

## Relational Transducers for Declarative Networking

TOM J. AMELOOT, Hasselt University & Transnational University of Limburg  
 FRANK NEVEN, Hasselt University & Transnational University of Limburg  
 JAN VAN DEN BUSSCHE, Hasselt University & Transnational University of Limburg

Motivated by a recent conjecture concerning the expressiveness of declarative networking, we propose a formal computation model for “eventually consistent” distributed querying, based on relational transducers. A tight link has been conjectured between coordination-freeness of computations, and monotonicity of the queries expressed by such computations. Indeed, we propose a formal definition of coordination-freeness and confirm that the class of monotone queries is captured by coordination-free transducer networks. Coordination-freeness is a semantic property, but the syntactic class of “oblivious” transducers we define also captures the same class of monotone queries. Transducer networks that are not coordination-free are much more powerful.

Categories and Subject Descriptors: H.2 [Database Management]: Languages; H.2 [Database Management]: Systems—*Distributed databases*; F.1 [Computation by Abstract Devices]: Models of Computation

General Terms: languages, theory

Additional Key Words and Phrases: distributed database, relational transducer, monotonicity, expressive power, cloud programming

### ACM Reference Format:

AMELOOT, T. J., NEVEN, F. and VAN DEN BUSSCHE, J. 2011. Relational Transducers for Declarative Networking. *J. ACM* V, N, Article A (January YYYY), 37 pages.  
 DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

### 1. INTRODUCTION

Declarative networking [Loo et al. 2009] is a recent approach by which distributed computations and networking protocols are modeled and programmed using formalisms based on Datalog. In his keynote speech at PODS 2010 [Hellerstein 2010a; Hellerstein 2010b], Hellerstein made a number of intriguing conjectures concerning the expressiveness of declarative networking. In the present paper, we are focusing on the CALM conjecture (Consistency And Logical Monotonicity). This conjecture suggests a strong link between, on the one hand, “eventually consistent” and “coordination-free” distributed computations, and on the other hand, expressibility in monotonic Datalog (without negation or aggregate functions). The conjecture was not fully formalized, however; indeed, as Hellerstein notes himself, a proper treatment of this conjecture requires crisp definitions of eventual consistency and coordination, which have been lacking so far. Moreover, it also requires a formal model of distributed computation.

---

Tom J. Ameloot is a PhD Fellow of the Fund for Scientific Research, Flanders (FWO).

Author’s email addresses: [tom.ameloot@uhasselt.be](mailto:tom.ameloot@uhasselt.be), [frank.neven@uhasselt.be](mailto:frank.neven@uhasselt.be), [jan.vandenbussche@uhasselt.be](mailto:jan.vandenbussche@uhasselt.be).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).  
 © YYYY ACM 0004-5411/YYYY/01-ARTA \$15.00  
 DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

In the present paper, we investigate the CALM conjecture in the context of a model for distributed database querying. In the model we propose, the computation is performed on a network of relational transducers. The relational transducer model, introduced by Abiteboul and Vianu, is well established in database theory research as a model for data-centric agents reacting to inputs [Abiteboul et al. 2000; Deutsch et al. 2009; Deutsch et al. 2007; Deutsch et al. 2006; Spielmann 2003]. Relational transducers are firmly grounded in the theory of database queries [Abiteboul and Vianu 1991b; Abiteboul and Vianu 1995] and also have close connections with Abstract State Machines [Blass et al. 2002]. It thus seems natural to consider networks of relational transducers, as we will do here. We give a formal operational semantics for such networks, formally define “eventual consistency”, and formally define what it means for a network to compute a conventional database query, in order to address the expressiveness issues raised by Hellerstein.

It is less clear, however, how to formalize the intuitive notion of “coordination”. We do not claim to resolve this issue definitively, but we propose a new, non-obvious definition that appears workable. Distributed algorithms requiring coordination are viewed as less efficient than coordination-free algorithms. Hellerstein has identified *monotonicity* as a fundamental property connected with coordination-freeness. Indeed, monotonicity enables “embarrassing parallelism” [Hellerstein 2010b]: agents working on parts of the data in parallel can produce parts of the output independently, without the need for coordination.

One side of the CALM conjecture now states that any database query expressible in monotonic Datalog can be computed in a distributed setting in an eventually consistent, coordination-free manner. This is the easy side of the conjecture, and indeed we formally confirm it in the following broader sense: any monotone query  $Q$  can be computed by a network of “oblivious” transducers. Oblivious transducers are unaware of the network extent (in a sense that we will make precise), and every oblivious transducer network is coordination-free. Here, we should note that the transducer model is parameterized by the query language  $\mathcal{L}$  that the transducer can use to update its local state. Formally, the monotone query  $Q$  to be computed must be expressible in the while-closure of  $\mathcal{L}$  for the above confirmation to hold. If  $Q$  is defined in Datalog, for example, then  $\mathcal{L}$  can just be unions of conjunctive queries.

The other side of the CALM conjecture, that the query computed by an eventually consistent, coordination-free distributed program is always expressible in Datalog, is false when taken literally, as we will point out. Nevertheless, we do give an extended version of the conjecture that holds. More importantly, we confirm the conjecture in the following more general form: coordination-free networks of transducers can compute only monotone queries. Note that here we are using our newly proposed formal definition of coordination-free.

This paper is organized as follows. Preliminaries are in Section 2. Section 3 introduces networks of transducers. Section 4 investigates the use of transducer networks for expressing conventional database queries in a distributed fashion. Section 5 discusses the issue of coordination, and looks into the CALM conjecture and related results. Section 6 contains results about the expressiveness of transducer networks. Section 7 shortly looks into a variation of the transducer model, and Section 8 is the conclusion.

This paper is the extended version of our conference paper [Ameloot et al. 2011].

### 1.1. Related Work

The desire to better understand coordination in the field of declarative networking is evidenced by the steadily growing literature on this subject. First, Alvaro et al. [Alvaro et al. 2011a] look at coordination as a quantitative property that can be minimized. They describe a program analysis technique to spot syntactical locations in the code where coordination is not needed. The goal then, is to help the programmer iteratively reduce the number of locations where coordination is used.

The conference paper of this work has also inspired follow-up work by others. In particular, Zinn et al. [Zinn et al. 2012] have generalized our results to also include a “partitioning policy”, which is a strategy to initialize every node of a network with input data before the computation starts. The basic idea is that each node is given local relations that provide information about how data is distributed, and in particular what data each node can autonomously reason about, i.e., without coordination with other nodes. This allows a node to sometimes perform nonmonotone operations without the need for communication. It even turns out that in some variations of the model considered, all database queries are “coordination-free”. It can be expected however, that such a partitioning policy is quite expensive in terms of how much additional data each node should have.

One of our results is that a monotone query can in principle be computed without coordination, but it remains open in what exact way the best performance can be achieved in a practical scenario. The work of Loo et al. [Loo et al. 2006] and Nigam et al. [Nigam et al. 2012], however, provides concrete algorithms for the case of distributed Datalog programs. They want to efficiently update the state (i.e., the derivations) on nodes of the network whenever some input facts change. It would be too costly to completely recompute the state of every node when an update happens. Instead, they propose a technique that propagates only the incremental changes that have to be distributedly applied. This allows for sending less messages around the network, and does not require needless recomputations of data. Their algorithms require no coordination, can handle recursive Datalog rules, and can tolerate messages that are delivered out of order by the network.

## 2. PRELIMINARIES

### 2.1. Basic Notions

We first recall some basic notions from database theory [Abiteboul et al. 1995]. 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 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  $R^{(k)}$ . We assume some infinite universe  $\mathbf{dom}$  of atomic data values. An arity of zero is also called *nullary*. We write  $()$  to denote the nullary tuple.

A fact  $\mathbf{f}$  is a pair  $(R, \bar{a})$ , often denoted as  $R(\bar{a})$ , where  $R$  is a relation name and  $\bar{a}$  is a tuple of values over  $\mathbf{dom}$ . 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  $adom(I) \subseteq \mathbf{dom}$ , is the set of atomic data values that occur in  $I$ .

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

We recall the following query languages [Abiteboul et al. 1995]:

- UCQ: unions of conjunctive queries,
- UCQ<sup>−</sup>: UCQ with negation in the body,
- FO: first order logic (relational calculus),
- While: FO with iteration,
- Datalog,
- NrDatalog: nonrecursive Datalog,

—  $\text{NrDatalog}^-$ : NrDatalog with negation on body atoms.

The weakest query language among these is UCQ.

## 2.2. Transducers

The computation on a single node of a network is formalized by means of relational transducers [Abiteboul et al. 2000; Ameloot et al. 2011; Deutsch et al. 2009; Deutsch et al. 2007; Deutsch et al. 2006; Spielmann 2003]. A *transducer schema*  $\Upsilon$  is a tuple  $\langle \Upsilon_{\text{in}}, \Upsilon_{\text{out}}, \Upsilon_{\text{msg}}, \Upsilon_{\text{mem}}, \Upsilon_{\text{sys}} \rangle$  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  $\text{Id}$  and  $\text{All}$ . A *transducer state* for  $\Upsilon$  is a database instance over  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}} \cup \Upsilon_{\text{sys}}$ . For a transducer state  $I$  we use notations of the form  $I|_{(\text{in}, \text{sys})}$  to denote the facts in  $I$  whose predicates are in  $\Upsilon_{\text{in}}$  or  $\Upsilon_{\text{sys}}$ . This notation is extended to other identifiers among in, out, msg, mem, and sys.

An (*epidemic*) *relational transducer*  $\Pi$  over  $\Upsilon$  is a collection of queries:

- 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  $\{R^{(k)}\}$ ;

where each of these queries has the input schema  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{msg}} \cup \Upsilon_{\text{mem}} \cup \Upsilon_{\text{sys}}$ . These queries will form the internal mechanism that a node uses to update its local storage and to send messages. The transducer model is parameterized by a generic query language  $\mathcal{L}$ : this language is used to concretely specify the above queries, in which case we call  $\Pi$  an  $\mathcal{L}$ -transducer. We will often abbreviate “epidemic relational transducer” simply as “transducer”. The term “epidemic” will become clear in Section 3.1, where the transducer model is used on a network.

Let  $\Pi$  be a transducer over schema  $\Upsilon$ . A *message instance* for  $\Upsilon$  is a database instance over  $\Upsilon_{\text{msg}}$ . 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  $\Upsilon$ ,  $I_{\text{rcv}}$  and  $J_{\text{snd}}$  are message instances for  $\Upsilon$  such that (denoting  $I' = I \cup I_{\text{rcv}}$ ):

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

where, following the presentation in [Zinn et al. 2012],

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

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 [Spielmann 2003]. 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  $J_{\text{snd}} = J'_{\text{snd}}$ .

### 3. TRANSDUCER NETWORKS

We now formalize a network of computing nodes. In Section 4 we give example programs.

A *network*  $\mathcal{N}$  is a finite, connected, and undirected graph whose nodes are all in **dom**. We write  $\text{nodes}(\mathcal{N})$  and  $\text{edges}(\mathcal{N})$  to denote the nodes and edges of  $\mathcal{N}$  respectively. For  $x \in \text{nodes}(\mathcal{N})$ , we write  $\text{neighbor}(x, \mathcal{N})$  to denote the set  $\{y \mid (x, y) \in \text{edges}(\mathcal{N})\}$ .

