beta_i^j) = \text{num}(\mathbf{g}, \beta_i^{j+1}) \leq 1$ . If  $i = j$  then the property is satisfied because we have selected only one  $k$  such that  $(k, \mathbf{g}, j) \in \text{msg}_{\mathcal{R}}^j$ .

**Equality** We have to show  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = \text{num}(\mathbf{g}, \gamma_{i+1}^j \cap \delta_i^{\mathcal{R}})$ . Let  $k$  be as defined above. If  $i < k$  then  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = 0$  and  $\text{num}(\mathbf{g}, \gamma_{i+1}^j) = 0$  because  $k$  is the smallest origin transition of  $\mathbf{g}$  registered in  $\text{msg}_{\mathcal{R}}^j$ . If  $j \leq i$  then  $\mathcal{E}_i^j = \mathcal{E}_i^{j+1}$  and  $\gamma_{i+1}^j = \gamma_{i+1}^{j+1}$  because in  $\text{msg}_{\mathcal{R}}^j \setminus \text{msg}_{\mathcal{R}}^{j+1}$  we do not register the sending of messages in  $j$ . Next, we consider the case  $k \leq i < j$ . A first observation is that by choice of  $k$ , we have  $\text{num}(\mathbf{g}, \gamma_{i+1}^j) \geq 1$ . Hence, it suffices to show  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = \text{num}(\mathbf{g}, \delta_i^{\mathcal{R}})$ . If  $i = k$  then both  $\text{num}(\mathbf{g}, \delta_i^{\mathcal{R}}) = 1$  and  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = 1$  hold. Now only the more specific case  $k < i < j$  remains, which we divide in two subcases.

If there is no triple  $(k_0, \mathbf{g}, l) \in \text{msg}_{\mathcal{R}}^{j+1}$  with  $k_0 < j$ , then because  $k < i < j$ , by choice of  $k$ , the message  $\mathbf{g}$  is not sent in transition  $i$ . This gives  $\text{num}(\mathbf{g}, \delta_i^{\mathcal{R}}) = 0$ . Consequently  $\mathbf{g}$  was never registered as being sent from transition  $i$ , giving  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = 0$ , as desired.

Now suppose that  $k_0$  exists. If  $k < i < k_0$  then, again like the previous case, we have  $\text{num}(\mathbf{g}, \delta_i^{\mathcal{R}}) = 0$  and  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = 0$ . Suppose  $k_0 \leq i < j$ . We have  $\text{num}(\mathbf{g}, \gamma_{i+1}^{j+1}) \geq 1$  because  $(k_0, \mathbf{g}, l) \in \text{msg}_{\mathcal{R}}^{j+1}$  for some  $l$  with  $j < l$ . Moreover, since  $\text{num}(\mathbf{g}, \mathcal{E}_i^{j+1}) = \text{num}(\mathbf{g}, \gamma_{i+1}^{j+1} \cap \delta_i^{\mathcal{R}})$  by the induction hypothesis, we obtain  $\text{num}(\mathbf{g}, \mathcal{E}_i^{j+1}) = \text{num}(\mathbf{g}, \delta_i^{\mathcal{R}})$ . Lastly, we have  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = \text{num}(\mathbf{g}, \mathcal{E}_i^{j+1})$  because  $k < i$ . Hence,  $\text{num}(\mathbf{g}, \mathcal{E}_i^j) = \text{num}(\mathbf{g}, \delta_i^{\mathcal{R}})$ .

### B.1.3 Claims

**Claim B.1.** Suppose we are in transition  $j$  of the inductive construction, with  $\text{hist}_{\mathcal{R}}^{j+1}$  and  $\text{msg}_{\mathcal{R}}^{j+1}$  already defined, satisfying the induction properties. Let  $\mathbf{g} \in \beta$ . Suppose there is a transition index  $k_0$  of  $\mathcal{R}$  such that  $(k_0, \mathbf{g}, l) \in \text{msg}_{\mathcal{R}}^{j+1}$  and  $k_0 < j$ . Assume that  $k_0$  is the smallest such index. For each transition  $i \in \{j, j-1, \dots, k_0\}$ , we have  $\text{num}(\mathbf{g}, \gamma_i^{j+1}) < \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$ .

*Proof.* We show this by backward induction on  $i = j, j-1, \dots, k_0$ . To increase readability, we will abbreviate  $j+1$  as the prime symbol  $\iota$ . So,  $\gamma_i^{j+1}$ ,  $\mathcal{E}_i^{j+1}$ , and  $\text{msg}_{\mathcal{R}}^{j+1}$  become respectively  $\gamma_i'$ ,  $\mathcal{E}_i'$ , and  $\text{msg}_{\mathcal{R}}'$ .

**Base case** For the base case,  $i = j$ , we have to show  $\text{num}(\mathbf{g}, \gamma_i') < \text{num}(\mathbf{g}, b_i^{\mathcal{R}})$ . If we can show  $\text{num}(\mathbf{g}, \gamma_i') \leq \text{num}(\mathbf{g}, \gamma_{i+1}' \setminus \delta_i^{\mathcal{R}})$ , then by applying the induction

property  $\gamma'_{i+1} \sqsubseteq b_{i+1}^{\mathcal{R}}$  on  $msg'_{\mathcal{R}}$ , we obtain  $num(\mathbf{g}, \gamma'_i) \leq num(\mathbf{g}, b_{i+1}^{\mathcal{R}} \setminus \delta_i^{\mathcal{R}})$ . And using  $b_{i+1}^{\mathcal{R}} \setminus \delta_i^{\mathcal{R}} = b_i^{\mathcal{R}} \setminus m_i^{\mathcal{R}}$  (by the operational semantics), we get  $num(\mathbf{g}, \gamma'_i) \leq num(\mathbf{g}, b_i^{\mathcal{R}} \setminus m_i^{\mathcal{R}})$ . Lastly, because  $m_i^{\mathcal{R}} \sqsubseteq b_i^{\mathcal{R}}$  and  $num(\mathbf{g}, m_i^{\mathcal{R}}) > 1$  (indeed,  $\mathbf{g} \in \beta \sqsubseteq m_j^{\mathcal{R}} = m_i^{\mathcal{R}}$ ), we obtain  $num(\mathbf{g}, \gamma'_i) < num(\mathbf{g}, b_i^{\mathcal{R}})$ , as desired.

We are left to show  $num(\mathbf{g}, \gamma'_i) \leq num(\mathbf{g}, \gamma'_{i+1} \setminus \delta_i^{\mathcal{R}})$ . Because in  $msg'_{\mathcal{R}}$  no needed messages are registered for transition  $j$  (and smaller), it must be  $num(\mathbf{g}, \gamma'_i) = num(\mathbf{g}, \gamma'_{i+1} \setminus \mathcal{E}'_i)$ . If we can show  $num(\mathbf{g}, \mathcal{E}'_i) = num(\mathbf{g}, \delta_i^{\mathcal{R}})$ , then we are ready. It actually suffices to show  $\mathbf{g} \in \gamma'_{i+1}$ , because then  $num(\mathbf{g}, \mathcal{E}'_i) = num(\mathbf{g}, \delta_i^{\mathcal{R}})$  follows from the induction property  $\mathcal{E}'_i = \gamma'_{i+1} \cap \delta_i^{\mathcal{R}}$  of  $msg'_{\mathcal{R}}$ .

We show  $\mathbf{g} \in \gamma'_{i+1}$ . By definition of  $k_0$ , there is a triple  $(k_0, \mathbf{g}, l) \in msg'_{\mathcal{R}}$  for some  $l$ . Again, because in  $msg'_{\mathcal{R}}$  no needed messages are registered for transition  $j$  and smaller, it must be  $j < l$  or equivalently  $j+1 = i+1 \leq l$ . Hence,  $\mathbf{g} \in \gamma'_{i+1}$  by definition of  $\gamma'_{i+1}$ .

**Inductive step** For the induction hypothesis, suppose that  $num(\mathbf{g}, \gamma'_{i+1}) < num(\mathbf{g}, b_{i+1}^{\mathcal{R}})$ . We show  $num(\mathbf{g}, \gamma'_i) < num(\mathbf{g}, b_i^{\mathcal{R}})$ . We proceed similarly as in the base case, but the strictness “ $<$ ” is obtained differently.

First, by definition of  $k_0$ , we have  $(k_0, \mathbf{g}, l) \in msg_{\mathcal{R}}^{j+1}$  for some  $l$ . Like above, we have  $j < l$ . Hence,  $k_0 \leq i < l$  or equivalently  $k_0 < i+1 \leq l$  and thus  $num(\mathbf{g}, \gamma'_{i+1}) \geq 1$ . Because  $\delta_i^{\mathcal{R}}$  is a set, if we can show  $num(\mathbf{g}, \gamma'_i) \leq num(\mathbf{g}, \gamma'_{i+1} \setminus \delta_i^{\mathcal{R}})$ , then the induction hypothesis gives  $num(\mathbf{g}, \gamma'_i) < num(\mathbf{g}, b_{i+1}^{\mathcal{R}} \setminus \delta_i^{\mathcal{R}})$ . By the operational semantics we would further obtain  $num(\mathbf{g}, \gamma'_i) < num(\mathbf{g}, b_i^{\mathcal{R}} \setminus m_i^{\mathcal{R}}) \leq num(\mathbf{g}, b_i^{\mathcal{R}})$ , as desired.

Showing  $num(\mathbf{g}, \gamma'_i) \leq num(\mathbf{g}, \gamma'_{i+1} \setminus \delta_i^{\mathcal{R}})$  is like in the base case.  $\square$

**Claim B.2.** Let  $\mathcal{R}$  be a run of  $\mathcal{N}$  on  $I$ . Let  $hist_{\mathcal{R}}$  and  $msg_{\mathcal{R}}$  be as defined in Section 5.3. Let  $i$  be a transition index of  $\mathcal{R}$ . We have  $\gamma_{i+1} = (\gamma_i \setminus \beta_i) \cup \mathcal{E}_i$  (multiset difference and union).

*Proof.* Let  $\mathbf{g}$  be a fact. We show  $num(\mathbf{g}, \gamma_{i+1}) = num(\mathbf{g}, (\gamma_i \setminus \beta_i) \cup \mathcal{E}_i)$ .

First,  $num(\mathbf{g}, \gamma_{i+1})$  is, by definition of  $\gamma_{i+1}$ , the number of triples  $(j, \mathbf{g}, k) \in msg_{\mathcal{R}}$  for which  $j < i+1$  and  $i+1 \leq k$ . Hence,  $num(\mathbf{g}, \gamma_{i+1}) = e_1 + e_2$ , where

- $e_1$  is the number of triples  $(j, \mathbf{g}, k) \in msg_{\mathcal{R}}$  for which  $j < i$  and  $i+1 \leq k$ , and,
- $e_2$  is the number of triples  $(j, \mathbf{g}, k) \in msg_{\mathcal{R}}$  for which  $j = i$  and  $i+1 \leq k$ .

Regarding  $e_2$ , since always  $j < k$ , the equality  $j = i$  already implies  $i+1 \leq k$ . So,  $e_2$  simplifies to the number of triples  $(i, \mathbf{g}, k) \in msg_{\mathcal{R}}$ , or equivalently  $e_2 = num(\mathbf{g}, \mathcal{E}_i)$ . If we would know that  $e_1 = num(\mathbf{g}, \gamma_i \setminus \beta_i)$  then overall we would obtain, as desired:

$$\begin{aligned} num(\mathbf{g}, \gamma_{i+1}) &= num(\mathbf{g}, \gamma_i \setminus \beta_i) + num(\mathbf{g}, \mathcal{E}_i) \\ &= num(\mathbf{g}, (\gamma_i \setminus \beta_i) \cup \mathcal{E}_i). \end{aligned}$$

Now we show  $e_1 = num(\mathbf{g}, \gamma_i \setminus \beta_i)$ . Using that  $i+1 \leq k$  is equivalent to  $i < k$ , we have  $e_1 = f_1 - f_2$ , where

- $f_1$  is the number of triples  $(j, \mathbf{g}, k) \in msg_{\mathcal{R}}$  for which  $j < i$  and  $i \leq k$ , and,
- $f_2$  is the number of triples  $(j, \mathbf{g}, k) \in msg_{\mathcal{R}}$  for which  $j < i$  and  $i = k$  (or simply  $i = k$  because always  $j < k$ ).

By definition of  $\gamma_i$  and  $\beta_i$ , we have  $f_1 = \text{num}(\mathbf{g}, \gamma_i)$  and  $f_2 = \text{num}(\mathbf{g}, \beta_i)$ . Lastly, because  $\text{num}(\mathbf{g}, \beta_i) \leq \text{num}(\mathbf{g}, \gamma_i)$ , we obtain

$$\begin{aligned} e_1 &= \text{num}(\mathbf{g}, \gamma_i) - \text{num}(\mathbf{g}, \beta_i) \\ &= \text{num}(\mathbf{g}, \gamma_i \setminus \beta_i). \end{aligned}$$

□

## B.2 Details of Section 5.4

**Claim B.3.** Let the transitions of  $\mathcal{S}$  be defined up to and including transition  $i$ . If  $\gamma_i \sqsubseteq b_i^{\mathcal{S}}$  then  $\beta_i \subseteq \text{set}(m_i^{\mathcal{S}})$ .

*Proof.* By definition,  $m_i^{\mathcal{S}} = (b_i^{\mathcal{S}} \setminus (\gamma_i \setminus \beta_i)) \cap m_i^{\mathcal{R}}$ . Let  $\mathbf{g} \in \beta_i$ . It is sufficient to show that  $\text{num}(\mathbf{g}, b_i^{\mathcal{S}} \setminus (\gamma_i \setminus \beta_i)) \geq 1$  and  $\text{num}(\mathbf{g}, m_i^{\mathcal{R}}) \geq 1$ .

We show that  $\text{num}(\mathbf{g}, b_i^{\mathcal{S}} \setminus (\gamma_i \setminus \beta_i)) \geq 1$ . It is sufficient to show  $\text{num}(\mathbf{g}, b_i^{\mathcal{S}}) \geq 1$  and  $\text{num}(\mathbf{g}, \gamma_i \setminus \beta_i) < \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ . First, because  $\beta_i$  is a set (property of  $\text{msg}_{\mathcal{R}}$ ), and  $\mathbf{g} \in \beta_i$ , we have  $\text{num}(\mathbf{g}, \beta_i) = 1$ . Also, the given assumption  $\gamma_i \sqsubseteq b_i^{\mathcal{S}}$  implies  $\text{num}(\mathbf{g}, \gamma_i) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ .

- We show  $\text{num}(\mathbf{g}, b_i^{\mathcal{S}}) \geq 1$ . From the definition of  $\beta_i$  and  $\gamma_i$ , we have  $\text{num}(\mathbf{g}, \beta_i) \leq \text{num}(\mathbf{g}, \gamma_i)$ . And since  $\text{num}(\mathbf{g}, \beta_i) = 1$  and  $\text{num}(\mathbf{g}, \gamma_i) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ , we obtain  $\text{num}(\mathbf{g}, b_i^{\mathcal{S}}) \geq 1$ .
- We show  $\text{num}(\mathbf{g}, \gamma_i \setminus \beta_i) < \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ . Since  $\text{num}(\mathbf{g}, \beta_i) = 1$  and  $\text{num}(\mathbf{g}, \beta_i) \leq \text{num}(\mathbf{g}, \gamma_i)$ , we have  $\text{num}(\mathbf{g}, \gamma_i \setminus \beta_i) < \text{num}(\mathbf{g}, \gamma_i)$ . Combined with  $\text{num}(\mathbf{g}, \gamma_i) \leq \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ , we obtain  $\text{num}(\mathbf{g}, \gamma_i \setminus \beta_i) < \text{num}(\mathbf{g}, b_i^{\mathcal{S}})$ .

We are left to show that  $\text{num}(\mathbf{g}, m_i^{\mathcal{R}}) \geq 1$ . By definition of  $\mathbf{g} \in \beta_i$ , there is a triple  $(k, \mathbf{g}, l) \in \text{msg}_{\mathcal{R}}$  with  $l = i$ . Hence, by construction of  $\text{msg}_{\mathcal{R}}$ , we have  $\text{num}(\mathbf{g}, m_i^{\mathcal{R}}) \geq 1$ . □

**Claim B.4.** Let  $\mathcal{R}$  be a run of  $\mathcal{N}$  on input  $I$ . Suppose a run  $\mathcal{S}$  of  $\mathcal{N}$  on  $J$  has the properties that (i)  $\text{last}(\mathcal{S})$  and  $\text{last}(\mathcal{R})$  contain the same output and memory  $C$ -facts, and, (ii) the message buffer of  $\text{last}(\mathcal{S})$  is a submultiset of the message buffer in  $\text{last}(\mathcal{R})$ . Then, for every extension  $\mathcal{S}'$  of  $\mathcal{S}$ , there is an extension  $\mathcal{R}'$  of  $\mathcal{R}$  such that  $\text{last}(\mathcal{S}')$  and  $\text{last}(\mathcal{R}')$  again contain precisely the same output and memory  $C$ -facts.

*Proof.* Let  $\mathcal{S}'$  be an extension of  $\mathcal{S}$  that does  $m$  new transitions after those of  $\mathcal{S}$ , with  $m \geq 1$ . The idea is to extend  $\mathcal{R}$  by also doing  $m$  new transitions, in each of which we do the same message deliveries as in the corresponding transition in the extension of  $\mathcal{S}$ . This results in run  $\mathcal{R}'$ .

For each  $i \in \{1, \dots, m+1\}$ , let  $\rho_i = (s_i^{\mathcal{R}}, b_i^{\mathcal{R}})$  and  $\sigma_i = (s_i^{\mathcal{S}}, b_i^{\mathcal{S}})$  denote the  $i^{\text{th}}$  configuration in the extension of respectively  $\mathcal{R}$  and  $\mathcal{S}$ , with  $\rho_1 = \text{last}(\mathcal{R})$  and  $\sigma_1 = \text{last}(\mathcal{S})$ . We show by induction on  $i \in \{1, \dots, m+1\}$  that (i)  $\sigma_i$  and  $\rho_i$  contain the same output and memory  $C$ -facts, and, (ii) the message buffer of  $\sigma_i$  is a submultiset of the message buffer of  $\rho_i$ . This second property helps us deliver the same messages in the extension of  $\mathcal{R}$  as done in the extension of  $\mathcal{S}$ .

For the base case, these properties hold because  $\rho_1 = \text{last}(\mathcal{R})$  and  $\sigma_1 = \text{last}(\mathcal{S})$ . Assuming the properties hold for configuration  $i$  with  $i \geq 1$ , for the inductive step we show that they can be satisfied in configuration  $i+1$ . Recall

that transition  $i$  is responsible for transforming configuration  $i$  into configuration  $i + 1$ . Now, in transition  $i$  of  $\mathcal{R}'$  we deliver the same message multiset as in transition  $i$  of  $\mathcal{S}'$ , which is possible by induction property (ii).

**Output and memory** We show that  $\sigma_{i+1}$  and  $\rho_{i+1}$  have the same output and memory  $C$ -facts. To show that the  $C$ -facts of  $\sigma_{i+1}$  are a subset of those in  $\rho_{i+1}$ , we can apply Claim B.6 (property 1). To show the reverse inclusion, let  $\mathbf{g}$  be a newly derived  $C$ -fact in transition  $i$  of  $\mathcal{R}'$ . We show that  $\mathbf{g}$  is also created in transition  $i$  of  $\mathcal{S}'$ . Let  $(\varphi, V)$  be a derivation pair for  $\mathbf{g}$  in transition  $i$  of  $\mathcal{R}'$ . We show that  $V$  is also satisfying for  $\varphi$  in transition  $i$  of  $\mathcal{S}'$ .

- Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{in}}$ . We have to show  $\mathbf{h} \in J$ . Suppose we would know that  $adom(\mathbf{h}) \subseteq adom(J)$ . Then, since  $\mathbf{h} \in I$  (because  $V$  is satisfying for  $\varphi$  in  $\mathcal{R}'$ ) and  $J = I^{[adom(J)]}$  (Claim B.5), we have  $\mathbf{h} \in J$ , as desired.

Now we show that  $adom(\mathbf{h}) \subseteq adom(J)$ . Let  $\mathbf{a} \in pos^\varphi|_{\Upsilon_{in}}$  be an atom such that  $V(\mathbf{a}) = \mathbf{h}$ . A variable  $u$  in  $\mathbf{a}$  is either free or bound. If  $u$  is free then  $V(u) \in C$  because  $\mathbf{g}$  is a  $C$ -fact, and thus  $V(u) \in adom(J)$  because  $C \subseteq adom(K_1) \subseteq adom(J)$ . Next, if  $u$  is bound then by message-boundedness of  $\varphi$ , value  $V(u)$  occurs in a delivered message during transition  $i$  of  $\mathcal{R}'$ . But this message is also delivered during transition  $i$  of  $\mathcal{S}'$ , and because values in messages of  $\mathcal{S}'$  are restricted to  $adom(J)$ , value  $V(u)$  occurs in  $adom(J)$ .

- Let  $\mathbf{h} \in V(neg^\varphi)|_{\Upsilon_{in}}$ . We have to show  $\mathbf{h} \notin J$ . This follows from  $\mathbf{h} \notin I$  (since  $V$  is satisfying for  $\varphi$  in  $\mathcal{R}'$ ) and  $J \subseteq I$ .
- Recall that  $\varphi$  is message-positive. Because  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}'$ , each message  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{msg}}$  is delivered during that transition. By definition of the message deliveries in  $\mathcal{R}'$ , these messages are also delivered in transition  $i$  of  $\mathcal{S}'$ .
- Let  $\mathbf{h} \in V(pos^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$ . We have to show that  $\mathbf{h}$  is in  $\sigma_i$ . Because  $\mathbf{g}$  is a  $C$ -fact, the message-boundedness of  $\varphi$  implies that  $\mathbf{h}$  is a  $C$ -fact. And because  $V$  is satisfying for  $\varphi$  in  $\mathcal{R}'$ ,  $\mathbf{h}$  is in  $\rho_i$ . By the induction hypothesis,  $\rho_i$  and  $\sigma_i$  have the same output and memory  $C$ -facts. Hence,  $\mathbf{h}$  is in  $\sigma_i$ . Similarly we can show for each  $\mathbf{h} \in V(neg^\varphi)|_{\Upsilon_{out} \cup \Upsilon_{mem}}$  that  $\mathbf{h}$  is not in  $\sigma_i$ .
- Because the nonequalities of  $\varphi$  are satisfied under  $V$  in  $\mathcal{R}'$ , they are also satisfied in  $\mathcal{S}'$ .

We conclude that  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{S}'$ . Hence,  $\mathbf{g} \in \sigma_{i+1}$ .

**Message buffer** We show  $b_{i+1}^S \subseteq b_{i+1}^R$ . Let  $m$  denote the message multiset delivered in transition  $i$ . Let  $\delta_i^R$  and  $\delta_i^S$  denote the message sets sent in new transition  $i$  of  $\mathcal{R}'$  and  $\mathcal{S}'$  respectively. The operational semantics implies that  $b_{i+1}^R = (b_i^R \setminus m) \cup \delta_i^R$  and  $b_{i+1}^S = (b_i^S \setminus m) \cup \delta_i^S$  (multiset difference and union). The desired inclusion  $b_{i+1}^S \subseteq b_{i+1}^R$  follows from  $(b_i^S \setminus m) \subseteq (b_i^R \setminus m)$  (by the induction hypothesis) and  $\delta_i^S \subseteq \delta_i^R$  (by Claim B.6, property 2).  $\square$

**Claim B.5.** The instance  $J$  satisfies  $J = I^{[adom(J)]}$ .

*Proof.* This is because (i)  $J \subseteq I$  implies  $J \subseteq I^{[adom(J)]}$ , and (ii), since  $adom(J) \subseteq adom(K_1) \cup adom(K_2)$ , we have

$$I^{[adom(J)]} \subseteq I^{[adom(K_1) \cup adom(K_2)]} = J.$$

□

**Claim B.6.** Let  $\mathcal{R}$  be a run of  $\mathcal{N}$  on  $I$  and let  $\mathcal{S}$  be a run of  $\mathcal{N}$  on  $J$ . Let  $i$  and  $j$  be a transition index of respectively  $\mathcal{R}$  and  $\mathcal{S}$ . For transition  $i$  of  $\mathcal{R}$ , let  $\rho_i$ ,  $m_i^{\mathcal{R}}$ , and  $\rho_{i+1}$ , respectively denote the begin-configuration, the delivered messages, and the end-configuration. For transition  $j$  of  $\mathcal{S}$  we similarly define  $\sigma_j$ ,  $m_j^{\mathcal{S}}$ , and  $\sigma_{j+1}$ .

Suppose that (i)  $\rho_i$  and  $\sigma_j$  have the same output and memory  $C$ -facts, and, (ii)  $m_j^{\mathcal{S}} \sqsubseteq m_i^{\mathcal{R}}$ . The following properties hold:

1. The output and memory  $C$ -facts of  $\sigma_{j+1}$  are a subset of those in  $\rho_{i+1}$ .
2. The messages sent in transition  $j$  of  $\mathcal{S}$  are a subset of those sent in transition  $i$  of  $\mathcal{R}$ .

*Proof.* The two properties are shown below.

**Property 1** Let  $g$  be an output or memory  $C$ -fact that is newly derived during transition  $j$  of  $\mathcal{S}$ , by means of a derivation pair  $(\varphi, V)$ . We show that  $V$  is also satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}$ .

- Let  $h \in V(pos^\varphi)|_{\mathcal{R}_{in}}$ . We have to show  $h \in I$ . This follows from  $h \in J$  (since  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}$ ) and  $J \subseteq I$  (by construction of  $J$ ).
- Let  $h \in V(neg^\varphi)|_{\mathcal{R}_{in}}$ . We have to show  $h \notin I$ . Since  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}$ , we have  $h \notin J$ . Since  $V$  can only assign values from  $adom(J)$ , we have  $adom(h) \subseteq adom(J)$ . So, if  $h \in I$  then  $h \in I^{[adom(J)]} = J$  (Claim B.5), which is false. Hence,  $h \notin I$ .
- Recall that  $\varphi$  is message-positive. Let  $h \in V(pos^\varphi)|_{\mathcal{R}_{msg}}$ . We have to show that  $h \in m_i^{\mathcal{R}}$ . Because  $V$  is satisfying for  $\varphi$  in  $\mathcal{S}$ , we have  $h \in m_j^{\mathcal{S}} \sqsubseteq m_i^{\mathcal{R}}$ .
- Let  $h \in V(pos^\varphi)|_{\mathcal{R}_{out} \cup \mathcal{R}_{mem}}$ . We have to show that  $h$  is in  $\rho_i$ . Because  $g$  is a  $C$ -fact, the message-boundedness of  $\varphi$  implies that  $h$  is a  $C$ -fact. Moreover, because  $V$  is satisfying for  $\varphi$ , fact  $h$  is a  $C$ -fact in  $\sigma_j$  and thus by assumption also in  $\rho_i$ .

We can similarly show for each  $h \in V(neg^\varphi)|_{\mathcal{R}_{out} \cup \mathcal{R}_{mem}}$  that  $h \notin \rho_i$ .

- Lastly, because the nonequalities of  $\varphi$  are satisfied under  $V$  in  $\mathcal{S}$ , they are also satisfied under  $V$  in  $\mathcal{R}$ .

We obtain that  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}$ . Hence,  $g$  is in  $\rho_{i+1}$ .