A (*homogeneous*) *transducer network* is a triple  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$  where  $\mathcal{N}$  is a network,  $\Upsilon$  is a transducer schema, and  $\Pi$  is a transducer over  $\Upsilon$  whose copies will be running at each node of the network. One could also consider non-homogeneous transducer networks, where each node contains a different transducer, possibly over a different schema, but these types of networks are not considered in the present paper. We discuss our design choices in Section 3.4.

For a query language  $\mathcal{L}$ , we say that a transducer network is an  $\mathcal{L}$ -transducer network if all its transducers are  $\mathcal{L}$ -transducers.

We now formalize how data is distributed across a network. A *distributed database instance over a network  $\mathcal{N}$  and a database schema  $\mathcal{D}$*  is a total function that assigns to each node of  $\mathcal{N}$  an ordinary database instance over  $\mathcal{D}$ .

#### 3.1. Operational Semantics

Let  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$  be a transducer network. Any distributed database instance over  $\mathcal{N}$  and  $\Upsilon_{\text{in}}$  can be given as input to  $\mathcal{T}$ . Let  $H$  be such an instance. A *configuration of  $\mathcal{T}$  on  $H$*  is a pair  $\rho = (s, b)$  of functions where

- $s$  maps each  $x \in \text{nodes}(\mathcal{N})$  to a transducer state  $J = s(x)$  such that  $J|_{(\text{in})} = H(x)$  and  $J|_{(\text{sys})} = \{\text{Id}(x)\} \cup \{\text{All}(y) \mid y \in \text{nodes}(\mathcal{N})\}$ ; and,
- $b$  maps each  $x \in \text{nodes}(\mathcal{N})$  to a finite multiset of facts over  $\Upsilon_{\text{msg}}$ .

We call  $s$  the *state function* and  $b$  the *buffer function*. Intuitively, the instance  $H$  is used to initialize each node, and for each  $x \in \text{nodes}(\mathcal{N})$  the system relations **Id** and **All** in  $\Upsilon_{\text{sys}}$  provide the local transducer at  $x$  the identity of the node  $x$  it is running on and the identities of the other nodes. Next, the buffer function maps each  $x \in \text{nodes}(\mathcal{N})$  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  $\mathcal{T}$  on  $H$* , denoted  $\text{start}(\mathcal{T}, H)$ , is the configuration  $\rho = (s, b)$  of  $\mathcal{T}$  on  $H$  that for each  $x \in \text{nodes}(\mathcal{N})$  defines  $s(x)|_{(\text{out}, \text{mem})} = \emptyset$  and  $b(x) = \emptyset$ .

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

- $m$  is a submultiset of  $b_1(x)$  and there exists a message instance  $J_{\text{snd}}$  such that

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

- is a local transition of transducer  $\Pi$ ;
- for each  $y \in \text{nodes}(\mathcal{N}) \setminus \{x\}$  we have  $s_1(y) = s_2(y)$ ;
- $b_2(x) = b_1(x) \setminus m$ ; for each  $y \in \text{neighbor}(x, \mathcal{N})$  we have  $b_2(y) = b_1(y) \cup J_{\text{snd}}$ ; and for all other nodes  $y$  we have  $b_2(y) = b_1(y)$ .

We call  $x$  the *active node* (or *recipient*) and  $\text{set}(m)$  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 to all its neighbors.<sup>1</sup> The node does not send messages to itself. 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  $\mathcal{T}$  on a distributed input database instance  $H$  is an *infinite* sequence of global transitions  $\rho_i \xrightarrow{x_i, m_i} \rho_{i+1}$  for  $i = 1, 2, 3, \dots$ , where  $\rho_1 = \text{start}(\mathcal{T}, H)$ , and the  $i^{\text{th}}$  transition with  $i \geq 2$  operates on the resulting configuration of the previous transition. It follows from the semantics of local transitions that when a node during one global transition changes its output or memory, then these changes are visible to that node only starting from the next global transition in which that node is active. Note also that 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 receive a message multiset and do local transitions at the same time. The reason for not including this kind of global transition is that it can be simulated by multiple sequential global transitions: this is done by letting the previously concurrent nodes become active in some arbitrary order, and in each of those single-node transitions, the active node just reads its own message buffer like in the concurrent transition. Because local transitions are deterministic, the nodes will update their state and send out messages in the same way as during the concurrent transition.

### 3.2. Fairness

In the literature on process models it is customary to require certain “fairness” conditions on runs [Francez 1986; Apt et al. 1988; Lamport 2000]. Let  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$  be a transducer network. In this paper, a run of  $\mathcal{T}$  on some input distributed database instance is called *fair* if (i) every node of  $\mathcal{N}$  is active in an infinite number of transitions and (ii) if for some node a fact occurs in its message buffer in an infinite number of configurations, then this fact is delivered to that node during an infinite number of transitions. Intuitively, the last condition demands that no sent messages are infinitely delayed. We consider only fair runs.

Note that every transducer network has a fair run for every input because heartbeat transitions are still possible even when the message buffers have become empty.

### 3.3. Message Delivery Constraints

We may want to impose a size-constraint on the delivered message multisets. Indeed, for a transducer network  $\mathcal{T}$  and a natural number  $k \geq 1$ , we can restrict our attention to runs of  $\mathcal{T}$  where the sizes of the delivered message *multisets* are of size at most  $k$ . This is the *k-delivery* semantics for  $\mathcal{T}$ . In a previous paper [Ameloot et al. 2011], we restricted attention to 1-delivery semantics. In this paper we assume no such bound, unless explicitly mentioned.

### 3.4. Discussion

In this section we want to motivate the usefulness of our transducer model by comparing it with the literature. First, some of the main characteristics of our model are that (i) the same transducer program is replicated at all nodes; (ii) a node can only send messages to its neighbors; and (iii) messages are never lost. These three characteristics occur commonly in declarative networking literature [Loo et al. 2009; Grumbach and Wang 2010; Nigam et al. 2012].

<sup>1</sup>Hence the name “epidemic” [Eugster et al. 2004].

The aspect where our model deviates, however, is that the sender of a message cannot address the message to a particular neighbor. But this can be simulated in our model by designating a specific component of each message as the addressee. When a node receives a message, it can verify that it is the valid addressee by using local relation  $\text{Id}$ , in which case it will process the message, and otherwise the node forwards the message to (all) its neighbors. This can be seen as less efficient: depending on the network topology, some nodes might receive many messages not addressed to them, and this could lead to some messages being forwarded forever. But we will not be concerned with such efficiency issues in this paper. Instead, we believe that sending to all neighbors is sufficient to reason about distributed algorithms. For completeness, however, in Section 7 we also briefly consider a transducer variant in which a node can send messages to a specific neighbor, and we relate this model to our epidemic model.

As a last remark, although the network topology is fixed at the start of a run, we note that our model can be used to simulate a dynamic network, in which nodes and edges can be temporarily offline. Indeed, offline nodes can be simulated by having some nodes not doing local transitions for a while, and offline edges can be simulated by delaying the messages sent through them. Our fairness condition however requires that eventually every node keeps on doing local transitions and that message delays are bounded.

#### 4. EXPRESSING QUERIES

What does it take for a transducer network to compute some global query? Here we propose a formal definition based on the two properties of *consistency* and *network-independence*. This is discussed in the following subsections.

##### 4.1. Transducer Kinds

We will use the following terminology for transducers. We call a transducer *oblivious* if it does not read the relations  $\text{Id}$  and  $\text{All}$  in its queries. Intuitively, this means that the transducer is unaware of the network context, because it does not know about the node it is running on, and it does not know about the other nodes. A transducer is called *inflationary* if it never deletes facts from its memory relations. That is, the deletion queries for the memory relations return the empty set of facts on all inputs. A transducer is called *monotone* if all its queries are monotone. The later Example 4.4 describes a transducer that is at the same time oblivious, inflationary, and monotone.

##### 4.2. Input and Output

Let  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$  be a transducer network. Let  $I$  be an ordinary (non-distributed) database instance over schema  $\Upsilon_{\text{in}}$ . This instance can be given as *input* to  $\mathcal{T}$  by partitioning it across the nodes, where the same fact can be given to multiple nodes. Formally, a distributed database instance  $H$  over  $\mathcal{N}$  and  $\Upsilon_{\text{in}}$  is said to be a *horizontal partition* of  $I$  if  $I = \bigcup_{x \in \text{nodes}(\mathcal{N})} H(x)$ .

Let  $\rho = (s, b)$  be a configuration of  $\mathcal{T}$  on input  $H$ . Naturally,  $\rho$  defines an output database instance  $\text{out}(\rho)$  over the schema  $\Upsilon_{\text{out}}$  as follows:

$$\text{out}(\rho) = \bigcup_{x \in \text{nodes}(\mathcal{N})} s(x)|_{(\text{out})}.$$

Let  $\mathcal{R}$  be a run of  $\mathcal{T}$  on some input. We denote the sequence of configurations of  $\mathcal{R}$  as  $\rho_1, \rho_2, \dots$ . If there is a number  $i \geq 1$  such that  $\text{out}(\rho_i) = \text{out}(\rho_j)$  for all  $j > i$ , then we call  $i$  a *quiescence point* for  $\mathcal{R}$ . We call a configuration  $\rho_i$  of  $\mathcal{R}$  a *quiescence configuration* if  $i$  is a quiescence point. Only quiescence configurations can follow a quiescence configuration, and all quiescence configurations define the same output database instance. Only a finite number of distinct output facts are possible because we only consider finite input instances, and

because the queries of transducers are generic, and hence domain-preserving. The following property is now clear:

**PROPOSITION 4.1.** *For every transducer network, on every input, every run contains a quiescence configuration.*

The *output* of run  $\mathcal{R}$  is now defined as  $out(\rho_i)$  where  $\rho_i$  is a quiescence configuration of  $\mathcal{R}$ . Our notion of output does not specify the output at individual nodes and does not prevent messages from being sent once a quiescence point is reached.

### 4.3. Consistency

We say that a transducer network  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$  is *consistent* if for all database instances  $I$  over  $\Upsilon_{in}$ , on all horizontal partitions of  $I$  over  $\mathcal{N}$ , all fair runs of  $\mathcal{T}$  have the same output, denoted  $\mathcal{T}(I)$ .

When  $\mathcal{T}$  is consistent, this function  $\mathcal{T}(\cdot)$  has the characteristic of a query, except that it need not be generic. For example, the “query” that asks for all data elements in the input that are not node identifiers, can be computed by a consistent transducer network. We mainly focus on the computation of generic queries. Naturally,  $\mathcal{T}$  is said to *compute a query*  $\mathcal{Q}$  over input schema  $\Upsilon_{in}$  and output schema  $\Upsilon_{out}$  if  $\mathcal{T}$  is consistent and  $\mathcal{T}(I) = \mathcal{Q}(I)$  for every database instance  $I$  over  $\Upsilon_{in}$  on which  $\mathcal{Q}$  is defined.