**Property 2** Let  $g$  be a message sent in transition  $j$  of  $\mathcal{S}$ , by means of a derivation pair  $(\varphi, V)$ . We show that  $V$  is also satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{R}$ . Because send rules are static, we only have to reason about input and message body atoms of  $\varphi$ . For these body atoms, the proof of property 1 above can actually be applied verbatim to show (i) for each  $h \in V(pos^\varphi)|_{\mathcal{R}_{in}}$  and  $h \in V(neg^\varphi)|_{\mathcal{R}_{in}}$  that respectively  $h \in I$  and  $h \notin I$ ; and (ii) for each  $h \in V(pos^\varphi)|_{\mathcal{R}_{msg}}$  that  $h$  is delivered in transition  $i$  of  $\mathcal{R}$ . □



## C Decidability

### C.1 Details of Section 6.1.2

**Claim C.1.** Let  $f$  be an output fact created in some run of  $\mathcal{N}$  on an input  $I$ . Denote  $C = \text{adom}(f)$ . Let  $\mathcal{R}$  be an arbitrary run of  $\mathcal{N}$  on input  $I$ . There exists a run  $\mathcal{S}$  of  $\mathcal{N}$  on input  $I$  with at most **runLen** transitions and such that  $\text{last}(\mathcal{S})$  contains precisely the same output and memory  $C$ -facts as  $\text{last}(\mathcal{R})$ .

*Proof.* We start by sketching the approach. Like in Section 5.3, we can “mark” the transitions where the output and memory  $C$ -facts are created, and also the transitions where any message is sent that is recursively needed by such a  $C$ -fact. This gives us the function  $\text{hist}_{\mathcal{R}}$  and the set  $\text{msg}_{\mathcal{R}}$  as defined there (satisfying the properties of Section 5.3.2). Since each  $C$ -fact requires at most  $\mathbf{B}^{\mathbf{P}}$  messages by recursion-freeness, at most  $\mathbf{CB}^{\mathbf{P}} + \mathbf{C} = \mathbf{runLen}$  transitions are marked this way. The maximum would be reached if each  $C$ -fact requires a unique set of messages. Let  $\mathcal{M}$  denote the marked transition indices of  $\mathcal{R}$ . Intuitively, the new run  $\mathcal{S}$  does only the marked transitions, so  $|\mathcal{M}|$  in total.

We also need some extra notations. We write  $\rho_i = (s_i^{\mathcal{R}}, b_i^{\mathcal{R}})$  and  $\sigma_i = (s_i^{\mathcal{S}}, b_i^{\mathcal{S}})$  to denote the begin-configuration of transition  $i$  in  $\mathcal{R}$  and  $\mathcal{S}$  respectively. For transition  $i$  of  $\mathcal{R}$ , let  $\gamma_i$  be as defined in Section 5.3.2, based on  $\text{msg}_{\mathcal{R}}$ . Denote  $n = |\mathcal{M}|$ . We can order the transitions of  $\mathcal{M}$  in ascending order, and we write  $\mathcal{M}(i)$  to denote the transition index of  $\mathcal{M}$  at ordinal  $i$  in this ordering, with  $i \in \{1, \dots, n\}$ . For uniformity, we define  $\mathcal{M}(n+1) = n' + 1$ , with  $n'$  the last transition index of  $\mathcal{R}$ .

Now, by induction on the configurations, we construct  $\mathcal{S}$  so that each configuration index  $i \in \{1, \dots, n+1\}$  satisfies the following properties:

- $s_i^{\mathcal{S}}$  contains the same output and memory  $C$ -facts as  $s_{\mathcal{M}(i)}^{\mathcal{R}}$ ; and,
- $\gamma_{\mathcal{M}(i)}$  is a submultiset of  $b_i^{\mathcal{S}}$ .

Then, the last configuration  $s_{n+1}^{\mathcal{S}}$  contains the same output and memory  $C$ -facts as  $s_{\mathcal{M}(n+1)}^{\mathcal{R}} = s_{n'+1}^{\mathcal{R}}$ , which is the last configuration of  $\mathcal{R}$ , as desired. The second induction property helps in showing the first induction property.

For the base case ( $i = 1$ ), we have  $s_1^{\mathcal{S}} = \emptyset$  because  $\sigma_1$  is the start configuration of  $\mathcal{S}$ . Moreover,  $s_{\mathcal{M}(1)}^{\mathcal{R}}$  can not contain any output and memory  $C$ -facts because  $\mathcal{M}(1)$  is the first marked transition, and thus the  $C$ -facts are created *in* or *after* transition  $\mathcal{M}(1)$ . A similar reasoning applies to needed messages:  $\gamma_{\mathcal{M}(1)} = \emptyset$ , which is a submultiset of  $b_1^{\mathcal{S}}$ .

For the induction hypothesis, we assume that the properties hold for configuration  $\sigma_i$  of  $\mathcal{S}$ , with  $i \geq 1$  (and  $i \leq n$ ). Abbreviate  $j = \mathcal{M}(i)$  and let  $\beta_j$  be as in Section 5.3.2. We define transition  $i$  of  $\mathcal{S}$  to deliver *precisely* set  $\beta_j$ . Note, we can deliver  $\beta_j$  because  $\gamma_j \sqsubseteq b_i^{\mathcal{S}}$  (induction hypothesis) and  $\beta_j \sqsubseteq \gamma_j$  (follows from their definition).<sup>14</sup> We now show that the induction properties are satisfied for configuration  $\sigma_{i+1}$ .

**Output and memory** Abbreviate  $k = \mathcal{M}(i+1)$ . We have to show that  $s_{i+1}^{\mathcal{S}}$  and  $s_k^{\mathcal{R}}$  contain the same output and memory  $C$ -facts. We have  $j < k$  (because  $\mathcal{M}(i) < \mathcal{M}(i+1)$ ). Also, there are no other marked transitions between  $j$  and

<sup>14</sup>We deliver no more than  $\beta_j$  to avoid unwanted fact derivations.

$k$ , so no new output and memory  $C$ -facts are created between  $j$  and  $k$ . Finally, inflationarity implies that  $s_{j+1}^{\mathcal{R}}$  and  $s_k^{\mathcal{R}}$  contain precisely the same output and memory  $C$ -facts. Hence, it is sufficient to show that  $s_{i+1}^{\mathcal{S}}$  and  $s_{j+1}^{\mathcal{R}}$  contain the same output and memory  $C$ -facts.

First, let  $\mathbf{g}$  be an output or memory  $C$ -fact in  $s_{i+1}^{\mathcal{S}}$ . We show that  $\mathbf{g} \in s_{j+1}^{\mathcal{R}}$ . If  $\mathbf{g} \in s_i^{\mathcal{S}}$  then by the induction hypothesis  $\mathbf{g} \in s_j^{\mathcal{R}} \subseteq s_{j+1}^{\mathcal{R}}$ . Now suppose  $\mathbf{g} \in s_{i+1}^{\mathcal{S}} \setminus s_i^{\mathcal{S}}$ . Let  $(\varphi, V)$  be a derivation pair for  $\mathbf{g}$  in transition  $i$  of  $\mathcal{S}$ . We show that  $V$  is also satisfying for  $\varphi$  in transition  $j$  of  $\mathcal{R}$ .

- Since  $\mathcal{S}$  and  $\mathcal{R}$  are given the same input, the input literals in the body of  $\varphi$  are satisfied under  $V$  in transition  $j$  of  $\mathcal{R}$  as well.
- Let  $\mathbf{h} \in V(\text{pos}^\varphi)|_{\Upsilon_{\text{msg}}}$ . Since  $V$  is satisfying for  $\varphi$  in transition  $i$  of  $\mathcal{S}$ , it must be  $\mathbf{h} \in \beta_j$ . By construction of  $\text{msg}_{\mathcal{R}}$ , the set  $\beta_j$  is delivered in transition  $j$  of  $\mathcal{R}$ , as desired.
- Since  $s_i^{\mathcal{S}}$  and  $s_j^{\mathcal{R}}$  contain the same output and memory  $C$ -facts (induction hypothesis), message-boundedness of  $\varphi$  implies that the output and memory literals of  $\varphi$  are satisfied under  $V$  in transition  $j$  of  $\mathcal{R}$ .
- Finally, the nonequalities of  $\varphi$  under  $V$  are also satisfied in transition  $j$  of  $\mathcal{R}$  because they are satisfied in transition  $i$  of  $\mathcal{S}$ .

Let  $\mathbf{g}$  be an output or memory  $C$ -fact in  $s_{j+1}^{\mathcal{R}}$ . Similarly to the above, if  $\mathbf{g} \in s_j^{\mathcal{R}}$  then by the induction hypothesis  $\mathbf{g} \in s_i^{\mathcal{S}} \subseteq s_{i+1}^{\mathcal{S}}$ . Because  $\mathbf{g}$  is an output or memory  $C$ -fact, the mapping  $\text{hist}_{\mathcal{R}}(j, \mathbf{g}) = (\varphi, V)$  is defined. We show that  $V$  is also satisfying for  $\varphi$  in transition  $i$  of  $\mathcal{S}$ . The reasoning for nonequalities and input, output, and memory literals of  $\varphi$  is the same as above for the case  $\mathbf{g} \in s_{i+1}^{\mathcal{S}} \setminus s_i^{\mathcal{S}}$ . Let  $\mathbf{h} \in V(\text{pos}^\varphi)|_{\Upsilon_{\text{msg}}}$ . Then  $\mathbf{h}$  is a message needed by  $(\varphi, V)$ , and thus  $\mathbf{g} \in \beta_j$  by construction of  $\text{msg}_{\mathcal{R}}$ . Hence,  $\mathbf{h}$  is delivered in transition  $i$  of  $\mathcal{S}$ .

**Buffer** We have to show  $\gamma_{\mathcal{M}(i+1)} \sqsubseteq b_{i+1}^{\mathcal{S}}$ . Abbreviate  $j = \mathcal{M}(i)$  and  $k = \mathcal{M}(i+1)$ . We have  $j+1 \leq k$  because  $j < k$ . We start by showing  $\gamma_{j+1} = \gamma_k$ , so it becomes sufficient to show  $\gamma_{j+1} \sqsubseteq b_{i+1}^{\mathcal{S}}$ .

Let  $\mathbf{g}$  be a fact. We show  $\text{num}(\mathbf{g}, \gamma_{j+1}) \leq \text{num}(\mathbf{g}, \gamma_k)$ . By definition of  $\gamma_{j+1}$ , expression  $\text{num}(\mathbf{g}, \gamma_{j+1})$  is the number of triples  $(a, \mathbf{g}, b) \in \text{msg}_{\mathcal{R}}$  for which  $a < j+1 \leq b$ . Let  $(a, \mathbf{g}, b)$  be such a triple. It is sufficient to show that  $a < k \leq b$ . We have  $a < k$  because  $a < j+1$  and  $j+1 \leq k$ . Secondly, if  $b < k$  then a needed message is delivered at transition  $b$  of  $\mathcal{R}$ , implying  $b \in \mathcal{M}$ , which is impossible because  $j < b < k$  and there are no marked transitions between  $j$  and  $k$ . Hence,  $k \leq b$ .

Let  $\mathbf{g}$  be a fact. We show  $\text{num}(\mathbf{g}, \gamma_k) \leq \text{num}(\mathbf{g}, \gamma_{j+1})$ . This is similar to the previous direction, but there are also some differences. By definition of  $\gamma_k$ , expression  $\text{num}(\mathbf{g}, \gamma_k)$  is the number of triples  $(a, \mathbf{g}, b) \in \text{msg}_{\mathcal{R}}$  for which  $a < k \leq b$ . Let  $(a, \mathbf{g}, b)$  be such a triple. It is sufficient to show that  $a < j+1 \leq b$ . We have  $j+1 \leq b$  because  $j+1 \leq k$  and  $k \leq b$ . Secondly, if  $j+1 \leq a$  then a needed message would be sent at transition  $a$  of  $\mathcal{R}$ , implying  $a \in \mathcal{M}$ , which is impossible because  $j < a < k$  and there are no marked transitions between  $j$  and  $k$ . Hence,  $a < j+1$ .

Lastly, we show that  $\gamma_{j+1} \sqsubseteq b_{i+1}^S$ . Using Claim B.2, we have  $\gamma_{j+1} = (\gamma_j \setminus \beta_j) \cup \mathcal{E}_j$ . Let  $\delta_i^S$  denote the set of messages sent during transition  $i$  of  $\mathcal{S}$ . The operational semantics implies  $b_{i+1}^S = (b_i^S \setminus \beta_j) \cup \delta_i^S$ . It is sufficient to show  $\gamma_j \setminus \beta_j \sqsubseteq b_i^S \setminus \beta_j$  and  $\mathcal{E}_j \subseteq \delta_i^S$ . The first inclusion follows from the induction hypothesis  $\gamma_j \sqsubseteq b_i^S$ . Now, let  $\mathbf{g} \in \mathcal{E}_j$ . We show  $\mathbf{g} \in \delta_i^S$ . By definition of  $\mathcal{E}_j$ , there is a triple  $(j, \mathbf{g}, b) \in \text{msg}_{\mathcal{R}}$ . So,  $\mathbf{g}$  is a needed message that should be sent in transition  $j$  of  $\mathcal{R}$ . Hence,  $\text{hist}_{\mathcal{R}}(j, \mathbf{g}) = (\varphi, V)$  is defined. We show that  $V$  is satisfying for  $\varphi$  during transition  $i$  of  $\mathcal{S}$ , so that  $\mathbf{g} \in \delta_i^S$ . Because  $\varphi$  is static, we only consider the input and message literals, where the latter are positive by message-positivity. The input literals of  $\varphi$  are satisfied under  $V$  in transition  $i$  of  $\mathcal{S}$ , because they are satisfied in transition  $j$  of  $\mathcal{R}$  and because both runs have the same input. Now, let  $\mathbf{h} \in V(\text{pos}^\varphi)|_{\Upsilon_{\text{msg}}}$ . We have to show that  $\mathbf{h}$  is delivered in transition  $i$  of  $\mathcal{S}$ . Because  $\mathbf{h}$  is delivered in transition  $j$  of  $\mathcal{R}$  (since  $V$  is satisfying for  $\varphi$ ),  $\mathbf{h}$  is a needed message for transition  $j$ ; hence,  $\mathbf{h} \in \beta_j$  and this set is delivered in transition  $i$  of  $\mathcal{S}$ .  $\square$

**Claim C.2.** Let  $I$  be an input for  $\mathcal{N}$ . Let  $\mathcal{R}$  be a run of  $\mathcal{N}$  on  $I$ . Let  $\mathcal{R}'$  be  $\mathcal{R}$  extended by doing  $\mathbf{P} + 1$  additional transitions in each of which we deliver the entire message buffer. Let  $\mathbf{g}$  be a message that is sent in some run  $\mathcal{S}$  of  $\mathcal{N}$  on  $I$ . Message  $\mathbf{g}$  is delivered in the last transition of  $\mathcal{R}'$ .

*Proof.* Recall the definitions and notations regarding derivation trees from Section 2.6. Let  $\mathcal{T}$  be a derivation tree for  $\mathbf{g}$  extracted from  $\mathcal{S}$ . Let  $\kappa^{\mathcal{T}}$  be the canonical scheduling of  $\mathcal{T}$ . Let  $n$  denote the height of  $\mathcal{T}$ , measured as the number of edges on the longest path from the root to a leaf. For  $i \in \{1, \dots, n\}$ , define the following message set  $M_i$ :

$$M_i = \bigcup_{\substack{x \in \text{int}^{\mathcal{T}}, \\ \kappa^{\mathcal{T}}(x) = i}} \text{body}^{\mathcal{T}}(x)|_{\Upsilon_{\text{msg}}}.$$

Because the rules of  $\Pi$  are message-positive,  $\text{body}^{\mathcal{T}}(x)|_{\Upsilon_{\text{msg}}}$  contains only facts. Intuitively,  $M_i$  is the union of all message facts needed by rules scheduled at transition  $i$  by  $\kappa^{\mathcal{T}}$ . Since  $n \leq \mathbf{P}$ , we can consider the transition index  $j$  of  $\mathcal{R}'$  such that  $j+1, \dots, j+n, j+n+1$  are the last  $n+1$  transitions of  $\mathcal{R}'$ . If we can show that  $\mathbf{g}$  is sent in transition  $j+n$ , then  $\mathbf{g}$  is delivered in the last transition  $j+n+1$  (because the entire buffer is delivered), as desired.

Because sending rules are static and message-positive, and  $\mathcal{R}'$  and  $\mathcal{S}$  have the same input  $I$ , it is sufficient to show that  $M_n$  is delivered in transition  $j+n$ , so that the root rule and valuation of  $\mathcal{T}$  derive  $\mathbf{g}$ . Specifically, we show by induction on  $i \in \{1, \dots, n\}$  that  $M_i$  is delivered in transition  $j+i$  of  $\mathcal{R}'$ . The property holds for the base case because  $M_1 = \emptyset$ .<sup>15</sup> For the induction hypothesis, we assume that  $M_i$  can be delivered in transition  $j+i$  of  $\mathcal{R}'$ . We now show that  $M_{i+1}$  can be delivered in transition  $j+i+1$  of  $\mathcal{R}'$ . Let  $\mathbf{h} \in M_{i+1}$ . By definition of  $M_{i+1}$ , there is an internal node  $x$  of  $\mathcal{T}$  with  $\kappa^{\mathcal{T}}(x) = i+1$  and  $\mathbf{h} \in \text{body}^{\mathcal{T}}(x)|_{\Upsilon_{\text{msg}}}$ . We show that  $\mathbf{h}$  is sent in transition  $j+i$  of  $\mathcal{R}'$ , so that  $\mathbf{h}$  is delivered in transition  $j+i+1$ . By message-positivity of  $\text{rule}^{\mathcal{T}}(x)$ , there is a child node  $y \in \text{int}^{\mathcal{T}}$  of  $x$  such that  $\text{fact}^{\mathcal{T}}(y) = \mathbf{h}$ . By definition of  $\kappa^{\mathcal{T}}$ , we have

<sup>15</sup>Indeed, if an internal node  $x$  needs child messages then the corresponding child nodes are scheduled earlier, making  $\kappa^{\mathcal{T}}(x) > 1$ .

Relation	Purpose
$s^{(1)}$ with $s \in \Gamma$	one relation for each tape symbol
$q^{(1)}$ with $q \in Q$	one relation for each state symbol
$0^{(1)}, 1^{(1)}, 01^{(1)}$	relations providing the numbers 0 and 1

Table 2: Input relations for  $M$

$\kappa^{\mathcal{T}}(y) = i$ . We show that  $val^{\mathcal{T}}(y)$  is satisfying for  $rule^{\mathcal{T}}(y)$  during transition  $j + i$  of  $\mathcal{R}'$ . Like above, because sending rules are static and message-positive, and  $\mathcal{R}'$  and  $\mathcal{S}$  have the same input  $I$ , it is sufficient to show that  $M_i$  is delivered in transition  $j + i$ , which holds by the induction hypothesis.  $\square$

## C.2 Complexity Lower Bound

Here we complete the specification of transducer  $\Pi$  over schema  $\Upsilon$  from Section 6.2. We assume that  $\Upsilon_{\text{in}}$  contains the additional relations of Table 2. All rules we specify below are *sending* rules.

Let  $w$  denote the input word for  $M$  under consideration, and let  $n = |w|$ . We can select a constant  $k \in \mathbb{N}$  such that if  $M$  accepts  $w$  then  $M$  has an accepting computation trace on  $w$  with at most  $2^{n^k}$  transitions.

### C.2.1 Binary addresses

Abbreviate  $z = n^k$ . Note,  $z$  is polynomial in  $n$ . Because we are only concerned with accepting computation traces of length at most  $2^{n^k}$ , the address of a reachable tape cell can be represented as a binary number with  $z$  bits. We denote such a number as  $(a_1 \dots a_z)$  where each  $a_i$  is 0 or 1 and  $a_z$  is the least significant bit. Note,  $z$  bits actually allow us to represent addresses larger than  $2^{n^k}$ , but the accepting computation trace will never reach these tape cells, hence, we will ignore those addresses in the following.

We will use messages of the form  $\text{succ}(a_1, \dots, a_z; b_1, \dots, b_z)$  to say that address  $(b_1 \dots b_z)$  is the successor of address  $(a_1 \dots a_z)$ , i.e.,  $(b_1 \dots b_z)$  is obtained from  $(a_1 \dots a_z)$  by adding 1.<sup>16</sup> Similarly, we use messages of the form  $\text{less}(a_1, \dots, a_z; b_1, \dots, b_z)$  and  $\text{diff}(a_1, \dots, a_z; b_1, \dots, b_z)$  to say respectively that  $(a_1 \dots a_z)$  is smaller than  $(b_1 \dots b_z)$  and that  $(a_1 \dots a_z)$  and  $(b_1 \dots b_z)$  are different. To specify these messages, we add the following rules for each  $p = 1, \dots, z$ :

<sup>16</sup>The semicolon in the fact only serves to better separate the two binary numbers visually.

$$\begin{aligned}
\text{succ}(a_1, \dots, a_{p-1}, a_p, \dots, a_z; a_1, \dots, a_{p-1}, b_p, \dots, b_z) &\leftarrow \\
&01(a_1), \dots, 01(a_{p-1}), 0(a_p), 1(b_p), \\
&1(a_{p+1}), \dots, 1(a_z), 0(b_{p+1}), \dots, 0(b_z). \\
\\
\text{less}(a_1, \dots, a_{p-1}, a_p, \dots, a_z; a_1, \dots, a_{p-1}, b_p, \dots, b_z) &\leftarrow \\
&01(a_1), \dots, 01(a_{p-1}), 0(a_p), 1(b_p), \\
&01(a_{p+1}), \dots, 01(a_z), 01(b_{p+1}), \dots, 01(b_z). \\
\\
\text{diff}(a_1, \dots, a_{p-1}, a_p, \dots, a_z; b_1, \dots, b_{p-1}, b_p, \dots, b_z) &\leftarrow \\
&01(a_1), \dots, 01(a_z), 01(b_1), \dots, 01(b_z), a_p \neq b_p.
\end{aligned}$$

Here, if  $p = 1$  then the variables  $a_1$  to  $a_{p-1}$  are nonexistent, and if  $p = z$  then the variables  $a_{p+1}$  to  $a_z$  and  $b_{p+1}$  to  $b_z$  are nonexistent. Note, the number and size of these above rules is polynomial in  $n$ , and they have no cyclic dependencies (leads to recursion-freeness).

### C.2.2 Sending error

The message **error** is sent when some crucial properties of the input relations are violated.

First, we demand that for each configuration at most one state and head position is specified, and also that each tape cell has at most one symbol:

$$\begin{aligned}
\text{error}() &\leftarrow \text{state}(i, q_1), \text{state}(i, q_2), q_1 \neq q_2. \\
&\leftarrow \text{head}(i, h_1, \dots, h_z), \text{head}(i, k_1, \dots, k_z), \\
&\quad \text{diff}(h_1, \dots, h_z; k_1, \dots, k_z). \\
&\leftarrow \text{tape}(i, a_1, \dots, a_z, s_1), \text{tape}(i, a_1, \dots, a_z, s_2), \\
&\quad s_1 \neq s_2.
\end{aligned}$$

For the relations providing the binary numbers, we demand that relations 0 and 1 are disjoint, contain at most one value, and that relation 01 is the union of 0 and 1:

$$\begin{aligned}
\text{error}() &\leftarrow 0(v), 1(v). \\
&\leftarrow 0(v), 0(w), v \neq w. \\
&\leftarrow 1(v), 1(w), v \neq w. \\
&\leftarrow 0(v), \neg 01(v). \\
&\leftarrow 1(v), \neg 01(v). \\
&\leftarrow 01(v), \neg 0(v), \neg 1(v).
\end{aligned}$$

For the relations providing symbols of  $\Gamma$ , we demand that they are pairwise disjoint and that each contains at most one symbol. We demand the same properties of the relations providing symbols of  $Q$ . Formally, for each  $(s_1, s_2) \in (\Gamma \times \Gamma) \cup (Q \times Q)$  with  $s_1 \neq s_2$ , we add the rule

$$\text{error}() \leftarrow s_1(v), s_2(v).$$

And for each  $s \in \Gamma \cup Q$ , we add the rule

$$\text{error}() \leftarrow s(v), s(w), v \neq w.$$

### C.2.3 Sending accept

We give the rules to send messages of the form  $\mathbf{reach}_0(i, j)$  and  $\mathbf{start}(i)$ , where  $\mathbf{reach}_0(i, j)$  indicates that configuration  $j$  can be reached by a valid Turing machine transition from configuration  $i$ , and where  $\mathbf{start}(i)$  indicates that configuration  $i$  has the properties of the start configuration.

**Sending  $\mathbf{reach}_0$**  We will send messages of the form  $\mathbf{tapeCOK}(i, j, a_1, \dots, a_z)$  to say that in configuration  $j$ , the tape cell at address  $(a_1 \dots a_z)$  can be explained by a Turing machine transition applied to configuration  $i$ .<sup>17</sup> To send  $\mathbf{reach}_0(i, j)$ , we have to check that such messages can be sent for *all* tape cells. We will simultaneously enforce that the state and head position of  $j$  can follow from the state and head position of  $i$ .

To send  $\mathbf{tapeCOK}(i, j, a_1, \dots, a_z)$ , we consider three cases, where  $(h_1 \dots h_z)$  denotes the head position of configuration  $i$ :

- $(a_1 \dots a_z) < (h_1 \dots h_z)$ , in which case the cell contents at  $(a_1 \dots a_z)$  should be unaltered in  $j$  with respect to  $i$ ;
- the symmetric case  $(h_1 \dots h_z) < (a_1 \dots a_z)$ , with the same constraint;
- $(a_1 \dots a_z) = (h_1 \dots h_z)$ , in which case a transition of Turing machine  $M$  has to explain the symbol at cell  $(a_1 \dots a_z)$  in  $j$ .

The first case is implemented by the following rule:

$$\mathbf{tapeCOK}(i, j, a_1, \dots, a_z) \leftarrow \mathbf{head}(i, h_1, \dots, h_z), \mathbf{less}(a_1, \dots, a_z; h_1, \dots, h_z), \\ \mathbf{tape}(i, a_1, \dots, a_z, s), \mathbf{tape}(j, a_1, \dots, a_z, s).$$

The second case is done with a similar rule, except that  $\mathbf{less}(a_1, \dots, a_z; h_1, \dots, h_z)$  is replaced by  $\mathbf{less}(h_1, \dots, h_z; a_1, \dots, a_z)$ .

The third case is split further depending on whether the head moves left or right. Let  $\delta$  denote the transition function of Turing machine  $M$ . For each mapping  $(q_1, s_1 \mapsto q_2, s_2, L) \in \delta$ , add the rule:

$$\mathbf{tapeCOK}(i, j, h_1, \dots, h_z) \leftarrow \mathbf{head}(i, h_1, \dots, h_z), \mathbf{head}(j, k_1, \dots, k_z), \\ \mathbf{succ}(k_1, \dots, k_z; h_1, \dots, h_z), \\ \mathbf{state}(i, q_1), \mathbf{tape}(i, h_1, \dots, h_z, s_1), \\ \mathbf{state}(j, q_2), \mathbf{tape}(j, h_1, \dots, h_z, s_2), \\ q_1(q_1), s_1(s_1), q_2(q_2), s_2(s_2).$$

Regarding relations  $q_1, s_1, q_2$  and  $s_2$ , it does not matter what precise values they contain by genericity of the rules (as long as the conditions enforced in Section C.2.2 hold). A similar rule is added for each mapping  $(q_1, s_1 \mapsto q_2, s_2, R) \in \delta$ , except that  $\mathbf{succ}(k_1, \dots, k_z; h_1, \dots, h_z)$  is replaced by  $\mathbf{succ}(h_1, \dots, h_z; k_1, \dots, k_z)$ . Note, the nondeterminism of Turing machine  $M$  is implemented by having multiple rules in  $\Pi$  of these last two forms. Also, the number of rules for relation  $\mathbf{tapeCOK}$  is *constant* because  $M$  is fixed, but their size is polynomial in  $n$ .