Because the individual queries that make up a transducer are generic, we can make the following observation:

**PROPOSITION 4.2.** *The function  $\mathcal{T}(\cdot)$  is generic for each consistent transducer network  $\mathcal{T}$  in which the transducer is oblivious.*

### 4.4. Examples

First, we explain our notational conventions for specifying concrete transducers. Because FO is equivalent to  $\text{NrDatalog}^-$  [Abiteboul et al. 1995], we will frequently use the more readable rule-based syntax of  $\text{NrDatalog}^-$  to specify FO-transducers. The answer relations of  $\text{NrDatalog}^-$  programs will be clearly marked. For example, for a memory-insertion query  $\mathcal{Q}_{ins}^R$ , the answer relation of the corresponding  $\text{NrDatalog}^-$  program is  $R_{ins}$ ; for an output query  $\mathcal{Q}_{out}^T$ , the answer relation is  $T_{out}$ ; for a message-sending query  $\mathcal{Q}_{snd}^S$ , the answer relation is  $S_{snd}$ . We leave a blank line between the  $\text{NrDatalog}^-$  rules that belong to different queries. Unmentioned queries of a transducer are assumed to always return the empty set of facts.

We now give some examples of transducer networks.

*Example 4.3.* For a simple example of a consistent transducer network, let the input be a binary relation  $R$ . Each node outputs the identical pairs from its part of the input. No messages are sent. This network computes the equality selection  $\sigma_{\$1=\$2}(R)$ . This is easily programmed on an FO-transducer, which is specified as follows. The transducer schema is  $\Upsilon_{in} = \{R^{(2)}\}$ ,  $\Upsilon_{out} = \{T^{(2)}\}$ ,  $\Upsilon_{msg} = \emptyset$ ,  $\Upsilon_{mem} = \emptyset$ , and the single transducer rule is:

$$T_{out}(\mathbf{u}, \mathbf{u}) \leftarrow R(\mathbf{u}, \mathbf{u}).$$

□

*Example 4.4.* To give an example of a consistent transducer network that involves communication, we compute the transitive closure of a binary relation  $R$  in a well-known distributed manner [Loo et al. 2009]. We present here a naive, unoptimized version. Each node sends its part of the input to its neighbors. Specifically, each node also forwards all messages it receives to its neighbors. This way, the input is flooded to all nodes. Each node accumulates the input facts it receives in a binary memory relation  $S$ . Finally, each node has an output relation  $T$  in which we repeatedly insert  $R \cup S \cup (T \circ T)$ , where  $\circ$  stands for relational

composition. Thanks to the monotonicity of the transitive closure, this transducer network is consistent. We can implement this idea with an UCQ-transducer. The transducer schema is  $\Upsilon_{\text{in}} = \{R^{(2)}\}$ ,  $\Upsilon_{\text{out}} = \{T^{(2)}\}$ ,  $\Upsilon_{\text{msg}} = \{U^{(2)}\}$ ,  $\Upsilon_{\text{mem}} = \{S^{(2)}\}$ , and the transducer rules are:

$$\begin{aligned} U_{\text{snd}}(\mathbf{u}, \mathbf{v}) &\leftarrow R(\mathbf{u}, \mathbf{v}). \\ U_{\text{snd}}(\mathbf{u}, \mathbf{v}) &\leftarrow U(\mathbf{u}, \mathbf{v}). \\ S_{\text{ins}}(\mathbf{u}, \mathbf{v}) &\leftarrow U(\mathbf{u}, \mathbf{v}). \\ T_{\text{out}}(\mathbf{u}, \mathbf{v}) &\leftarrow R(\mathbf{u}, \mathbf{v}). \\ T_{\text{out}}(\mathbf{u}, \mathbf{v}) &\leftarrow S(\mathbf{u}, \mathbf{v}). \\ T_{\text{out}}(\mathbf{u}, \mathbf{v}) &\leftarrow T(\mathbf{u}, \mathbf{w}), T(\mathbf{w}, \mathbf{v}). \end{aligned}$$

Note that the transducer is oblivious. There is no need to reason explicitly about node identifiers, because all we need is let the nodes steadily accumulate all input facts across the network and incrementally produce chunks of output. The transducer is also inflationary and monotone, reflecting the essential nature of the transitive closure computation.  $\square$

*Example 4.5.* Let us see a simple example of a transducer network that is *not* consistent. Consider a network with at least two nodes. The input is a unary relation  $R$ . Each node sends its part of  $R$  to its neighbors. Next, each node outputs the received messages on condition that the output is still empty. When there are at least two nodes and at least two different  $R$ -facts, different runs may deliver the messages in different orders, so different outputs can be produced, even for the same input distributed database instance. We can write an FO-transducer  $\Pi$  to implement this idea. The transducer schema  $\Upsilon$  is  $\Upsilon_{\text{in}} = \{R^{(1)}\}$ ,  $\Upsilon_{\text{out}} = \{T^{(1)}\}$ ,  $\Upsilon_{\text{msg}} = \{U^{(1)}\}$ ,  $\Upsilon_{\text{mem}} = \emptyset$ , and the transducer rules are:

$$\begin{aligned} U_{\text{snd}}(\mathbf{u}) &\leftarrow R(\mathbf{u}). \\ \text{block}() &\leftarrow T(\mathbf{u}). \\ T_{\text{out}}(\mathbf{u}) &\leftarrow \neg \text{block}(), U(\mathbf{u}). \end{aligned}$$

$\square$

Undecidability for testing consistency of a transducer network readily follows from undecidability of satisfiability of FO. The proof is in Appendix A.

#### 4.5. Network-Independence

We are mainly interested in the case where a query can be correctly computed by a transducer regardless of the network.

Let  $\Pi$  be a transducer over a schema  $\Upsilon$ . We say that  $\Pi$  is *network-independent* if for all networks  $\mathcal{N}$ , the transducer networks  $\langle \mathcal{N}, \Upsilon, \Pi \rangle$  are consistent and compute the same query  $\mathcal{Q}$ . We say that  $\mathcal{Q}$  is the query *distributedly computed* by  $\Pi$ . The transducer from Example 4.4 is network-independent. Now consider the following example.

*Example 4.6.* We give a simple example of a transducer that gives rise to consistent transducer networks but that is *not* network-independent. Suppose we have a unary input relation  $R$ . Each node sends its own identifier to its neighbors. This way the edges of the network can be discovered. The discovered edges are forwarded to every node, and when a node detects that the collected edges together form a complete graph, then the node outputs its local input for relation  $R$ . If the network is indeed a complete graph, by fairness

eventually all nodes will detect this, and then the transducer network computes the identity query. But on other network topologies the empty query is computed.

For completeness, we specify an FO-transducer  $\Pi$  to implement this idea. We define the transducer schema  $\Upsilon$  as  $\Upsilon_{\text{in}} = \{R^{(1)}\}$ ,  $\Upsilon_{\text{out}} = \{T^{(1)}\}$ ,  $\Upsilon_{\text{msg}} = \{A^{(1)}, B^{(2)}\}$ ,  $\Upsilon_{\text{mem}} = \{S^{(2)}\}$ . The rules of the transducer are:

$$A_{\text{snd}}(\mathbf{u}) \leftarrow \text{Id}(\mathbf{u}).$$

$$B_{\text{snd}}(\mathbf{u}, \mathbf{v}) \leftarrow A(\mathbf{u}), \text{Id}(\mathbf{v}).$$

$$B_{\text{snd}}(\mathbf{u}, \mathbf{v}) \leftarrow B(\mathbf{u}, \mathbf{v}).$$

$$S_{\text{ins}}(\mathbf{u}, \mathbf{v}) \leftarrow B(\mathbf{u}, \mathbf{v}).$$

$$\text{missing}() \leftarrow \text{All}(\mathbf{u}), \text{All}(\mathbf{v}), \mathbf{u} \neq \mathbf{v}, \neg S(\mathbf{u}, \mathbf{v}).$$

$$T_{\text{out}}(\mathbf{u}) \leftarrow R(\mathbf{u}), \neg \text{missing}().$$

□

Testing network-independence for FO-transducers is undecidable. See Appendix A for the proof.

#### 4.6. Preliminary Observations

We now give several preliminary results about expressing queries with transducers, that are important for deriving later results.

First, we present two lemmas which show that at each node a transducer can always assemble a local copy of all input facts available on the network.

**LEMMA 4.7.** *Let  $\mathcal{D}$  be a database schema. There is a transducer schema  $\Upsilon$  with  $\Upsilon_{\text{in}} = \mathcal{D}$  and an oblivious, inflationary, monotone UCQ-transducer  $\Pi$  over  $\Upsilon$  such that for every transducer network for  $\Pi$ , for every instance  $I$  of  $\mathcal{D}$ , on every horizontal partition of  $I$ , every fair run reaches a configuration where every node has a local copy of the entire instance  $I$  in its memory.*

**PROOF.** Because the construction is straightforward, we only provide a sketch. The idea is that all nodes will send out their local input facts and forward any message they receive. The local inputs, together with the received inputs, are accumulated in local memory relations. In any fair run, eventually all nodes will have received all input facts. Relations  $\text{Id}$  and  $\text{All}$  are not needed. We also do not need deletions on the memory relations. This technique has already been illustrated by Example 4.4. □

We can refine the technique of Lemma 4.7 to let a node know *when* it has collected every input fact in memory:

**LEMMA 4.8.** *Let  $\mathcal{D}$  be a database schema. There is a transducer schema  $\Upsilon$  with  $\Upsilon_{\text{in}} = \mathcal{D}$  and an UCQ<sup>-</sup>-transducer  $\Pi$  over  $\Upsilon$  such that for every transducer network for  $\Pi$ , for every instance  $I$  of  $\mathcal{D}$ , on every horizontal partition of  $I$ , every fair run reaches a configuration where every node has a local copy of the entire instance  $I$  in its memory, and an additional flag ‘ready’ is true (implemented by a nullary memory relation). Moreover, the flag ‘ready’ does not become true at a node before that node has the entire instance  $I$  in its memory.*

*The transducer  $\Pi$  is not oblivious, but can be made inflationary when using locally the language NrDatalog<sup>-</sup> instead of UCQ<sup>-</sup>.<sup>2</sup>*

<sup>2</sup>This is because a NrDatalog<sup>-</sup> program allows auxiliary relations to be declared, to which negation can be applied.