<sup>17</sup>The name  $\mathbf{tapeCOK}$  stands for “tape cell ok”.

Next, we send messages of the form  $\mathbf{tapeOK}_m(i, j, a_1, \dots, a_z; b_1, \dots, b_z)$ , with  $m = 0, \dots, z$  and  $(a_1 \dots a_z) \leq (b_1 \dots b_z)$ , to say that interval  $[(a_1 \dots a_z), (b_1 \dots b_z)]$  contains  $2^m$  tape cells and that the message  $\mathbf{tapeCOK}(i, j, c_1, \dots, c_z)$  can be sent for *all* addresses  $(c_1 \dots c_z)$  in this interval. The goal is to eventually send a message  $\mathbf{tapeOK}_z(i, j, a_1, \dots, a_z; b_1, \dots, b_z)$  where  $(a_1 \dots a_z)$  is the first tape cell. To start, we generate  $\mathbf{tapeOK}_0$ -messages:

$$\mathbf{tapeOK}_0(i, j, a_1, \dots, a_z; a_1, \dots, a_z) \leftarrow \mathbf{tapeCOK}(i, j, a_1, \dots, a_z).$$

And we add the following rule for each  $m = 1, \dots, z$ :

$$\begin{aligned} \mathbf{tapeOK}_m(i, j, a_1, \dots, a_z; b_1, \dots, b_z) \leftarrow \\ \mathbf{tapeOK}_{m-1}(i, j, a_1, \dots, a_z; c_1, \dots, c_z), \\ \mathbf{tapeOK}_{m-1}(i, j, d_1, \dots, d_z; b_1, \dots, b_z), \\ \mathbf{succ}(c_1, \dots, c_z; d_1, \dots, d_z). \end{aligned}$$

Note, the number and size of such rules is polynomial in  $n$ .

Finally, the  $\mathbf{reach}_0$ -messages are sent with the following rule:

$$\begin{aligned} \mathbf{reach}_0(i, j) \leftarrow \mathbf{tapeOK}_z(i, j, a_1, \dots, a_z; b_1, \dots, b_z), \\ 0(a_1), \dots, 0(a_z). \end{aligned}$$

Note, we constrain attention to the range  $[0, 2^z]$ .

**Sending start** To send a message  $\mathbf{start}(i)$ , we have to check that configuration  $i$  has the properties of the start configuration: (i) the tape contains the input word  $w$  starting at the first tape cell, with the other tape cells blank; (ii) the state is  $q_0$ ; and, (iii) the head is at tape cell 0. The last two properties are easily checked.

To check property (i), we send messages of the form  $\mathbf{startTapeCOK}(i, a_1, \dots, a_z)$  to indicate that the contents of tape cell  $(a_1 \dots a_z)$  in configuration  $i$  is as required by the start configuration. We add the following rule for all addresses  $a \in [0, n-1]$ , where  $(a_1 \dots a_z)$  is the binary representation of  $a$  and  $w_a$  is the symbol of word  $w$  at (zero-based) index  $a$ :

$$\begin{aligned} \mathbf{startTapeCOK}(i, a_1, \dots, a_z) \leftarrow \\ a_1(a_1), \dots, a_z(a_z), \mathbf{tape}(i, a_1, \dots, a_z, s), w_a(s). \end{aligned}$$

We also add one rule to demand that the other tape cells contain blanks, where  $\sqcup \in \Gamma$  denotes the blank symbol and  $(b_1 \dots b_z)$  is the binary representation of  $n-1$ :

$$\begin{aligned} \mathbf{startTapeCOK}(i, a_1, \dots, a_z) \leftarrow \\ b_1(b_1), \dots, b_z(b_z), \mathbf{less}(b_1, \dots, b_z; a_1, \dots, a_z), \\ \mathbf{tape}(i, a_1, \dots, a_z, s), \sqcup(s). \end{aligned}$$

Note, the number and size of rules for relation  $\mathbf{startTapeCOK}$  is polynomial in  $n$ .

Next, similarly to the relations  $\mathbf{tapeOK}_m$  above, we send messages of the form  $\mathbf{startTapeOK}_m(i, a_1, \dots, a_z; b_1, \dots, b_z)$ , with  $m = 0, \dots, z$  and  $(a_1 \dots a_z) \leq (b_1 \dots b_z)$ , to say that the interval  $[(a_1 \dots a_z), (b_1 \dots b_z)]$  contains  $2^m$  tape cells and that message  $\mathbf{startTapeCOK}(i, c_1, \dots, c_z)$  can be sent for all addresses  $(c_1 \dots c_z)$  in this interval. We do not explicitly give the rules, because they are very similar to the rules of the relations  $\mathbf{tapeOK}_m$ . The number and size of the added rules is also polynomial in  $n$ .

Finally, we can send the **start**-messages:

$$\begin{aligned} \mathbf{start}(i) \leftarrow & \mathbf{startTapeOK}_z(i, \mathbf{a}_1, \dots, \mathbf{a}_z; \mathbf{b}_1, \dots, \mathbf{b}_z), \\ & 0(\mathbf{a}_1), \dots, 0(\mathbf{a}_z), \mathbf{head}(i, \mathbf{a}_1, \dots, \mathbf{a}_z), \\ & \mathbf{state}(i, \mathbf{q}), q_0(\mathbf{q}). \end{aligned}$$

#### C.2.4 Correctness

Here we argue the correctness of the reduction.

**First Direction** Suppose that  $M$  has an accepting computation trace on input word  $w$ . We have to show that the transducer network  $\mathcal{N}$  for  $w$  is inconsistent.

The accepting computation trace of  $M$  is a sequence of configurations, and we identify each configuration by their (one-based) ordinal. We always have  $i \leq 2^z$ . Let  $I$  be the input instance for  $\mathcal{N}$  consisting of the following facts:

- facts  $\mathbf{state}(i, q_i)$  and  $\mathbf{head}(i, h_1, \dots, h_z)$  for each configuration  $i$ , where  $q_i$  and  $(h_1 \dots h_z)$  are respectively the state and head position of  $i$ ;
- fact  $\mathbf{tape}(i, a_1, \dots, a_z, s)$  for each configuration  $i$  and each address  $(a_1 \dots a_z) \in [0, 2^z]$ , where  $s \in \Gamma$  is the contents of cell  $(a_1 \dots a_z)$  in configuration  $i$ ;
- fact  $s(s)$  for each  $s \in \Gamma$ ; fact  $q(q)$  for each  $q \in Q$ ; facts  $0(0)$ ,  $1(1)$ ,  $01(0)$ , and  $01(1)$ ; and, fact  $A(a)$ .

Note, no **error**-message can be sent on this instance (cf. Section C.2.2). Hence, it is sufficient to show that  $\mathbf{accept}()$  can be sent, so that input fact  $A(a)$  gives rise to the messages  $A_{\text{msg}}(a)$  and  $B_{\text{msg}}(a)$ . Then there exist two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  so that  $T(a)$  is created in  $\mathcal{R}_1$  and not in  $\mathcal{R}_2$  or any extension thereof.

Let  $e$  denote the last configuration of the computation trace. The state of  $e$  is  $q_{\text{accept}}$ . Looking at the rules for sending **accept**-messages (Section 6.2), since  $I$  contains  $\mathbf{state}(e, q_{\text{accept}})$  and  $q_{\text{accept}}(q_{\text{accept}})$ , we are left to show that the following messages can be sent:  $\mathbf{start}(1)$  and  $\mathbf{reach}_m(1, e)$  for some  $m \in [0, z]$ . Because configuration 1 is the start configuration of the computation trace, and because we have accurately described this configuration in the input relations, we can see that  $\mathbf{start}(1)$  can be sent. Similarly, we can see that for each pair  $(i, j)$  of subsequent configurations in the trace, the message  $\mathbf{reach}_0(i, j)$  can be sent. And because the  $\mathbf{reach}_m$ -rules with  $m \in [0, z]$  allow us to connect configurations over arbitrary distances within  $[1, 2^z]$ , we can also send  $\mathbf{reach}_m(1, e)$  for some  $m \in [0, z]$ .



**Second Direction** Suppose that the transducer network  $\mathcal{N}$  for  $w$  is inconsistent. We have to show that  $M$  has an accepting computation trace on  $w$ .

First, because  $\mathcal{N}$  is inconsistent, there exists an input instance  $I$  for  $\mathcal{N}$ , and two runs  $\mathcal{R}_1$  and  $\mathcal{R}_2$  of  $\mathcal{N}$  on  $I$ , such that  $\text{last}(\mathcal{R}_1)$  contains an output fact  $T(a)$  that is not in  $\text{last}(\mathcal{R}_2)$ , and  $T(a)$  can not be created in any extension of  $\mathcal{R}_2$ .

We first show that **accept**() can be sent on input  $I$  and that **error**() can not. The presence of  $T(a)$  in  $\text{last}(\mathcal{R}_1)$  implies that the message  $A_{\text{msg}}(a)$  can be sent. This in turn implies that **accept**() can be sent. Now, since by static send rules the message  $A_{\text{msg}}(a)$  can also be sent in an extension of  $\mathcal{R}_2$ , the reason why  $T(a)$  can not be created in that extension is that the memory fact  $B(a)$  is present *and* that the message **error**() can never be delivered, and hence can never be sent.

Looking at the sending rules for relation **accept**, the sending of **accept**() in  $\mathcal{R}_1$  must have been caused by the joint occurrence of the following four facts during some transition of  $\mathcal{N}$ : the message facts **start**( $x$ ) and **reach** <sub>$m$</sub> ( $x, y$ ) for some  $x, y \in \text{adom}(I)$  and  $m \in [0, z]$ , and the input facts **state**( $y, q$ ) and  $q_{\text{accept}}(q)$ . The input facts together already imply that  $y$  could describe an accepting configuration. Now we have to look at the derivation histories of the two messages to construct a full accepting computation trace.

As a general remark, because **error**() can never be sent, the input satisfies the restrictions enforced in Section C.2.2. In particular, each configuration has at most one state and at most one head position in relations **state** and **head** respectively, and each configuration has at most one symbol for each tape cell in relation **tape**. So, the presence of the message **start**( $x$ ) implies that  $x$  not only has *precisely* one state, one head position and one symbol in each tape cell, but also that  $x$  satisfies the additional properties of a valid start configuration. Hence,  $x$  is a fully specified start configuration.

The presence of the message **reach** <sub>$m$</sub> ( $x, y$ ) implies there is a sequence of configurations  $c_1, \dots, c_e$  in the input with  $c_1 = x$  and  $c_e = y$  and such that the message **reach**<sub>0</sub>( $i, j$ ) can be sent for each pair  $(i, j)$  of subsequent configurations. Again using the absence of **error**(), the presence of the message **reach**<sub>0</sub>( $i, j$ ) implies that configurations  $i$  and  $j$  each have *precisely* one state, one head position, and one symbol in each tape cell, and that there exists a valid transition rule of Turing machine  $M$  to explain how configuration  $j$  follows from configuration  $i$ . Finally, using that  $y$  is accepting (see above), we have found an accepting computation trace of  $M$  on  $w$ .

## D Expressivity Upper Bound

### D.1 Correctness Part 1

Let  $\Phi$  be as constructed in Section 7.2.3. Let  $H$  be an arbitrary distributed database instance over  $\text{in}^{\mathcal{N}}$ . Abbreviate  $I = \langle H \rangle^{\mathcal{N}}$ . Let  $\mathbf{f} \in \Phi(I)$ . We have to show that  $\mathbf{f}$  is output at node  $x$  when  $\mathcal{N}$  is run on  $H$ . It is sufficient to show that  $\mathbf{f}$  is output by  $\mathcal{M}$  on input  $I$ .

We remind that Section 7.2.2 contains common concepts and notations. Helper claims can be found in Section D.3.

### D.1.1 Satisfying valuation

Since  $\mathbf{f} \in \Phi(I)$ , program  $\Phi$  contains a UCQ<sup>-</sup>-program  $derive_{G, \mathcal{T}_0}$  such that  $\mathbf{f} \in derive_{G, \mathcal{T}_0}(I)$ . Hence, there exists a subset  $G_0 \subseteq forest_R$  and an equivalence relation  $E$  on  $adom(G_0)$  such that  $G = E(G_0)$  and  $\mathcal{T}_0 \in G$ .

Like before, we regard  $derive_{G, \mathcal{T}}$  as an  $\exists$ FO-formula, where  $\mathcal{T}$  is the truncated version of  $\mathcal{T}_0$  and  $\kappa$  is the canonical scheduling of  $\mathcal{T}_0$ :

$$derive_{G, \mathcal{T}_0} := \exists \bar{z} (diffVal_G \wedge sndMsg_G \wedge succeed_{G, \mathcal{T}, \kappa}).$$

Here, free variables are constituted by the tuple  $\bar{x}$  of values occurring in the root fact of  $\mathcal{T}_0$ , and  $\bar{z}$  are the values in  $adom(G)$  that are not in  $\bar{x}$ . Since  $\mathbf{f} \in derive_{G, \mathcal{T}_0}(I)$ , there exists a valuation  $Val : adom(G) \rightarrow adom(I)$  that makes the following quantifier-free formula true:

$$diffVal_G \wedge sndMsg_G \wedge succeed_{G, \mathcal{T}, \kappa}.$$

The part  $diffVal_G$  makes  $Val$  injective.

### D.1.2 Concrete run

For each tree  $\mathcal{T}' \in G$ , for each internal node  $x$  of  $\mathcal{T}'$ , we can apply the function  $Val$  after valuation  $val^{\mathcal{T}}(x)$ . The resulting valuations still satisfy the nonequalities of the rules, because these nonequalities are satisfied under  $val^{\mathcal{T}}(x)$  and  $Val$  is injective. Let  $F$  denote the forest of (structurally equivalent) derivation trees obtained from  $G$  in this way. Following the principle of canonical runs of Section 7.2.3, we will concurrently execute all trees in  $F$  by their canonical scheduling. This results in a run  $\mathcal{R}$ , whose length is the largest height of any tree in  $F$ . We now show that  $\mathbf{f}$  is derived in  $\mathcal{R}$ .

Let  $\mathcal{T}_0$  be as above. Let  $\mathcal{S}_0 \in F$  be the structurally equivalent tree. We first show that  $fact^{\mathcal{S}_0}(root^{\mathcal{S}_0}) = \mathbf{f}$ . The tuple of values in  $fact^{\mathcal{T}_0}(root^{\mathcal{T}_0})$  are the free variables of  $derive_{G, \mathcal{T}_0}$ . Thus  $Val(fact^{\mathcal{T}_0}(root^{\mathcal{T}_0})) = \mathbf{f}$ . And by construction of  $F$ , we have  $fact^{\mathcal{S}_0}(root^{\mathcal{S}_0}) = Val(fact^{\mathcal{T}_0}(root^{\mathcal{T}_0}))$ .

Henceforth, we will focus on the truncated trees  $\mathcal{T}$  and  $\mathcal{S}$  of  $\mathcal{T}_0$  and  $\mathcal{S}_0$  respectively. The canonical scheduling  $\kappa$  of  $\mathcal{T}_0$  is also defined on  $\mathcal{S}$ . Now, using the order implied by  $\kappa$ , we show by induction on  $x \in \alpha^{\mathcal{S}}$  that  $fact^{\mathcal{S}}(x)$  is derived in transition  $\kappa(x)$  of  $\mathcal{R}$ . So, let  $x \in \alpha^{\mathcal{S}}$  be a node such that for all alpha child nodes  $y$  of  $x$ , the fact  $fact^{\mathcal{S}}(y)$  is derived in transition  $\kappa(y)$  of  $\mathcal{R}$ .<sup>18</sup> We show that  $val^{\mathcal{S}}(x)$  is satisfying for  $rule^{\mathcal{S}}(x)$  in transition  $\kappa(x)$ . The nonequalities of  $rule^{\mathcal{S}}(x)$  are satisfied because they are satisfied under  $val^{\mathcal{T}}(x)$  and because  $Val$  is injective. Next, we differentiate between the different kinds of atoms in the body of  $rule^{\mathcal{S}}(x)$ .

**Input** Let  $\mathbf{l} \in body^{\mathcal{S}}(x)|_{\Upsilon_{in}}$ . We have to show  $I \models \mathbf{l}$ . Let  $\mathbf{l}' \in body^{\mathcal{T}}(x)|_{\Upsilon_{in}}$  be such that  $\mathbf{l} = Val(\mathbf{l}')$ . By construction,  $\mathbf{l}'$  occurs in the conjunction  $succeed_{G, \mathcal{T}, \kappa}^{in}$ , and since this formula is true under  $Val$  with respect to  $I$ , we have  $I \models Val(\mathbf{l}')$  or equivalently  $I \models \mathbf{l}$ , as desired.

<sup>18</sup>This property is automatically satisfied in the base case, where  $x$  has no alpha child nodes.

**Messages** Let  $\mathbf{l} \in \text{body}^{\mathcal{S}}(x)|_{\Upsilon_{\text{msg}}}$ . Abbreviate  $i = \kappa(x)$ . We have to show that  $\mathbf{l}$  is delivered in transition  $i$  of  $\mathcal{R}$ . Because  $\text{rule}^{\mathcal{S}}(x)$  is message-positive,  $\mathbf{l}$  is a fact. Let  $\mathbf{g} \in \text{body}^{\mathcal{T}}(x)|_{\Upsilon_{\text{msg}}}$  be such that  $\mathbf{l} = \text{Val}(\mathbf{g})$ . Because  $\kappa$  is an alignment for  $\mathcal{T}$  with respect to the abstract canonical run  $\mathcal{R}^G$ , we have  $\mathbf{g} \in M_i^G$ . By Claim D.1, the fact  $\mathbf{l} = \text{Val}(\mathbf{g})$  is delivered during transition  $i$  of  $\mathcal{R}$ , as desired.

**Positive output and memory** Let  $\mathbf{l} \in \text{body}^{\mathcal{S}}(x)|_{\Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}}}$  be such that  $\mathbf{l}$  is positive. There is an alpha child  $y$  of  $x$  such that  $\text{fact}^{\mathcal{S}}(y) = \mathbf{l}$ . By assumption on  $x$ ,  $\text{fact}^{\mathcal{S}}(y)$  is derived during transition  $\kappa(y)$  of  $\mathcal{R}$ , and thus  $\mathbf{l}$  is available during transition  $\kappa(x)$ , as desired.

**Negative output and memory** Let  $\mathbf{l} \in \text{body}^{\mathcal{S}}(x)|_{\Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}}}$  be such that  $\mathbf{l}$  is negative. Denote  $\mathbf{l} = \neg \mathbf{g}$ . We show that  $\mathbf{g}$  is not derived before transition  $\kappa(x)$  of  $\mathcal{R}$ . To relate back to  $\mathcal{T}$ , there is also a fact  $\mathbf{h}$  such that  $\mathbf{g} = \text{Val}(\mathbf{h})$  and  $\neg \mathbf{h} \in \text{body}^{\mathcal{T}}(x)$ .

Towards a proof by contradiction, suppose that  $\mathbf{g}$  is derived in some transition  $j < \kappa(x)$  of  $\mathcal{R}$ . Then it is possible to extract a truncated derivation tree  $\mathcal{S}'$  from  $\mathcal{R}$  with  $\text{fact}^{\mathcal{S}'}(\text{root}^{\mathcal{S}'}) = \mathbf{g}$ , together with an alignment  $\kappa'$  of  $\mathcal{S}'$  such that for all alpha nodes  $z$  of  $\mathcal{S}'$ , the fact  $\text{fact}^{\mathcal{S}'}(z)$  is derived during transition  $\kappa'(z)$  of  $\mathcal{R}$  because  $\text{val}^{\mathcal{S}'}(z)$  is satisfying for  $\text{rule}^{\mathcal{S}'}(z)$ . Note that  $\text{Val}^{-1}$  is defined because  $\text{Val}$  is injective. Let  $\mathcal{T}'$  be the truncated derivation tree obtained from  $\mathcal{S}'$  by applying for each alpha node  $z$ , the function  $\text{Val}^{-1}$  after the valuation  $\text{val}^{\mathcal{S}'}(z)$ . The tree  $\mathcal{T}'$  has root fact  $\text{Val}^{-1}(\mathbf{g}) = \mathbf{h}$ .

There exists  $y \in \beta^{\mathcal{T}}(x)$  with  $\text{fact}^{\mathcal{T}}(y) = \mathbf{h}$ . Suppose we would also know that  $(\mathcal{T}', \kappa') \in \text{align}^G(\mathbf{h})$  (shown below). Then the subformula  $\text{succeed}_{G, \mathcal{T}', \kappa'}^{\text{deny}}$  contains the subformula  $\neg \text{succeed}_{G, \mathcal{T}', \kappa'}$ , which is true under  $\text{Val}$ . Equivalently,  $\text{succeed}_{G, \mathcal{T}', \kappa'}$  is false under  $\text{Val}$ . We will use this information to show that at least one alpha node  $z$  of  $\mathcal{T}'$  exists for which valuation  $\text{Val} \circ \text{val}^{\mathcal{T}'}(z)$  is not satisfying for  $\text{rule}^{\mathcal{T}'}(z)$  during transition  $\kappa'(z)$  of  $\mathcal{R}$ , or equivalently, valuation  $\text{Val} \circ \text{Val}^{-1} \circ \text{val}^{\mathcal{S}'}(z) = \text{val}^{\mathcal{S}'}(z)$  is not satisfying for  $\text{rule}^{\mathcal{S}'}(z)$  during transition  $\kappa'(z)$ . This gives the desired contradiction.

Since  $\text{succeed}_{G, \mathcal{T}', \kappa'}$  is false under  $\text{Val}$ , it must be that either  $\text{succeed}_{G, \mathcal{T}', \kappa'}^{\text{in}}$  is false or  $\text{succeed}_{G, \mathcal{T}', \kappa'}^{\text{deny}}$  is false. In the first case, there is an alpha node  $z$  of  $\mathcal{T}'$  and a literal  $\mathbf{l} \in \text{body}^{\mathcal{T}'}(z)|_{\Upsilon_{\text{in}}}$  such that  $I \not\models \text{Val}(\mathbf{l})$ . This immediately gives that  $\text{Val} \circ \text{val}^{\mathcal{T}'}(z)$  is not satisfying for  $\text{rule}^{\mathcal{T}'}(z)$  during any transition of  $\mathcal{R}$ , hence, not in transition  $\kappa'(z)$ , as desired.

Now suppose that  $\text{succeed}_{G, \mathcal{T}', \kappa'}^{\text{deny}}$  is false under  $\text{Val}$ . Thus,  $\text{succeed}_{G, \mathcal{T}', \kappa'}^{\text{deny}}$  contains a subformula  $\neg \text{succeed}_{G, \mathcal{T}'', \kappa''}$  where  $\text{succeed}_{G, \mathcal{T}'', \kappa''}$  is true under  $\text{Val}$ . Hence, there is an alpha node  $z$  of  $\mathcal{T}'$ , with a beta child  $u$ , letting  $\mathbf{i} = \text{fact}^{\mathcal{T}'}(u)$ , and there is a pair  $(\mathcal{T}'', \kappa'') \in \text{align}^G(\mathbf{i})$  with  $\kappa''(\text{root}^{\mathcal{T}''}) < \kappa'(z)$ . Let  $\mathcal{S}''$  be the (truncated) derivation tree obtained from  $\mathcal{T}''$  by applying  $\text{Val}$  after all valuations. Now, using the natural recursion on  $\text{succeed}_{G, \mathcal{T}'', \kappa''}$ , it is possible to show that  $(\mathcal{S}'', \kappa'')$  derives  $\text{Val}(\mathbf{i})$  during earlier transition  $\kappa''(\text{root}^{\mathcal{T}''}) < \kappa'(z)$ . This reasoning ends, because in each recursive step we come strictly closer to the beginning of  $\mathcal{R}$ , and eventually we only use formulas of the form  $\text{succeed}_{G, \_, \_}^{\text{in}}$ . Since valuation  $\text{Val} \circ \text{val}^{\mathcal{T}'}(z)$  requires the absence of  $\text{Val}(\mathbf{i})$  during  $\kappa'(z)$ , and

$Val(i)$  is present in  $\kappa'(z)$ , this valuation is not satisfying during transition  $\kappa'(z)$  of  $\mathcal{R}$ , as desired.

Let  $\mathcal{T}'$  and  $\kappa'$  be as above. We are left to show that  $(\mathcal{T}', \kappa') \in align^G(\mathbf{h})$ . First, because  $\kappa'$  is an alignment for  $\mathcal{S}'$ , and because  $\mathcal{T}'$  and  $\mathcal{S}'$  are structurally equivalent,  $\kappa'$  is a scheduling for  $\mathcal{T}'$ . Next, let  $z$  be an internal (alpha) node of  $\mathcal{T}'$ . Let  $\mathbf{l} \in body^{\mathcal{T}'}(z)|_{\Upsilon_{msg}}$ , where  $\mathbf{l}$  is a fact by message-positivity of  $rule^{\mathcal{T}'}(z)$ . We have to show that  $\mathbf{l} \in M_j^G$  where  $j = \kappa'(z)$ . Since  $val^{\mathcal{T}'}(z) = Val^{-1} \circ val^{\mathcal{S}'}(z)$ , we can consider the fact  $\mathbf{i} \in body^{\mathcal{S}'}(z)|_{\Upsilon_{msg}}$  such that  $\mathbf{l} = Val^{-1}(\mathbf{i})$ . Now, since  $\kappa'$  is an alignment for  $\mathcal{S}'$  with respect to  $\mathcal{R}$ , we know that  $\mathbf{i}$  is delivered in transition  $j$  of  $\mathcal{R}$ . Then, by Claim D.1, there is a fact  $\mathbf{l}' \in M_j^G$  such that  $Val(\mathbf{l}') = \mathbf{i}$ . But by injectivity of  $Val$ , this means  $\mathbf{l}' = Val^{-1}(\mathbf{i}) = \mathbf{l}$ , as desired.

## D.2 Correctness Part 2

Let  $H$  be an arbitrary input over  $in^{\mathcal{N}}$ . Abbreviate  $I = \langle H \rangle^{\mathcal{N}}$ . Let  $\mathbf{f}$  be an  $R$ -fact output at node  $x$  when  $\mathcal{N}$  is run on  $H$ . This implies that  $\mathcal{M}$  outputs  $\mathbf{f}$  on input  $I$ . We have to show that  $\mathbf{f} \in \Phi(I)$ , with  $\Phi$  as constructed in Section 7.2.3.

Let  $\Pi$  denote the transducer of  $\mathcal{M}$ . We remind that Section 7.2.2 contains common concepts and notations. Additionally, for two structurally equivalent derivation trees  $\mathcal{T}$  and  $\mathcal{S}$ , we write  $map_{\mathcal{T}, \mathcal{S}}$  to denote the structural bijection from nodes of  $\mathcal{T}$  to nodes of  $\mathcal{S}$ . Lastly, helper claims can be found in Appendix D.3.

### D.2.1 Collecting trees

On input  $I$ , from each run of  $\mathcal{M}$  in which  $\mathbf{f}$  is output, we can extract a derivation tree for  $\mathbf{f}$ . Now, let  $F$  be a maximal set of derivation trees for  $\mathbf{f}$  extracted from all possible runs of  $\mathcal{M}$  on  $I$ , such that no two trees are structurally equivalent. Set  $F$  is finite because  $\Pi$  is recursion-free.

### D.2.2 Canonical run

Following the principle of canonical runs from Section 7.2.3, we can concurrently execute all trees of  $F$ . This results in a run  $\mathcal{R}$  whose length is the height of the largest tree in  $F$ .