PROOF. We provide a sketch, and the full construction can be found in Appendix A.2. The idea is that a node  $x$  will send its local input facts over relation  $R^{(k)} \in \mathcal{D}$  to every other node, with an additional last component that contains the identifier of  $x$ , to indicate the origin of the fact. We call this last component the “tag”. Next, when a node  $y$  receives a tagged input fact, it removes the tag and stores the fact in its memory. This already lets each node incrementally accumulate all inputs across the network. Now, for each fact that  $y$  receives from  $x$ , node  $y$  also sends an acknowledgment back to  $x$ . The node  $x$  checks whether  $y$  has (eventually) acknowledged all the input facts of  $x$ . If yes, then  $x$  sends out  $\text{done}(x, y)$ . From the viewpoint of  $y$ , if  $y$  has received  $\text{done}(x, y)$  from all other nodes  $x$  then it knows that it has accumulated all the input facts on the network, and the **ready**-flag is created at  $y$ . The relations **Id** and **All** are used heavily in this protocol.  $\square$

The following theorem indicates that our transducer model has enough expressive power to study queries in the distributed context:

**THEOREM 4.9.** *Let  $\mathcal{L}$  be a language containing  $UCQ^\neg$ . Then every query expressible in  $\mathcal{L}$  can be distributedly computed by an  $\mathcal{L}$ -transducer. In particular, if  $\mathcal{L}$  is a computationally complete query language, every partial computable query can be distributedly computed by an  $\mathcal{L}$ -transducer.*

PROOF. Let  $Q$  be a query expressible in  $\mathcal{L}$ . Let  $\mathcal{D}$  and  $\mathcal{D}'$  be respectively the input and output schema of  $Q$ . We construct an  $\mathcal{L}$ -transducer  $\Pi$  to compute  $Q$  in two steps. In the first step, we use the partial specification of  $\Pi$  from Lemma 4.8 to obtain the entire input instance at every node. The language  $UCQ^\neg$  suffices for this. This transducer has input schema  $\Upsilon_{\text{in}} = \mathcal{D}$  but does not produce any output yet. In the second step, we define the output schema of this transducer to be  $\Upsilon_{\text{out}} = \mathcal{D}'$ . Now, because  $Q$  is expressible in  $\mathcal{L}$ , once the flag **ready** becomes true, we can output  $Q$  in the next local transition, by implementing for each output relation an  $\mathcal{L}$ -query that reads only the collected input facts.  $\square$

In the context of the CALM conjecture, monotone queries will play an important role. For now, we observe that oblivious transducers are sufficient to compute them:

**THEOREM 4.10.** *Let  $\mathcal{L}$  be a query language containing  $UCQ$ . Then every monotone query expressible in  $\mathcal{L}$  can be distributedly computed by an oblivious  $\mathcal{L}$ -transducer. In particular, if  $\mathcal{L}$  is computationally complete, every partial computable monotone query can be distributedly computed by an oblivious  $\mathcal{L}$ -transducer. Moreover, these oblivious transducers can be made inflationary and monotone.*

PROOF. Let  $Q$  be a monotone query expressible in  $\mathcal{L}$ . The idea is the same as in the proof of Theorem 4.9, but we now use the oblivious, inflationary, monotone transducer from Lemma 4.7, to let every node gradually collect all inputs facts available on the network. Now, because  $Q$  is expressible in  $\mathcal{L}$ , in every local transition we can execute  $\mathcal{L}$ -queries for the output relations that read the part of the input already accumulated in memory. Since  $Q$  is monotone, no incorrect tuples are output this way. Eventually, all nodes have accumulated all the input across the network, and no new outputs will be produced.  $\square$

## 5. THE CALM CONJECTURE

The following was conjectured by Hellerstein:

**CONJECTURE 5.1 (CALM CONJECTURE [HELLERSTEIN 2010B]).** *A program has an eventually consistent, coordination-free execution strategy if and only if it is expressible in (monotonic) Datalog.*

Before we can rigorously investigate this conjecture, we want to formalize the notion of “coordination-freeness”. This is presented in Section 5.1. Next, we will present our formal

CALM conjecture and its associated results in Section 5.2. Additional results are in Section 5.3.

### 5.1. Coordination-free

The CALM conjecture hinges on an intuitive notion of “coordination” of certain distributed computations. We illustrate this notion with a few examples.

In the well-known two-phase commit protocol [Gray 1978], each node is responsible for executing some part of a distributed transaction. To keep the distributed database consistent in the face of runtime crashes, either all parts should be committed or none is. To this purpose, after executing its part of the distributed transaction, but before actually committing the results, a node checks that *every* node can commit its results. This way, the distributed commit can proceed only if all individual nodes can commit. Naturally, the nodes have to exchange messages to determine if they can commit or not.

As another example, the multicast protocol of Lemma 4.8 also relies on heavy coordination: the nodes exchange many messages, including acknowledgments, before they all obtain the flag ‘ready’.

Generalizing both examples, the main idea behind coordination is that a large set of nodes needs to obtain a consensus. For two-phase commit this is the global decision whether all nodes should commit or not, and for Lemma 4.8 the consensus is that all nodes have the same data. Reaching a consensus is known to be difficult in the distributed context [Attiya and Welch 2004]. Because of the complexity of consensus, the involved nodes sometimes have to wait relatively long before they can continue with the actual computation. This is called a “global barrier” [Hellerstein 2010b].

It should be clear that coordination typically decreases the efficiency of distributed computations, because while the coordination is under way, the nodes are just waiting. So, it seems useful to understand precisely when coordination can be avoided, for which we will use the term “coordination-freeness”. This is what the CALM conjecture is all about. It seems hard to give a definitive formalization of coordination-freeness. Still, we offer here a nontrivial definition that appears interesting. A very drastic, too drastic, definition of coordination-free would be to disallow any communication. Our definition is much less severe and only requires that the computation can succeed without communication on “suitable” horizontal partitions. It actually does not matter what a suitable partition is, as long as it exists. Even under this liberal definition, the only-if direction of (our formalization of) the CALM conjecture will turn out to hold.<sup>3</sup>

Formally, let  $\Pi$  be a transducer over a schema  $\Upsilon$ . Let  $\mathcal{T}$  be a transducer network for  $\Pi$ . We call  $\mathcal{T}$  *coordination-free* if for every database instance  $I$  over  $\Upsilon_{\text{in}}$ , there exists a horizontal partition  $H$  of  $I$  and a run of  $\mathcal{T}$  on  $H$  in which a quiescence configuration is already reached by performing only heartbeat transitions (zero or more). Intuitively, if the horizontal partition is right, then no communication is required to correctly compute the query. The property of coordination-freeness is mainly interesting for consistent transducer networks, because then at the quiescence configuration that was reached with only heartbeat transitions, the produced output is the same as produced by any other fair run. We call transducer  $\Pi$  *coordination-free* if for every network its corresponding transducer network is coordination-free.

*Example 5.2.* Consider again the transitive closure computation from Example 4.4. When every node already has the full input, they can each individually compute the transitive closure with only heartbeats. Hence, this transducer is coordination-free.  $\square$

<sup>3</sup>Of course, under the drastic definition of coordination-freeness, the if-direction of the CALM conjecture (which is the easy direction) as formulated below in Proposition 5.3, will no longer hold.

The transitive closure query is monotone and this example can actually be generalized in the following proposition. This proposition is implicit in the literature on embarrassingly parallel computation [Hellerstein 2010b; Loo et al. 2006; Nigam et al. 2012], and our main result (Theorem 5.11) will provide a converse to it.

**PROPOSITION 5.3.** *Let  $\mathcal{L}$  be a query language containing UCQ. Every monotone query  $Q$  expressible in  $\mathcal{L}$  can be distributedly computed by a coordination-free  $\mathcal{L}$ -transducer.*

**PROOF.** Recall from the proof of Theorem 4.10 that there is an oblivious  $\mathcal{L}$ -transducer that distributedly computes  $Q$ . Using the same intuition as in Example 5.2, this transducer is coordination-free.  $\square$

The reader should not be lulled into believing that with a coordination-free program it is always sufficient to give the full input at all nodes, as the following example shows:

*Example 5.4.* Consider the following query  $Q$ , having as input two nullary relations  $A$  and  $B$ , and a nullary output relation  $T$ : create the non-empty output (representing “true”) if at least one of  $A$  and  $B$  is nonempty. This query is monotone. Consider the following (contrived) transducer  $\Pi$  to compute  $Q$ . If the network has only one node (which can be tested by looking at the relation  $\text{All}$ ), the transducer simply outputs the answer to the query. Otherwise, it first tests if its local input fragments of  $A$  and  $B$  are *both* nonempty. If this is the case, nothing is output locally yet, but a nullary fact  $C$  is sent out. Any node that receives the message  $C$  will output it. When precisely one of  $A$  and  $B$  is nonempty locally, the transducer simply outputs the correct output directly. The transducer is network-independent. Also, the transducer is coordination-free, because on networks with at least two nodes there always is a partition of the data under which no node has both  $A$  and  $B$  locally nonempty, and the query can be computed without communication. Moreover, when  $A$  and  $B$  are both nonempty, and every node has the entire input, no run will reach a quiescence configuration without communication.  $\square$

The following two examples show that network-independence for a transducer does not guarantee coordination-freeness, and vice versa.

*Example 5.5.* We provide an example of a transducer that is network-independent but not coordination-free, i.e., requires communication. Let  $Q$  be the following “emptiness” query, having a nullary input relation  $R$ , and a nullary output relation  $T$ : create the non-empty output (representing “true”) iff  $R$  is empty. This query is nonmonotone. We now describe a transducer to distributedly compute  $Q$ . Since every node can have a part of the input, the nodes coordinate with each other to be certain that  $R$  is empty at every node. Every node sends out its identifier (using the relation  $\text{Id}$ ) on condition that its local relation  $R$  is empty. Received messages are forwarded, so that if  $R$  is globally empty, eventually all nodes will have received the identifiers of all nodes, which can be checked using the relation  $\text{All}$ . When this happens, the transducer at each node outputs a nullary fact.

For completeness, we specify an FO-transducer  $\Pi$  to implement this idea. The transducer schema  $\Upsilon$  is as follows:  $\Upsilon_{\text{in}} = \{R^{(0)}\}$ ,  $\Upsilon_{\text{out}} = \{T^{(0)}\}$ ,  $\Upsilon_{\text{msg}} = \{U^{(1)}\}$  and  $\Upsilon_{\text{mem}} = \{S^{(0)}\}$ . The rules are:

$$U_{\text{snd}}(\mathbf{u}) \leftarrow \text{Id}(\mathbf{u}), \neg R().$$

$$U_{\text{snd}}(\mathbf{u}) \leftarrow U(\mathbf{u}).$$

$$S_{\text{ins}}(\mathbf{u}) \leftarrow \text{Id}(\mathbf{u}), \neg R().$$

$$S_{\text{ins}}(\mathbf{u}) \leftarrow U(\mathbf{u}).$$

$$\text{missing}() \leftarrow \text{All}(\mathbf{u}), \neg S(\mathbf{u}).$$

$$T_{\text{out}}() \leftarrow \neg \text{missing}().$$

□

*Example 5.6.* We give a transducer that is coordination-free, and that is consistent on every network, but is not network-independent. The transducer has two unary input relations  $R$  and  $S$ , and it has a unary output relation  $T$ . Using relations  $\text{Id}$  and  $\text{All}$ , the transducer can detect if there is only one node, or if there are more nodes. If there is just one node, the single node outputs the union of  $R$  and  $S$ . If there are at least two nodes, then all nodes will copy their local inputs into their memory; they also broadcast their input facts to each other, so that all nodes accumulate all inputs of the network; and, the nodes will continuously output the intersection of the accumulated  $R$ -facts with the accumulated  $S$ -facts.

First, we see that on each network this transducer is consistent. Indeed, on a single-node network the union of  $R$  and  $S$  is output, and on a multi-node network the intersection of  $R$  and  $S$  is output. This different output behavior prevents the transducer from being network-independent. Finally, the transducer is coordination-free because on a single-node network the output is always computed with only heartbeats, and on a multi-node network we can consider the partition where each node has the entire input, and then the intersection of  $R$  and  $S$  can already be computed with only heartbeats. □

Coordination-freeness seems a useful property for a transducer to have. However, it cannot be decided automatically in general:

**PROPOSITION 5.7.** *Coordination-freeness is undecidable for FO-transducers.*

**PROOF.** We reduce the finite satisfiability problem for FO to deciding coordination-freeness for FO-transducers. Let  $\varphi$  be an FO-sentence over a database schema  $\mathcal{D}$ . We construct an FO-transducer  $\Pi$  that is coordination-free iff  $\varphi$  is *not* finitely satisfiable.