We now show that  $\mathbf{f}$  is derived in  $\mathcal{R}$ . Because  $\mathcal{M}$  outputs  $\mathbf{f}$  on input  $I$ , consistency of  $\mathcal{M}$  implies that  $\mathcal{R}$  can always be extended to a run  $\mathcal{R}'$  in which  $\mathbf{f}$  is output. From  $\mathcal{R}'$ , we can extract a pair  $(\mathcal{T}, \kappa)$  of a concrete derivation tree for  $\mathbf{f}$  and a scheduling for this tree, such that for each  $x \in int^{\mathcal{T}}$  the fact  $fact^{\mathcal{T}}(x)$  is derived during transition  $\kappa(x)$  of  $\mathcal{R}'$  by applying  $val^{\mathcal{T}}(x)$  to  $rule^{\mathcal{T}}(x)$ . There is some tree  $\mathcal{S} \in F$  structurally equivalent to  $\mathcal{T}$ . Using the order implied by canonical scheduling  $\kappa^{\mathcal{S}}$ , we show by induction on the alpha nodes  $x \in \alpha^{\mathcal{S}}$  that  $fact^{\mathcal{S}}(x)$  is derived during transition  $\kappa^{\mathcal{S}}(x)$  by applying valuation  $val^{\mathcal{S}}(x)$  to  $rule^{\mathcal{S}}(x)$ . Let  $x \in \alpha^{\mathcal{S}}$ , assuming for each descendant  $y \in \alpha^{\mathcal{S}}$  of  $x$  that  $fact^{\mathcal{S}}(y)$  is derived during transition  $\kappa^{\mathcal{S}}(y)$ .

**Input** Since  $\mathcal{S} \in F$ , the tree  $\mathcal{S}$  was extracted from a run, and hence, the input literals of  $rule^{\mathcal{S}}(x)$  must be satisfied under  $val^{\mathcal{S}}(x)$ .

**Messages** Moreover, because sending rules are message-positive and static, it can be shown that the messages needed by  $rule^S(x)$  under  $val^S(x)$  are delivered in  $\mathcal{R}$  during transition  $\kappa^S(x)$  (details omitted).

**Output and memory** Using the assumption on descendant alpha nodes of  $x$ , the positive output and memory facts required by  $val^S(x)$  are also satisfied.

As the last step, we show that the negative output and memory literals under  $val^S(x)$  are absent during transition  $\kappa^S(x)$ . Let us abbreviate  $n = map_{\mathcal{S}, \mathcal{T}}$  (defined in Section 7.2.2). Since  $\mathcal{S}$  and  $\mathcal{T}$  are structurally equivalent and both derive the root fact  $\mathbf{f}$ , we can apply Claim D.2 to know that the valuations  $val^S(x)$  and  $val^T(n(x))$  assign the same values to the free variables of  $rule^S(x)$ . By selection of  $(\mathcal{T}, \kappa)$ , the output and memory facts that rule  $rule^S(x)$  tests for absence under  $val^T(n(x))$ , are effectively absent during transition  $\kappa(n(x))$  of  $\mathcal{R}'$ . Now, because  $\Pi$  is inflationary, if we would know  $\kappa^S(x) \leq \kappa(n(x))$ , then these same output and memory facts must also be absent during transition  $\kappa^S(x)$ , as desired. We are left to show that  $\kappa^S(x) \leq \kappa(n(x))$ . By definition of canonical scheduling  $\kappa^S$ , transition  $\kappa^S(x)$  is the earliest transition of  $\mathcal{R}$  in which the rule  $rule^S(x)$  can be executed if the derivation strategy represented by  $\mathcal{S}$  must be followed.<sup>19</sup> Now, since the subtree under  $x$  in  $\mathcal{S}$  is structurally equivalent to the subtree under  $n(x)$  in  $\mathcal{T}$ , we have  $\kappa^S(x) \leq \kappa(n(x))$ .

### D.2.3 Create valuation

From Section 7.2.3, recall the set  $forest_R$ , in which no two trees are structurally equivalent. For each tree  $\mathcal{T} \in F$ , there is a unique tree  $\mathcal{S} \in forest_R$  that is structurally equivalent to  $\mathcal{T}$ . Let  $G_0 \subseteq forest_R$  be all these trees. We define a function  $Val_0 : adom(G_0) \rightarrow adom(F)$ , giving rise to an equivalence relation on  $adom(G_0)$ .

First, let  $\mathcal{S} \in G_0$ . We can uniquely identify a *component* of a positive atom in  $\mathcal{S}$  by a triple  $(p, \mathbf{a}, i)$ , where  $p$  is a path followed from the root towards an internal node  $x$  of  $\mathcal{S}$ ;  $\mathbf{a}$  is the head or a positive body atom of  $rule^S(x)$ ; and,  $i$  is a component index in  $\mathbf{a}$ . Here,  $p$  can be uniquely specified as the sequence of atoms  $lit^S(x)$  labelling the encountered internal nodes  $x$ . Two components  $(p_1, \mathbf{a}_1, i_1)$  and  $(p_2, \mathbf{a}_2, i_2)$  belong to the same rule if  $p_1 = p_2$ . Now, we define an equivalence relation over the components in a bottom-up way, as follows. Starting at an internal node  $x$  without other internal nodes as children, two components in  $rule^S(x)$  are equivalent if they contain the same variable. Going to the parent  $y$  of  $x$ , two components  $c_1$  and  $c_2$  in  $rule^S(y)$  are equivalent if (i) they contain the same variable; or (ii) they occur together in a positive body atom  $\mathbf{a}$  of  $rule^S(y)$ , and for the child  $x$  of  $y$  with  $lit^S(x) = \mathbf{a}$ , the components in the head of  $rule^S(x)$  corresponding to  $c_1$  and  $c_2$  are equivalent. The equivalence relation on the components of  $\mathcal{S}$  is unique, and its number of equivalence classes upper bounds the active domain size of  $\mathcal{S}$ .

Now we define function  $Val_0 : adom(G_0) \rightarrow adom(F)$ . Let  $\mathcal{S} \in G_0$  and let  $\mathcal{T} \in F$  denote the structurally equivalent tree. Because  $\mathcal{S}$  and  $\mathcal{T}$  are structurally equivalent, the equivalence classes on components of  $\mathcal{S}$  transfer naturally to equivalence classes on the components of  $\mathcal{T}$ . Because  $\mathcal{S}$  is general, its valuations

<sup>19</sup>Indeed, for each subtree, the minimum number of transitions required to derive its root fact is precisely the height of this tree, and this is expressed in the canonical scheduling.

assign a different value to each equivalence class, so we can define a function  $V_S : \text{adom}(\mathcal{S}) \rightarrow \text{adom}(\mathcal{T})$  that contains for each equivalence class  $e$  of  $\mathcal{S}$  the mapping  $(a \mapsto b)$ , where  $a$  and  $b$  are the values assigned to  $e$  by  $\mathcal{S}$  and  $\mathcal{T}$  respectively. For the entire set  $G_0$ , we take the union of all mappings  $V_S$  with  $\mathcal{S} \in G_0$ . The result is denoted  $Val_0$ , and this is a function because each tree in  $G_0$  has a disjoint active domain. We can now define an equivalence relation  $E$  on  $\text{adom}(G_0)$ : two values are equivalent if their image under  $Val_0$  is the same. Assuming an order on **dom** (the same order as in Section 7.2.3), we can replace each value in  $\text{adom}(G_0)$  by the smallest value in its equivalence relation. This results in a set  $G$  of derivation trees, in which still as many structurally different trees occur as in  $G_0$ , and with  $\text{adom}(G) \subseteq \text{adom}(G_0)$ .<sup>20</sup>

Let  $Val$  denote the restriction of  $Val_0$  to  $\text{adom}(G)$ ; this function is injective.

#### D.2.4 Satisfying valuation

Let  $F$ ,  $G$ , and  $Val$  be as previously defined. For each tree  $\mathcal{S} \in G$ , if we would apply  $Val$  after each valuation in  $\mathcal{S}$ , we obtain a tree in  $F$ . So, if we would consider  $\text{adom}(\mathcal{S})$  to be variable symbols, then we can see  $Val$  as an assignment to these variables. This will be used below to show that  $\mathbf{f} \in \Phi(I)$ .

As shown above, there is a derivation tree  $\mathcal{T} \in F$  that derives  $\mathbf{f}$  in  $\mathcal{R}$ , when executed according to its canonical scheduling. Let  $\mathcal{S}_0 \in G$  be the tree that is structurally equivalent to  $\mathcal{T}$ . As remarked above, applying  $Val$  to  $\mathcal{S}_0$  gives  $\mathcal{T}$ . Let  $\mathcal{S}$  denote the truncated version of  $\mathcal{S}_0$ , and let  $\kappa$  denote the restriction of the canonical scheduling of  $\mathcal{S}_0$  to the remaining nodes. Recalling the construction in Section 7.2.3, we have added to the  $\text{UCQ}^-$ -program  $\Phi$  the  $\text{UCQ}^-$ -program  $\text{derive}_{G,\mathcal{S}}$ , given by the following equivalent  $\exists\text{FO}$ -formula:

$$\text{derive}_{G,\mathcal{S}} := \exists \bar{z} (\text{diffVal}_G \wedge \text{sndMsg}_G \wedge \text{succeed}_{G,\mathcal{S},\kappa}),$$

where  $\bar{z}$  is an ordering of the values in  $\text{adom}(G)$  not occurring in the tuple  $\bar{x}$  in the root fact of  $\mathcal{S}$ . So,  $\bar{x}$  are the free variables. Now, denoting  $\mathbf{f} = R(\bar{a})$ , to show  $\mathbf{f} \in \Phi(I)$ , it suffices to show that if  $\bar{x}$  is assigned  $\bar{a}$  then the resulting sentence is true with respect to  $I$ . This amounts to showing that the following quantifier-free formula is true under  $Val$  with respect to  $I$ :

$$\text{diffVal}_G \wedge \text{sndMsg}_G \wedge \text{succeed}_{G,\mathcal{S},\kappa}.$$

**Diffval and sndMsg** The subformula  $\text{diffVal}_G$  is true because  $Val$  is injective on  $\text{adom}(G)$ . Next, the subformula  $\text{sndMsg}_G$  is a large conjunction of input literals from the sending rules in  $G$ . Let  $\mathbf{l}$  be such a literal. We have to show  $I \models Val(\mathbf{l})$ . There exists a tree  $\mathcal{S}' \in G$  and an internal node  $x$  of  $\mathcal{S}'$  such that  $\text{rule}^{\mathcal{S}'}(x)$  is a sending rule and  $\mathbf{l} \in \text{body}^{\mathcal{S}'}(x)|_{\Upsilon_{\text{in}}}$ . Let  $\mathcal{T}' \in F$  be the tree structurally equivalent to  $\mathcal{S}'$ , and abbreviate  $n' = \text{map}_{\mathcal{S}',\mathcal{T}'}$ . By construction of  $Val$ , we have  $Val(\mathbf{l}) \in \text{body}^{\mathcal{T}'}(n'(x))$ . Since  $\text{val}^{\mathcal{T}'}(n'(x))$  was satisfied during some run, which follows from  $\mathcal{T}' \in F$ , and all runs have the same input facts, we obtain  $I \models Val(\mathbf{l})$ .

<sup>20</sup>Nonequalities in rules of  $G$  are satisfied under their valuations because they are satisfied in  $F$ .

**Succeed input** Now consider the subformula  $succeed_{G,S,\kappa}$ . This formula is specified as

$$succeed_{G,S,\kappa} := succeed_{G,S,\kappa}^{\text{in}} \wedge succeed_{G,S,\kappa}^{\text{deny}}.$$

Let  $\mathcal{S}_0$  and  $\mathcal{T} \in F$  be as above:  $\mathcal{S}$  is the truncated version of  $\mathcal{S}_0$  and  $\mathcal{T}$  is structurally equivalent to  $\mathcal{S}_0$ . Abbreviate  $n = \text{map}_{\mathcal{S}_0, \mathcal{T}}$ .

Similarly to  $\text{sndMsg}_G$ , the subformula  $succeed_{G,S,\kappa}^{\text{in}}$  is a conjunction of input literals. Let  $\mathbf{l}$  be such a literal. We have to show  $I \models \text{Val}(\mathbf{l})$ . There exists a node  $x \in \alpha^{\mathcal{S}}$  such that  $\mathbf{l} \in \text{body}^{\mathcal{S}}(x)|_{\Upsilon_{\text{in}}}$ . By construction of  $\text{Val}$ , we have  $\text{Val}(\mathbf{l}) \in \text{body}^{\mathcal{T}}(n(x))$ . And similarly to our reasoning for  $\text{sndMsg}_G$ , we can now obtain that  $I \models \text{Val}(\mathbf{l})$ .

**Succeed deny** Consider the subformula  $succeed_{G,S,\kappa}^{\text{deny}}$ . Let  $x \in \alpha^{\mathcal{S}}$ ,  $y \in \beta^{\mathcal{S}}(x)$ , denoting  $\mathbf{g} = \text{fact}^{\mathcal{S}}(y)$ , and  $(\mathcal{S}', \lambda) \in \text{align}^G(\mathbf{g})$  with  $\lambda(\text{root}^{\mathcal{S}'}) < \kappa(x)$ . We have to show that  $\neg \text{succeed}_{G,S',\lambda}$  is true under  $\text{Val}$ , which amounts to showing that  $\text{succeed}_{G,S',\lambda}$  is false under  $\text{Val}$ . The main strategy will be to use that  $\mathcal{S}'$  extended with  $\text{Val}$  fails in  $\mathcal{R}$  when executed according to  $\lambda$ . The reasons for failure make (parts of) formula  $\text{succeed}_{G,S',\lambda}$  false.