Consider the transducer  $\Pi$  in Example 5.5, that is over schema  $\Upsilon$ . We may assume without loss of generality that the relation names of  $\Upsilon$  do not occur in  $\mathcal{D}$ . We obtain a new transducer schema  $\Upsilon'$  from  $\Upsilon$  by adding  $\mathcal{D}$  to  $\Upsilon_{\text{in}}$ ; by adding new message relations  $\{(C^{\text{msg}}, k) \mid C^{(k)} \in \mathcal{D}\}$ ; and, by adding new memory relations  $\{(C^{\text{mem}}, k) \mid C^{(k)} \in \mathcal{D}\}$ . We obtain a new transducer  $\Pi'$  over  $\Upsilon'$  by modifying  $\Pi$  to let all nodes gradually accumulate all input facts by means of message forwarding. Moreover, besides keeping the old output condition “ $\neg \text{missing}$ ”, we will only produce an output if additionally  $\varphi$  is satisfied on the accumulated  $\mathcal{D}$ -facts so far (in memory).

For the first direction, suppose that  $\varphi$  is finitely satisfiable on a database instance  $I$  over  $\mathcal{D}$ . We show that  $\Pi'$  is not coordination-free. We can regard  $I$  as a database instance over  $\Upsilon'_{\text{in}}$ , where relation  $R$  is empty. Let  $\mathcal{N}$  be a network containing two nodes  $x$  and  $y$ . Let  $\mathcal{T}$  denote the transducer network based on  $\mathcal{N}$  and  $\Pi'$ . Suppose that there is some horizontal partition  $H$  of  $I$  over  $\mathcal{N}$  and a run  $\mathcal{R}$  of  $\mathcal{T}$  on input  $H$  in which a first quiescence configuration is already reached by doing only heartbeat transitions. Because  $I$  does not contain  $R()$ , the nodes send the messages  $U(x)$  and  $U(y)$ . Because of fairness, these messages must be delivered to  $y$  and  $x$  respectively, which can happen only after the first quiescence point

because before the quiescence point there are only heartbeat transitions. Eventually, every node will find  $\neg\text{missing}()$  to be true. The same reasoning can be applied to the relations of  $\mathcal{D}$ : whether  $I$  is empty or not, there must be a configuration after the first quiescence point, in which all nodes have accumulated  $I$  in the memory relations. Then  $\varphi$  also becomes true, and thus we know that every node eventually outputs  $T()$ . Note that this fact cannot be in the first quiescence configuration because it requires the delivery of at least one of the messages  $U(x)$  or  $U(y)$ . So, the initial quiescence configuration that was reachable by only heartbeat transitions cannot exist. Thus, the network  $\mathcal{N}$  and input  $I$  are a proof that  $\Pi'$  is not coordination-free.

For the other direction, suppose that  $\varphi$  is not finitely satisfiable. Then no transducer network based on  $\Pi'$  can produce output, no matter what the input instance over  $\Upsilon'_{\text{in}}$  or horizontal partition of that instance is. Hence, the start configuration of every run is already a quiescence configuration, and  $\Pi'$  is coordination-free.  $\square$

Although coordination-freeness is undecidable for FO-transducers (and by extension more powerful transducers), we can identify a syntactic class of transducers that is guaranteed to be coordination-free, and that will prove to have the same expressive power as the class of coordination-free transducers. Importantly, the syntactic restriction does not guarantee network-independence. Recall from Section 4.1 that an *oblivious* transducer does not read the system relations  $\text{Id}$  and  $\text{A11}$ . For now we observe:

**PROPOSITION 5.8.** *Let  $\mathcal{L}$  be a query language. Every network-independent, oblivious  $\mathcal{L}$ -transducer is coordination-free.*

**PROOF.** Let  $\Pi$  be a network-independent, oblivious  $\mathcal{L}$ -transducer over a schema  $\Upsilon$ . Let  $\mathcal{Q}$  be the query distributedly computed by  $\Pi$ .

First, on a single-node network, the single node is always given the entire input and there can only be heartbeat transitions. Then, for an input instance  $I$  over  $\Upsilon_{\text{in}}$ , a quiescence configuration containing  $\mathcal{Q}(I)$  is always reached by doing only heartbeat transitions.

Now consider any other network  $\mathcal{N}$ , any instance  $I$  over  $\Upsilon_{\text{in}}$ , and the horizontal partition  $H$  that places the entire instance  $I$  at every node. Since  $\Pi$  is oblivious, nodes cannot detect that they are on a network with multiple nodes unless they receive a message. So, by doing only heartbeat transitions initially, every node will act the same as if in a single-node network and will already output the entire  $\mathcal{Q}(I)$ . Because  $\Pi$  is network-independent, the nodes cannot output more than  $\mathcal{Q}(I)$  when they receive messages afterwards.  $\square$

## 5.2. Main Results

Now we can formalize the original Conjecture 5.1. We will take the terms “program” and “to have an execution strategy” to mean “query” and “to be distributedly computed by a transducer”, respectively. The term “eventually consistent” is then formalized by our notions of consistency and network-independence. Under this interpretation, the conjecture becomes:

**CONJECTURE 5.9.** *A query can be distributedly computed by a coordination-free transducer if and only if it is expressible in Datalog.*

Let us immediately get the if-side of this conjecture out of the way. It holds, because a query in Datalog is monotone, and then by Theorem 4.10 there exists an oblivious transducer to compute the query, but we have seen in Proposition 5.8 that oblivious transducers are coordination-free.

As to the only-if side, the explicit mention of Datalog is a bit of a nuisance because Datalog is limited to polynomial time whereas there certainly are monotone queries outside PTIME. We also mention the celebrated paper [Afrati et al. 1995] where Afrati, Cosmadakis

and Yannakakis show that even within PTIME there exist queries that are monotone but not expressible in Datalog.

But Datalog aside, however, the true emphasis of the CALM Conjecture clearly lies in the monotonicity aspect. Indeed, we confirm it in this sense:

**THEOREM 5.10.** *Let  $\mathcal{L}$  be a query language. Every query that is distributedly computed by a coordination-free  $\mathcal{L}$ -transducer is monotone.*

**PROOF.** Let  $\Pi$  be a coordination-free  $\mathcal{L}$ -transducer over a schema  $\Upsilon$  that distributedly computes a query  $\mathcal{Q}$ . Let  $I$  and  $J$  be two database instances over the schema  $\Upsilon_{\text{in}}$  such that  $I \subseteq J$ . We must show that  $\mathcal{Q}(I) \subseteq \mathcal{Q}(J)$ . Consider a fact  $\mathbf{f} \in \mathcal{Q}(I)$ . Consider a network  $\mathcal{N}$  with at least two nodes. Let  $\mathcal{T}$  denote the transducer network based on  $\mathcal{N}$  and  $\Pi$ . Since  $\Pi$  is coordination-free and network-independent, there exists a horizontal partition  $H$  of  $I$  and a run  $\mathcal{R}$  of  $\mathcal{T}$  on input  $H$  in which a quiescence configuration, containing the facts  $\mathcal{Q}(I)$ , is already reached by letting the nodes do only heartbeat transitions. Let  $x$  be a node where  $\mathbf{f}$  is output in the quiescence configuration. Let  $y$  be a node different from  $x$  and consider a horizontal partition  $H'$  of  $J$  where  $H'(x) = H(x)$  and  $H'(y) = H(y) \cup (J \setminus I)$ . Let  $n$  be the number of initial heartbeat transitions with recipient  $x$  in run  $\mathcal{R}$  that were needed to output  $\mathbf{f}$  at  $x$ . Consider a prefix of a run of  $\mathcal{T}$  on input  $H'$  in which we initially do  $n$  heartbeat transitions, all with active node  $x$ . Because local transitions are deterministic, the node  $x$  goes through the same state changes as in run  $\mathcal{R}$  before  $\mathbf{f}$  is output and therefore  $\mathbf{f}$  is output again in this prefix. The prefix can be extended to a full fair run  $\mathcal{R}'$  of  $\mathcal{T}$  on input  $H'$ . Since  $\mathcal{T}$  is consistent, the fact  $\mathbf{f}$  will be output on any partition of  $J$ , during any fair run. Hence,  $\mathbf{f}$  belongs to the query computed by  $\mathcal{T}$  applied to  $J$ . Moreover,  $\Pi$  is network-independent, so  $\mathbf{f}$  belongs to  $\mathcal{Q}(J)$ .  $\square$

We can now obtain the following result:

**THEOREM 5.11.** *Let  $\mathcal{L}$  be a query language containing UCQ. For every query  $\mathcal{Q}$  that is expressible in  $\mathcal{L}$ , the following are equivalent:*

- (1)  $\mathcal{Q}$  can be distributedly computed by a coordination-free  $\mathcal{L}$ -transducer;
- (2)  $\mathcal{Q}$  can be distributedly computed by an oblivious  $\mathcal{L}$ -transducer; and,
- (3)  $\mathcal{Q}$  is monotone.

**PROOF.** Theorem 4.10 yields (3)  $\Rightarrow$  (2); Proposition 5.8 yields (2)  $\Rightarrow$  (1); Theorem 5.10 yields (1)  $\Rightarrow$  (3).  $\square$

In particular, if  $\mathcal{L}$  is computationally complete, then the previous equivalences hold for any computable query. As a small remark, now it is of no surprise that Example 5.5 required coordination; indeed, there we distributedly compute a non-monotone query.

**5.2.1. Discussion.** Theorem 5.11 can be used as follows in practice. Essentially, by restricting a language, its execution can in general be optimized more thoroughly than the unrestricted language. A well-known example is SQL versus a Turing-complete programming language. For our situation, the programmer of a distributed (query) algorithm can write a program in a high-level declarative formalism, like the transducer model presented in this paper, or a Datalog-variant like e.g. [Loo et al. 2006; Alvaro et al. 2011b; Abiteboul et al. 2011]. Suppose that the query is monotone. Then we know by Theorem 5.11 that it can be implemented in a coordination-free manner. Moreover, we can prevent the programmer from abusing coordination using the syntactic restriction of obliviousness. The main idea is that the programmer is given only a few communication primitives, like sending a message to its neighbors, and a syntactic restriction is imposed to prevent the programmer from using network relations like `Id` or `All` (or equivalent information). Next, the programmer, or a software tool, needs to assert that the program is network-independent, i.e., on every network, all fair runs produce the desired outcome. Then, using Theorem 5.11, if the runtime

is told that the program is oblivious and network-independent, the runtime can execute the program without any coordination. By contrast, if the programmer uses `Id` and `All`, then this semantic property is no longer guaranteed, and one would have to resort to a general execution strategy that has built-in coordination, which seems a waste if the program expresses a monotone computation. This way, obliviousness could be a useful guiding principle for distributed query evaluation. The works of Loo et al. [Loo et al. 2006] and Nigam et al. [Nigam et al. 2012] provide coordination-free distributed execution engines for Datalog.

### 5.3. Further Results

It is natural to wonder about variations of our model. One question may be about the system relations `Id` and `All`. Without them (the oblivious case), we know that we are always coordination-free and thus monotone.

What if we would read precisely one system relation; only `Id` or only `All`? As to coordination-freeness, the argument given in the proof of Proposition 5.8 still works when the transducer reads only `Id`, because then nodes still cannot detect that they are on a network with multiple nodes. However, the argument fails when the transducer reads only `All`, and indeed we have the following counterexample.

*Example 5.12.* We describe a transducer that is network-independent, reads only `All`, but that is not coordination-free. The query expressed is simply the identity query on a unary relation  $R$ . The transducer can observe the difference between a single-node and a multi-node network by looking at the relation `All`. If it is a single-node network, the node simply outputs the result directly. If it is a multi-node network, every node sends out a message. Only upon receiving a message will a node then output the result. Thus on a multi-node network, regardless of the horizontal partition, communication is needed for the transducer network to produce the required output.  $\square$

So, coordination-freeness is not guaranteed when reading only `All`, but yet, monotonicity is not lost.

**THEOREM 5.13.** *Let  $\mathcal{L}$  be a query language. Every query distributedly computed by an  $\mathcal{L}$ -transducer that reads only relation `All`, is monotone.*

**PROOF.** Let  $\Pi$  be a network-independent transducer that reads only `All`. As a technical convenience, we assume that runs can use *concurrent* global transitions, in which multiple nodes can be active at the same time, each receiving messages from their own message buffer. At the end of such a concurrent global transition, for each node, its message buffer is extended with the multiset union of all messages sent to it by its neighbors. These concurrent transitions can be simulated by a sequence of ordinary single-node transitions, as remarked at the end of Section 3.1.

Let  $\Upsilon$  be the schema of  $\Pi$ . Let  $\mathcal{Q}$  be the query distributedly computed by  $\Pi$ . Let  $I$  and  $J$  be two database instances over  $\Upsilon_{\text{in}}$  such that  $I \subseteq J$ . Let  $\mathbf{f} \in \mathcal{Q}(I)$ . We have to show that  $\mathbf{f} \in \mathcal{Q}(J)$ . The main trick used in this proof is that although  $\Pi$  can count the number of nodes of a network (using relation `All`), it cannot directly observe the *edges* of the network. So, when  $\mathbf{f}$  is output on input  $I$  in one network, we can fool the transducer to output  $\mathbf{f}$  on input  $J$  in another network that has only slightly different edges.

*Run on  $I$ .* Consider a network  $\mathcal{N}_1$  in the form of a ring, containing at least four nodes. See Figure 1 for an example. Let  $\mathcal{T}$  denote the transducer network based on  $\mathcal{N}_1$  and  $\Pi$ . Let  $H_1$  be the horizontal partition of  $I$  that places  $I$  on every node of  $\mathcal{N}_1$ .

We show now that there exists a run  $\mathcal{R}_1$  of  $\mathcal{T}$  on input  $H_1$  with sequence of configurations  $\rho_1 = (s_1, b_1)$ ,  $\rho_2 = (s_2, b_2)$ ,  $\dots$ , such that for each  $i \geq 1$  and each  $x, y \in \text{nodes}(\mathcal{N}_1)$  we have  $s_i(x) = s_i(y)$  and  $b_i(x) = b_i(y)$ . In words: in every configuration, all nodes have the same transducer state and the same message buffer. We inductively construct  $\mathcal{R}_1$ . For the base

case ( $i = 1$ ), configuration  $\rho_1$  satisfies the property because it is the start configuration: all nodes are given the entire input  $I$ , and all message buffers are empty. For the induction hypothesis, assume that the property holds for  $i$ . For the inductive step, we show how to continue the partially constructed run  $\mathcal{R}_1$  so that the property holds for  $i + 1$ . Denote  $m = b_i(x)$  for some node  $x$ . Possibly  $m = \emptyset$ . We next do a concurrent global transition in which *each* node is the recipient of delivered message multiset  $m$ . This is possible by using the induction hypothesis. So, we are delivering the entire message buffer at once to each node. Again by the induction hypothesis, all nodes have the same state in configuration  $\rho_i$ , and since local transitions are deterministic, all nodes will have the same state in configuration  $\rho_{i+1}$ . Also, if one node sends a message set  $J_{\text{snd}}$  on delivery of  $m$ , then all nodes will send this set on delivery of  $m$ . Hence, because  $\mathcal{N}_1$  is a ring, for each node, the messages of  $J_{\text{snd}}$  will have been added *twice* to its message buffer at the end of the concurrent global transition. Since all nodes emptied their message buffer at the beginning of the concurrent transition, we see that in  $\rho_{i+1}$  the nodes again have the same message buffer.

The run  $\mathcal{R}_1$  can be converted to a fair run  $\mathcal{R}'_1$  with only non-concurrent global transitions and that produces the same output as  $\mathcal{R}_1$ . Moreover, because  $\Pi$  is network-independent, we know that  $\mathcal{R}'_1$  outputs  $\mathcal{Q}(I)$ , and thus  $\mathcal{R}_1$  outputs  $\mathcal{Q}(I)$ . Therefore, we can consider a node  $u$  of  $\mathcal{N}_1$  and an index  $k \geq 1$  such that  $u$  outputs  $\mathbf{f}$  during the  $k^{\text{th}}$  concurrent global transition of  $\mathcal{R}_1$ .

*Run on  $J$ .* Let  $u$  be the node as previously defined. Let  $z$  be a node of  $\mathcal{N}_1$  that is not a neighbor of  $u$ . We obtain a new network  $\mathcal{N}_2$  from  $\mathcal{N}_1$  by adding an edge between the two neighbors of  $z$ . Because  $\mathcal{N}_1$  is a ring with at least four nodes, we know that this edge was not previously there and thus  $\mathcal{N}_2$  contains a smaller ring without node  $z$ . Let  $\mathcal{T}'$  denote the transducer network based on  $\mathcal{N}_2$  and  $\Pi$ . Importantly, note that  $\mathcal{N}_1$  and  $\mathcal{N}_2$  have precisely the same nodes. Let  $H_2$  be the horizontal partition of  $J$  that places  $I$  on every node except  $z$  and that places  $J \setminus I$  on  $z$ .

Let us abbreviate  $N = \text{nodes}(\mathcal{N}_2) \setminus \{z\}$ . Recall the sequence of configurations  $\rho_1, \rho_2, \dots$ , of run  $\mathcal{R}_1$  from above. We now show that there exists an (unfair) run  $\mathcal{R}_2$  of  $\mathcal{T}'$  on input  $H_2$  with sequence of configurations  $\rho'_1 = (s'_1, b'_1), \rho'_2 = (s'_2, b'_2), \dots$ , such that for each  $i \geq 1$  and each  $y \in N$  we have  $s'_i(y) = s_i(y)$  and  $b'_i(y) = b_i(y)$ . In words: the smaller ring of nodes  $N$  follows exactly the states and message buffers of run  $\mathcal{R}_1$ . We inductively construct  $\mathcal{R}_2$ . For the base case ( $i = 1$ ), the property is satisfied because input partition  $H_2$  initializes the nodes of  $N$  in the same way as input partition  $H_1$ . For the induction hypothesis, we assume that the property holds for index  $i$ . For the inductive step, we show that the property holds for index  $i + 1$ . As in the construction of  $\mathcal{R}_1$ , we next do a concurrent global transition in which we deliver to every node of  $N$  the contents of its entire message buffer. Using the induction hypothesis, this causes each node of  $N$  to send the same message instance  $J_{\text{snd}}$  to their neighbors. This message instance was also sent during the corresponding global transition of  $\mathcal{R}_1$ . Let  $y_1$  and  $y_2$  denote the two neighbors of node  $z$  in  $\mathcal{N}_1$ . We have  $\{y_1, y_2\} \subseteq N$ . Because we have added the extra edge between  $y_1$  and  $y_2$  in  $\mathcal{N}_2$ , node  $y_1$  sends  $J_{\text{snd}}$  to  $z$  and to  $y_2$ . This is similar for  $y_2$ . Node  $z$  does not send anything because it is ignored. So, like in  $\mathcal{R}_1$ , both  $y_1$  and  $y_2$  have  $J_{\text{snd}}$  added precisely twice to their message buffer at the end of the concurrent global transition. The rest of the reasoning is the same as in the inductive step for constructing  $\mathcal{R}_1$ . We obtain that the nodes of  $N$  have the same state and message buffers in configuration  $\rho'_{i+1}$  as in configuration  $\rho_{i+1}$ .

Consider again the run  $\mathcal{R}_2$ . Because  $u \in N$ , the fact  $\mathbf{f}$  is eventually output at  $u$  during  $\mathcal{R}_2$ , during some global transition  $k$ . But  $\mathcal{R}_2$  is clearly not fair because the node  $z$  is ignored. However, we can make a new run  $\mathcal{R}'_2$  by copying only the first  $k$  global transitions of  $\mathcal{R}_2$ , converting each of them to a sequence of ordinary (non-concurrent) global transitions and then extending this prefix arbitrarily to a full fair run. Thus, we obtain that  $\mathbf{f}$  is output in

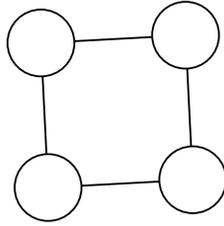


Fig. 1. Ring network topology

a fair run of  $\mathcal{T}'$  on input  $H_2$ . Since  $\Pi$  is network-independent, we obtain that  $\mathbf{f} \in \mathcal{Q}(J)$ , as desired.  $\square$

As a corollary, we can add two more statements to the three equivalent statements of Theorem 5.11:

**COROLLARY 5.14.** *Let  $\mathcal{L}$  be a query language containing UCQ. The following statements are equivalent for any query  $\mathcal{Q}$  expressible in  $\mathcal{L}$ :*

- (1)  $\mathcal{Q}$  can be distributedly computed by an oblivious  $\mathcal{L}$ -transducer;
- (2)  $\mathcal{Q}$  can be distributedly computed by an  $\mathcal{L}$ -transducer that is given only  $\text{Id}$ ; and,
- (3)  $\mathcal{Q}$  can be distributedly computed by an  $\mathcal{L}$ -transducer that is given only  $\text{All}$ .

**PROOF.** The directions (1)  $\Rightarrow$  (2) and (1)  $\Rightarrow$  (3) are immediate because an oblivious transducer is given neither of  $\text{Id}$  or  $\text{All}$ . For (2)  $\Rightarrow$  (1), when only  $\text{Id}$  is read, the query  $\mathcal{Q}$  is monotone as argued above. Then, by also using that  $\mathcal{Q}$  is expressible in  $\mathcal{L}$ , we can apply Theorem 4.10 to know that  $\mathcal{Q}$  is computable by an oblivious  $\mathcal{L}$ -transducer. The direction (3)  $\Rightarrow$  (1) is similar, but this time Theorem 5.13 is used.  $\square$

To conclude this section, we note that distributed algorithms involving a form of coordination typically require the participating nodes to have some knowledge about the other participating nodes [Attiya and Welch 2004]. This justifies our modeling of this knowledge in the form of the system relations  $\text{Id}$  and  $\text{All}$ . Importantly, we have shown that these relations are only necessary if one wants to compute a nonmonotone query in a distributed fashion.