First, we show that the fact  $\text{Val}(\mathbf{g})$  has to be absent during (and before) transition  $\kappa(x)$  of  $\mathcal{R}$ . By definition of  $y$ , we have  $\neg \mathbf{g} \in \text{body}^{\mathcal{S}}(x)$ . Let  $\mathcal{S}_0$ ,  $\mathcal{T} \in F$ , and mapping  $n$ , be as above for the case “succeed input”. We have  $\neg \text{Val}(\mathbf{g}) \in \text{Val}(\text{body}^{\mathcal{S}}(x)) = \text{body}^{\mathcal{T}}(n(x))$ . Now, because valuation  $\text{val}^{\mathcal{T}}(n(x))$  is satisfying during transition  $\kappa^{\mathcal{T}}(n(x)) = \kappa(x)$ ,  $\text{Val}(\mathbf{g})$  must be absent during  $\kappa(x)$ . By inflationarity of the transducer,  $\text{Val}(\mathbf{g})$  is thus also absent before  $\kappa(x)$ .

Let  $(\mathcal{S}', \lambda)$  be as above. There must be an alpha node  $z$  of  $\mathcal{S}'$  such that fact  $\text{Val}(\text{fact}^{\mathcal{S}'}(z))$  is not derived during transition  $\lambda(z)$  of  $\mathcal{R}$  because otherwise  $\text{Val}(\text{fact}^{\mathcal{S}'}(\text{root}^{\mathcal{S}'})) = \text{Val}(\mathbf{g})$  would be derived in transition  $\lambda(\text{root}^{\mathcal{S}'}) < \kappa(x)$ , which is false. Let  $z$  be the first of such failed nodes with respect to  $\lambda$ . Valuation  $\text{Val} \circ \text{val}^{\mathcal{S}'}(z)$  is not satisfying for  $\text{rule}^{\mathcal{S}'}(z)$  during transition  $\lambda(z)$  of  $\mathcal{R}$ , and each reason is used to show that some part of formula  $\text{succeed}_{G,S',\lambda}$  is false under  $\text{Val}$ . We consider the different kinds of literal in  $\text{rule}^{\mathcal{S}'}(z)$ :

**[Input]** Suppose there is a literal  $\mathbf{l} \in \text{body}^{\mathcal{S}'}(z)|_{\Upsilon_{\text{in}}}$  such that  $I \not\models \text{Val}(\mathbf{l})$ . Then the conjunction  $\text{succeed}_{G,S',\lambda}^{\text{in}}$ , and hence the entire formula  $\text{succeed}_{G,S',\lambda}$ , is false under  $\text{Val}$  because  $\text{succeed}_{G,S',\lambda}^{\text{in}}$  contains  $\mathbf{l}$ .

**[Messages]** Recall that  $\text{rule}^{\mathcal{S}'}(z)$  is message-positive. Suppose that there is a fact  $\mathbf{l} \in \text{body}^{\mathcal{S}'}(z)|_{\Upsilon_{\text{msg}}}$  such that  $\text{Val}(\mathbf{l})$  is not delivered in transition  $\lambda(z)$  of  $\mathcal{R}$ . We argue that this is actually not possible, so this case can not occur. First, because  $\lambda$  is an alignment of  $\mathcal{S}'$  to the abstract canonical run  $\mathcal{R}^G$ , fact  $\mathbf{l}$  is delivered in transition  $\lambda(z)$  of  $\mathcal{R}^G$ . Hence, by Claim D.1, fact  $\text{Val}(\mathbf{l})$  is delivered in transition  $\lambda(z)$  of  $\mathcal{R}$ .

**[Positive output and memory]** Suppose there is a positive literal  $\mathbf{l} \in \text{body}^{\mathcal{S}'}(z)|_{\Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}}}$  (i.e.,  $\mathbf{l}$  is a fact) such that  $\text{Val}(\mathbf{l})$  is not available during transition  $\lambda(z)$  of  $\mathcal{R}$ . We will again show that this case can not occur. The existence of  $\mathbf{l}$  implies that  $z$  has an alpha child-node  $z'$  in  $\mathcal{S}'$  with  $\text{fact}^{\mathcal{S}'}(z') = \mathbf{l}$ . This implies  $\lambda(z') < \lambda(z)$ . Since  $z$  is the first failed alpha node of  $\mathcal{S}'$  with respect to  $\lambda$ , it must be that the fact  $\text{Val}(\text{fact}^{\mathcal{S}'}(z')) = \text{Val}(\mathbf{l})$  is derived in transition  $\lambda(z')$ . Hence,  $\text{Val}(\mathbf{l})$  is available in transition  $\lambda(z)$  by inflationarity of  $\Pi$ .

**[Negative output and memory]** Suppose there is a negative literal  $\neg i \in \text{body}^{S'}(z)|_{\Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}}}$ , such that  $\mathbf{h} = \text{Val}(\mathbf{i})$  is present during transition  $\lambda(z)$  of  $\mathcal{R}$ . From  $\mathcal{R}$ , we can extract a pair  $(\mathcal{T}'', \lambda'')$  with  $\mathcal{T}''$  a truncated derivation tree for  $\mathbf{h}$ , and  $\lambda''$  an alignment of  $\mathcal{T}''$  to  $\mathcal{R}$  according to which  $\mathcal{T}''$  derives  $\mathbf{h}$ . Note that  $\text{Val}^{-1}$  exists because  $\text{Val}$  is injective. Now, let  $\mathcal{S}''$  denote the tree obtained from  $\mathcal{T}''$  by applying for each internal node  $u$  of  $\mathcal{T}''$  the function  $\text{Val}^{-1}$  after  $\text{val}^{\mathcal{T}''}(u)$ . Note,  $\text{adom}(\mathcal{S}'') \subseteq \text{adom}(G)$ .

Because in  $\mathcal{S}'$  there is a beta child node  $z'$  of  $z$  with  $\text{fact}^{S'}(z') = \mathbf{i}$ , if we could show  $(\mathcal{S}'', \lambda'') \in \text{align}^G(\mathbf{i})$ , then formula  $\text{succeed}_{G, \mathcal{S}', \lambda}^{\text{deny}}$  contains the subformula  $\neg \text{succeed}_{G, \mathcal{S}'', \lambda''}$ . Then, we can recursively show that  $\text{succeed}_{G, \mathcal{S}'', \lambda''}$  is true under  $\text{Val}$ , making  $\text{succeed}_{G, \mathcal{S}', \lambda}^{\text{deny}}$ , and by extension  $\text{succeed}_{G, \mathcal{S}', \lambda}$ , false under  $\text{Val}$ , as desired. This is similar to our current proof where we show that  $\text{succeed}_{G, \mathcal{S}, \kappa}$  is true under  $\text{Val}$ , but we would replace  $(\mathcal{S}, \kappa)$  by  $(\mathcal{S}'', \lambda'')$ . This recursive step always ends, as we argued at the end of Section 7.2.3.

We are left to show that  $(\mathcal{S}'', \lambda'') \in \text{align}^G(\mathbf{i})$ . First,  $\mathcal{S}''$  derives the fact  $\text{Val}^{-1}(\mathbf{h}) = \text{Val}^{-1}(\text{Val}(\mathbf{i})) = \mathbf{i}$ . Next, alignment  $\lambda''$  for  $\mathcal{S}''$  schedules nodes before their ancestors because it also does this for  $\mathcal{T}''$ . For the last step, let  $u$  be an internal node of  $\mathcal{S}''$ . We have to show that each  $\mathbf{e} \in \text{body}^{\mathcal{S}''}(u)|_{\Upsilon_{\text{msg}}}$  is delivered during transition  $\lambda''(u)$  of  $\mathcal{R}^G$ . By construction of  $\mathcal{S}''$  from  $\mathcal{T}''$ , there is some  $\mathbf{e}' \in \text{body}^{\mathcal{T}''}(u)|_{\Upsilon_{\text{msg}}}$  that is delivered in transition  $\lambda''(u)$  of  $\mathcal{R}$  and  $\mathbf{e} = \text{Val}^{-1}(\mathbf{e}')$ . But by Claim D.1, we have  $\mathbf{e}' \in \text{Val}(M_j^G)$  with  $j = \lambda''(u)$ . Hence,  $\mathbf{e} \in \text{Val}^{-1} \circ \text{Val}(M_j^G) = M_j^G$ , as desired.

### D.3 Claims

**Claim D.1.** Consider the symbols defined in Section 7.2.3. Let  $G \subseteq \text{forest}_R$ . Let  $F$  be a set of derivation trees of  $\Pi$  such that (i) no two trees are structurally equivalent; (ii) for each  $\mathcal{T} \in F$  there is a structurally equivalent tree  $\mathcal{S} \in G$ ; and, (iii) there is an *injective* function  $\text{Val} : \text{adom}(G) \rightarrow \text{adom}(F)$  such that when  $\text{Val}$  is applied after the valuations of a tree  $\mathcal{S} \in G$ , we obtain the structurally equivalent tree  $\mathcal{T} \in F$ . Finally, let  $I$  be an input for  $\mathcal{M}$  such that formula  $\text{sndMsg}_G$  is satisfied under  $\text{Val}$  with respect to  $I$ .

Let  $\mathcal{R}^G$  and  $\mathcal{R}$  denote the canonical runs based on  $G$  and  $F$  respectively, that both have the same length  $n$ . Let  $i \in \{1, \dots, n\}$  and let  $M_i^G$  denote the (abstract) message set delivered in transition  $i$  of  $\mathcal{R}^G$ . In transition  $i$  of  $\mathcal{R}$ , we deliver *precisely*  $\text{Val}(M_i^G)$ .

*Proof.* We show this by induction on  $i$ . For the base case ( $i = 1$ ), the property holds because  $M_1^G = \emptyset$  and no messages are delivered in the first transition of  $\mathcal{R}$  (as no messages were previously sent).

For the induction hypothesis, assume the property holds for transitions  $j = 1, \dots, i - 1$  with  $i > 1$ . For the inductive step, we show that the property is satisfied for transition  $i$ . First, note that at most  $\text{Val}(M_i^G)$  can be delivered in transition  $i$  of  $\mathcal{R}$ , because this transition only delivers the messages needed by rules in  $F$  scheduled at  $i$ , and because the trees in  $F$  are obtained from those in  $G$  by concatenating  $\text{Val}$  to their valuations.

For the second direction, let  $\mathbf{g} \in M_i^G$  and denote  $\mathbf{h} = \text{Val}(\mathbf{g})$ . We show that  $\mathbf{h}$  is delivered in transition  $i$  of  $\mathcal{R}$ . Since  $\mathbf{g} \in M_i^G$ , there is a tree  $\mathcal{S}' \in G$ , and an internal node  $x$  of  $\mathcal{S}'$ , such that  $\kappa^{\mathcal{S}'}(x) = i$  and  $\mathbf{g} \in \text{body}^{\mathcal{S}'}(x)|_{\Upsilon_{\text{msg}}}$ . By



message-positivity of  $rule^{S'}(x)$ , there is a child  $y$  of  $x$  such that  $fact^{S'}(y) = \mathbf{g}$ . From the definition of the canonical scheduling, we have  $\kappa^{S'}(y) = \kappa^{S'}(x) - 1$ . Denoting  $j = \kappa^{S'}(y)$ , we have  $j = i - 1$ . We show that  $Val \circ val^{S'}(y)$  is satisfying for  $rule^{S'}(y)$  during transition  $j$ , such that  $Val(\mathbf{g}) = \mathbf{h}$  is sent in transition  $j$ , and can be delivered in (the next) transition  $i$ . The nonequalities of  $rule^{S'}(y)$  are satisfied because they are satisfied under  $val^{S'}(y)$  (by construction of  $G$ ) and because  $Val$  is injective. Next, because  $rule^{S'}(y)$  is static, we only have to consider input and message atoms:

- Let  $\mathbf{l} \in body^{S'}(y)|_{\Upsilon_{in}}$ . We have,  $I \models Val(\mathbf{l})$ , as desired, because  $\mathbf{l}$  is added to  $sndMsg_G$ , which is true under  $Val$  with respect to  $I$ .
- Let  $\mathbf{l} \in body^{S'}(y)|_{\Upsilon_{msg}}$ . Because  $rule^{S'}(y)$  is message-positive,  $\mathbf{l}$  is a fact. Moreover, we have  $\mathbf{l} \in M_j^G$ . By applying the induction hypothesis to transition  $j$ , we know that  $Val(\mathbf{l})$  is delivered during transition  $j$ , as desired.

□

**Claim D.2.** Let  $\mathcal{T}$  and  $\mathcal{S}$  be two structurally equivalent derivation trees of  $\Pi$ , that derive the same output or memory fact  $\mathbf{f}$ . Abbreviate  $n = map_{\mathcal{S}, \mathcal{T}}$ . For each  $x \in \alpha^{\mathcal{S}}$ , the valuations  $val^{\mathcal{S}}(x)$  and  $val^{\mathcal{T}}(n(x))$  assign the same values to the free variables of the rule  $rule^{\mathcal{S}}(x) = rule^{\mathcal{T}}(n(x))$ .<sup>21</sup>

*Proof.* We show the property by induction on the length of the path from the root to the node  $x \in \alpha^{\mathcal{S}}$  in question. In the base case, simply  $x = root^{\mathcal{S}}$  and  $n(x) = root^{\mathcal{T}}$ . We are given that  $fact^{\mathcal{S}}(root^{\mathcal{S}}) = fact^{\mathcal{T}}(root^{\mathcal{T}})$ . Hence, valuations  $val^{\mathcal{S}}(root^{\mathcal{S}})$  and  $val^{\mathcal{T}}(root^{\mathcal{T}})$  assign the same values to free variables. Moreover, because  $\mathbf{f}$  is an output or memory fact,  $rule^{\mathcal{S}}(root^{\mathcal{S}})$  is message-bounded, and thus any variable occurring in an output or memory literal in the body must be a free variable. Hence, for every alpha child  $y$  of  $root^{\mathcal{S}}$ , we have  $fact^{\mathcal{S}}(y) = fact^{\mathcal{T}}(n(y))$ . The reasoning can now be repeated for  $y$ . □

---

<sup>21</sup>Node  $x$  and  $n(x)$  have the same rule by structural equivalence.