## 6. EXPRESSIVENESS ANALYSIS

In this section we want to better understand the transducer model itself. The main question we would like to address is how the transducer model can be combined with a local query language to express a global query. It is not obvious what peculiarities of the model can be exploited in the local queries, and how. It will turn out actually that the global query language expressed by the transducer is the while-closure of its local query language. Intuitively, this is because each node can do multiple local transitions in a run, which can be seen as iterations of an implicit while-loop. This is very natural, and we believe this shows that our (distributed) transducer model is relatively elegant, because it respects previous results about well-known query languages [Abiteboul et al. 1995].

Table I summarizes the expressiveness results.

### 6.1. While versus FO

We first show the following property, and although the result might not sound very surprising, writing out the details turned out to be rather intricate.

**LEMMA 6.1.** *A query is expressible in While if and only if it is computable by an FO-transducer on a single-node network.*

Table I. Expressiveness Summary

Queries expressible in While = queries computable by FO-transducers = queries computable by $UCQ^-$ -transducers
Monotone queries expressible in While = queries computable by oblivious FO-transducers
Queries expressible in Datalog = queries computable by inflationary NrDatalog-transducers
Queries in PSPACE = queries computable by multi-node FO-transducer networks under 1-delivery semantics

PROOF (SKETCH). For the only-if direction, we have to simulate a While-program on a single-node FO-transducer network. A While-program can be simulated by iterated heartbeats using well-known techniques [Abiteboul and Vianu 1991a]. The main idea is that the loops in the while-program are rewritten with explicit “goto” statements. The statements of this rewritten program can then be simulated by an FO-transducer that keeps track of which statement is to be executed next, and goto-statements can make the simulation jump back to a previous statement (simulating a loop). We illustrate this technique in Appendix B.1.

For the if-direction, let  $\Pi$  be an FO-transducer over a schema  $\Upsilon$  that computes a query  $\mathcal{Q}$  on a single-node network  $\mathcal{T}$ . A While-program that computes the query  $\mathcal{Q}$  has to use exactly the same input and output schema as  $\Pi$ , namely,  $\Upsilon_{in}$  and  $\Upsilon_{out}$  respectively. The While-program is however allowed to declare any number of temporary relations. We may assume that  $\Pi$  does not read message relations in its internal queries, because no messages can be received on a single-node network. As a first case, let us additionally assume that the internal FO-queries of  $\Pi$  do not read relations  $Id$  and  $A11$  (the oblivious case). Now, because the memory relations of  $\Pi$  start empty, and temporary relations declared in the While-program also start empty, we can easily construct a While-program  $P$  that consists of one big loop, of which one iteration performs the same state changes as  $\Pi$  during one heartbeat transition. We provide an example in Appendix B.2. In order to terminate,  $P$  must detect repetition of transducer states, because this implies that  $\Pi$  has repeated a state and will output no new output facts. Detecting such a repetition is possible by using the technique of Abiteboul and Simon [Abiteboul and Simon 1991].

Let us now consider the second case where  $\Pi$  reads  $Id$  or  $A11$  (or both) in its internal queries. These relations can not be simulated by the While-program. Indeed, these relations are always non-empty from the perspective of  $\Pi$ , and a While-program can not create temporary relations to represent them: when the input is empty, the While-program can not invent a value to store in  $Id$  and  $A11$ , and when the input is nonempty, the While-program can in general not choose one value to store in  $Id$  and  $A11$ . Therefore, we will first eliminate the use of  $Id$  and  $A11$  from the queries of  $\Pi$ . Once this is done, we can apply the above translation for the oblivious case.

*Remove relation A11.* Note that in the FO-queries of  $\Pi$  we can replace the use of relation  $A11$  by  $Id$  because, on a single-node network, both relations have the same contents. Formally, in a transducer state there is a fact  $Id(a)$  iff there is a fact  $A11(a)$ .

*Remove relation Id.* Assume that relation  $A11$  is not used in  $\Pi$ . Next, we remove the use of relation  $Id$  from  $\Pi$ . We will only sketch the approach, and the details can be found in Appendix B.3. We use the work of Van den Bussche and Cabibbo [Van den Bussche and Cabibbo 1998], who have shown how to convert an ordinary (untyped) FO-formula to a *typed* formula that computes the same query. In typed formulas, each variable is of a specific *sort*, meaning that it ranges over an isolated domain of values. In our case, we distinguish between two sorts: (*i*) values in the active domain of an input database instance

over  $\Upsilon_{\text{in}}$ ; and, (ii) the identifier  $x$  of the single node in  $\mathcal{T}$  (with  $\mathcal{T}$  as defined above). We will denote these sorts as respectively *atom* and *id*. A *type*  $\tau$  is a tuple of sort symbols, like  $(\text{atom}, \text{id}, \text{id})$ .

Based on  $\Pi$ , we construct a second transducer  $\Pi^2$  as follows. For each relation  $R^{(k)} \in \Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}}$  of  $\Pi$  and each type  $\tau$  of arity  $k$ , transducer  $\Pi^2$  has a relation  $R_\tau^{(k)}$ . Transducer  $\Pi^2$  also has a memory relation **Adom** in which it stores all values from its input. We now describe how  $\Pi^2$  updates such a relation  $R_\tau^{(k)}$ . Let  $\varphi$  denote the FO-formula used by  $\Pi$  to insert tuples in relation  $R$  (deletion is similar). If for example  $\tau = (\text{atom}, \text{atom}, \dots, \text{id})$ , then transducer  $\Pi^2$  will use a formula of the form

$$\psi(\mathbf{u}_1, \dots, \mathbf{u}_k) \wedge \mathbf{Adom}(\mathbf{u}_1) \wedge \mathbf{Adom}(\mathbf{u}_2) \wedge \dots \wedge \mathbf{Id}(\mathbf{u}_k)$$

to insert into  $R_\tau$  the tuples of type  $\tau$  that are computed by  $\varphi$ . The formula  $\psi$  is basically the formula  $\varphi$ , but modified to cope with the separation of tuples by their type: each time  $\varphi$  reads a tuple from a relation  $S^{(l)}$ , formula  $\psi$  reads a tuple from the union  $\bigcup_{\tau \in \alpha} S_\tau^{(l)}$ , where  $\alpha$  is all types of arity  $l$ . This way,  $\Pi^2$  also computes the same query as  $\Pi$ .

Now, we can apply Proposition 1 of [Van den Bussche and Cabibbo 1998] to the formulas in  $\Pi^2$  to obtain new formulas in which there is no explicit reference to relations **Adom** and **Id**. Instead, the converted formulas use variables of two sorts (the *atom* and *id* sorts). In a last step, we can syntactically eliminate any reference to *id* variables, and obtain back normal FO-formulas. These can be used in a new transducer  $\Pi^3$ , which is oblivious, to compute the same query as  $\Pi$ .  $\square$

Now we can obtain the following result:

**THEOREM 6.2.** *A query is expressible in While if and only if it can be distributedly computed by an FO-transducer.*

**PROOF.** For the if-direction, let  $\Pi$  be an FO-transducer that distributedly computes a query  $\mathcal{Q}$ . Because  $\Pi$  is network-independent, the query  $\mathcal{Q}$  must also be computed when executing  $\Pi$  on a single-node network. Then, by using Lemma 6.1, there is a While-program that computes  $\mathcal{Q}$ .

For the only-if direction, let  $\mathcal{Q}$  be a query that can be computed by a While-program. We specify an FO-transducer to compute  $\mathcal{Q}$  in two steps. First we use Lemma 4.8 to obtain the entire input instance at every node. Every node can then act as if it was alone, ignoring any further messages, and simulate the While-program again using Lemma 6.1.  $\square$

For monotone queries we have the following, more specific result:

**THEOREM 6.3.** *Every monotone query expressible in While can be distributedly computed by an oblivious FO-transducer.*

**PROOF.** Let  $\mathcal{Q}$  be a monotone query expressible in While. We construct an oblivious FO-transducer to compute  $\mathcal{Q}$ . Note that Theorem 4.10 is not applicable, because that would give us an oblivious While-transducer, and not an oblivious FO-transducer. But the proof idea of the theorem can still be used.

First, we use the simple UCQ-protocol of Lemma 4.7 to let all nodes accumulate all input facts in memory. This does not require **Id** or **All**. Next, every time a node receives a new input fact, it starts or restarts a simulation of the While-program for  $\mathcal{Q}$ . The simulation uses the techniques of the proof of Lemma 6.1 (only-if direction), where specifically the output facts are first computed in temporary memory relations before being officially output. Checking whether a new input fact is received is done by comparing a received input fact with the previously accumulated input facts in memory. The restarting of the simulation of the While-program is done by emptying all memory relations, and restarting the program counter. The restart can happen at the moment a simulation is busy, in which

case the temporary output is discarded. The restart can also happen after a simulation was already successfully ended. Since the query  $\mathcal{Q}$  is monotone, no incorrect facts were output by previous simulations.

Eventually, every node will have accumulated all input facts, so the simulation can surely run to completion on all input facts. We also do not need relations `Id` and `All` to simulate the While-program. Hence, the transducer is oblivious.  $\square$

Note that the converse of Theorem 6.3, to the effect that every query distributedly computed by an oblivious FO-transducer is monotone and expressible in While, holds by combining Theorems 5.11 and 6.2 that give respectively the monotonicity of the query and the expressibility in While.

For our next result, we will use that FO is equivalent to  $\text{NrDatalog}^\neg$  [Abiteboul et al. 1995]. Basically, a program in  $\text{NrDatalog}^\neg$  is a sequence of  $\text{UCQ}^\neg$  statements. The following proposition shows that transducers can simulate this sequential composition of simpler statements:

PROPOSITION 6.4.

- (i) *Every query that can be distributedly computed by an FO-transducer can be distributedly computed by an  $\text{UCQ}^\neg$ -transducer.*
- (ii) *Every monotone query that can be distributedly computed by an FO-transducer can be distributedly computed by an oblivious  $\text{UCQ}^\neg$ -transducer.*

PROOF. First, we make a general observation. For every query  $\mathcal{Q}$  that is distributedly computed by an FO-transducer, we can apply Theorem 6.2 to know that  $\mathcal{Q}$  is expressible with a While-program  $P$ . Moreover, since the language FO is equivalent to  $\text{NrDatalog}^\neg$  [Abiteboul et al. 1995], every FO-statement in  $P$  can be replaced by a sequence of  $\text{UCQ}^\neg$ -statements, to obtain a new program  $P'$ . Then, it is clear that program  $P'$  can be simulated by an  $\text{UCQ}^\neg$ -transducer on a single-node network using iterated heartbeats, very similar to the proof of the only-if direction for Lemma 6.1.

For result (i), we let each node first collect a local copy of the entire input by using the protocol of Lemma 4.8, which can be done with a  $\text{UCQ}^\neg$ -transducer. After collecting the input, each node can simulate the program  $P'$  in isolation.

For result (ii), where  $\mathcal{Q}$  is monotone, we use instead Lemma 4.7 to let each node gradually accumulate all input, and we restart the simulation of  $P'$  when new inputs arrive.  $\square$

## 6.2. Datalog versus NrDatalog

What if we are only interested in Datalog? Between the languages Datalog and NrDatalog, a similar relation exists as between While and FO:

THEOREM 6.5. *A query is expressible in Datalog if and only if it can be distributedly computed by an inflationary NrDatalog-transducer.*

PROOF. First we consider the only-if direction. We construct an oblivious, inflationary transducer to simulate a Datalog program. The basic idea is the same as in the proof of Theorem 4.10. The input tuples are sent out and accumulated on every node. During every transition, we apply the immediate consequence operator of the Datalog program [Abiteboul et al. 1995], that can be expressed by NrDatalog. The relations `Id` and `All` are not needed, and the transducer can be made oblivious. Also, by the monotone nature of Datalog evaluation, deletions are never needed, and the transducer can be made inflationary.

Now we consider the if-direction. Let  $\mathcal{Q}$  be a query distributedly computed by an inflationary NrDatalog-transducer  $\Pi$  over a schema  $\Upsilon$ . We show that  $\mathcal{Q}$  can be expressed in Datalog. Because of network-independence, it is sufficient to look at the behavior of  $\Pi$  on a single-node network. We simulate this behavior with a Datalog program  $P$  as follows. We assume that the logical “and” and the universal quantifier are not core primitives of FO,

since these can be simulated by negation together with respectively the logical “or” and the existential quantifier. We call an FO-formula *positive* if each atom and existential quantifier occurs under an even number of negation symbols. The language NrDatalog is equivalent to positive FO. So,  $\Pi$  is just an inflationary FO-transducer, in which the internal FO-queries are positive. Now, the same transformation as in the proof of the if-direction for Lemma 6.1 can be applied to transform  $\Pi$  into a new FO-transducer  $\Pi'$  that computes  $\mathcal{Q}$  without reading relations  $\text{Id}$  and  $\text{A11}$ . Moreover, this transformation preserves the positivity of the formula. Hence,  $\Pi'$  can be immediately seen as an inflationary NrDatalog transducer that does not read  $\text{Id}$  and  $\text{A11}$ . Next, we unite all NrDatalog rules of  $\Pi'$  in a Datalog program  $P$ . Because  $P$  by the nature of Datalog can only accumulate its generated facts,  $P$  has at least the opportunities of  $\Pi'$  to join facts, and  $P$  outputs at least the output of  $\Pi'$ . Moreover, because  $\Pi'$  is inflationary itself,  $\Pi'$  eventually has the same opportunities to join facts as  $P$ . In conclusion,  $P$  computes exactly the original query  $\mathcal{Q}$ .  $\square$

It remains open if we can drop the word “inflationary” from Theorem 6.5.

### 6.3. Restrict Delivery

It is well-known that providing an order on the active domain increases the expressiveness of a query language [Abiteboul et al. 1995]. This result transfers nicely to our transducer model. By guaranteeing that only one message is delivered during every global transition, referred to as *1-delivery semantics* (cf. Section 3.3), an order can be established on each node:

**PROPOSITION 6.6.** *Under 1-delivery semantics, every PSPACE query can be computed by an FO-transducer network with at least two nodes.*

**PROOF.** In a network with at least two nodes, under 1-delivery semantics, each node can establish a linear order on the active domain by cooperating with the other nodes as follows. When a node has collected all inputs of the network (by means of Lemma 4.8), it sends out the elements of the active domain, that get forwarded by other nodes. Eventually, all these elements arrive back at the node, and the order can be established because at most one value is received at once. Then, each node can simulate a While-program on the collected input, that uses the established order. The transducer involved is not truly network-independent, as this only works when there are at least two nodes.  $\square$

### 6.4. Specialized CALM Properties

Using our previous results about expressivity, we obtain the following variants of Theorem 5.11. Especially, the second variant, which deals with Datalog, may come closest to the CALM conjecture as originally imagined by Hellerstein [Hellerstein 2010b].

**COROLLARY 6.7.** *Within each of the following two groups, the statements are equivalent, for any query  $\mathcal{Q}$ :*

- (1) (a)  $\mathcal{Q}$  can be distributedly computed by a coordination-free FO-transducer.  
 (b)  $\mathcal{Q}$  can be distributedly computed by an oblivious FO-transducer.  
 (c)  $\mathcal{Q}$  is monotone and expressible in the language While.
- (2) (a)  $\mathcal{Q}$  can be distributedly computed by a coordination-free, inflationary NrDatalog-transducer.  
 (b)  $\mathcal{Q}$  can be distributedly computed by an oblivious, inflationary NrDatalog-transducer.  
 (c)  $\mathcal{Q}$  is expressible in Datalog.

**PROOF.** Regarding (1), for (c)  $\Rightarrow$  (b) use Theorem 6.3; for (b)  $\Rightarrow$  (a) use Proposition 5.8; for (a)  $\Rightarrow$  (c) use Theorems 5.10 and 6.2 to obtain respectively the properties of “ $\mathcal{Q}$  is monotone” and “expressible in the language While”.

Regarding (2), for  $(c) \Rightarrow (b)$  use (proof of if-direction in) Theorem 6.5; for  $(b) \Rightarrow (a)$  use Proposition 5.8; for  $(a) \Rightarrow (c)$  use Theorem 6.5.  $\square$

## 7. VARIATION OF THE MODEL

In the literature on declarative networking, a seemingly common language feature seems to be that nodes do not simply send each message to all of their neighbors, but instead to a specifically addressed neighbor [Loo et al. 2009; Grumbach and Wang 2010; Alvaro et al. 2011b; Nigam et al. 2012]. We call this the *addressing model*. One could argue that this model lies closer to how real networks operate, and that is why we devote a small section to this model.

### 7.1. Addressing Transducers

Recall our original epidemic transducer model that was presented in Sections 2.2 and 3. An *addressing transducer*  $\Pi$  over a transducer schema  $\Upsilon$  is the same as an epidemic transducer over  $\Upsilon$  with the only difference that for a message relation  $R^{(k)} \in \Upsilon_{\text{msg}}$ , the sending query will produce facts of arity  $k+1$  instead of  $k$ . The extra component will contain the addressee of each message, which is by convention the first component. Now we look at how the operational semantics must be changed accordingly. With  $\Pi$  as above, consider a transducer network  $\mathcal{T} = \langle \mathcal{N}, \Upsilon, \Pi \rangle$ . We define how an active node  $x \in \mathcal{N}$  does a global transition. Similarly to the original definition of global transition (in Section 3), we let  $x$  receive some messages from its message buffer. Then,  $x$  does a local transition in which it generates a set of newly sent messages  $J_{\text{snd}}$ , each having the addressee specified as their first component. Now, the messages that are effectively added to the message buffer of another node  $y$ , denoted  $K^{\rightarrow y}$ , is defined as: if  $y$  is a neighbor of  $x$  then  $K^{\rightarrow y} = \{R(\bar{a}) \mid R(y, \bar{a}) \in J_{\text{snd}}\}$  and otherwise  $K^{\rightarrow y} = \emptyset$ , i.e., we select precisely the messages that are sent to  $y$  when  $y$  is a neighbor. An addressee value that is not a neighbor of  $x$  will result in the loss of the corresponding message. Note that the message buffers contain facts without an explicit addressee-component, like in the operational semantics for epidemic transducers.

As a special case, if  $\mathcal{N}$  forms a complete graph, every node can send a message to every individual other node.

### 7.2. Properties

First, all our previous results that do not explicitly mention “oblivious” transducers still hold for addressing transducers because, when there is no restriction on using **Id** or **All**, addressing transducers and epidemic transducers are equivalent in terms of what queries they can compute. Indeed, it was already noted in Section 3.4 that the epidemic model can simulate the addressing model by manually adding an addressee-component to every message relation in the transducer schema, and by comparing for each received message the addressee component with the value in the local relation **Id**. The other direction is also possible, namely that an addressing transducer can simulate an epidemic one. It suffices for the addressing transducer to send each message explicitly to *every* neighbor.

Interestingly, a notion of obliviousness can also be defined for addressing transducers. Formally, we say that an addressing transducer is *oblivious* if the relations **Id** and **All** are only used in the message sending queries.<sup>4</sup>

Now, most of our results involving oblivious epidemic transducers also hold for oblivious *addressing* transducers, because of the following reasons. First, the proof techniques frequently use that every node sends out its local input facts, and these are forwarded so that eventually all nodes accumulate all inputs. This can be done with an oblivious addressing transducer as well. Second, these results are mostly about *network-independent* transducers, and a frequently occurring idea in those proofs is that we only focus on the behaviour of a

<sup>4</sup>Note, we can not completely forbid the use of node relations because we need to indicate addressees.

single node: an oblivious transducer can not distinguish between a single-node network and a multi-node network unless it receives a message, so on a single-node network it should exhibit predictable behaviour if it wants to be network-independent. This trick is also applicable to oblivious addressing transducers, because they too can not distinguish between single-node and multi-node networks unless they receive a message. We now explicitly give the results that are not transferable, and why this is the case.

First, Proposition 4.2 does not hold for oblivious addressing transducer networks, because this result talks about a concrete transducer network. The transducer may now exploit the number of nodes. In particular, if there are multiple nodes, the transducer may assume messages are eventually delivered. So, it is possible to construct a multi-node transducer network in which the oblivious addressing transducer smuggles node identifiers in the sent messages (by reading  $\mathbf{All}$ ), and when these arrive, it is possible to only output the input facts whose active domain is contained in the set of node identifiers. This would prevent the transducer network from computing a generic query.

Although not purely about obliviousness, the result of Theorem 5.13 is also not transferable to addressing transducers, as illustrated by the following example, where relation  $\mathbf{All}$  is used to make the nodes dependent on message arrival.

*Example 7.1.* We give an addressing transducer that reads only relation  $\mathbf{All}$  and that computes the nonmonotone emptiness query on a nullary input relation  $R$  (see also Example 5.5). Reading relation  $\mathbf{All}$  in output or memory queries, a node can know from the start if it is alone or not. If the node is alone, then it can immediately output the desired result by looking at the local relation  $R$ . But if there are multiple nodes, every node  $x$  sends each local fact  $\mathbf{All}(y)$  as a message  $A(y)$  to node  $y$ . Although the operational semantics drops the message when  $y$  is not a neighbor of  $x$ , because each network is connected,  $y$  has at least one neighbor from which it will receive  $A(y)$ . This way, each node can establish its own identity. Next, the same protocol as in Example 5.5 can be followed.  $\square$

## 8. CONCLUSION

Encouraged by Hellerstein [Hellerstein 2010a; Hellerstein 2010b], we have tried in this paper to formalize and prove the CALM Conjecture. We do not claim that our approach is the only one that works. Yet, we believe our approach is natural because it is firmly grounded in previous database theory practice, and delivers solid results. Much further work is possible; we list a few obvious topics:

- Look at Hellerstein’s other conjectures (e.g. the CRON conjecture [Ameloot and Van den Bussche 2012b]);
- Investigate the expressiveness of variations or extensions of the basic distributed computation model presented here; and,
- Identify special cases where essential semantic notions such as monotonicity, consistency [Ameloot and Van den Bussche 2012a], network-independence, coordination-freeness, etc, are decidable.

## APPENDIX

### A. EXPRESSING QUERIES

#### A.1. Undecidability

PROPOSITION A.1. *Consistency for FO-transducer networks is undecidable.*

PROOF. We reduce the finite satisfiability problem for FO to deciding consistency for FO-transducer networks. Let  $\varphi$  be an FO-sentence over a database schema  $\mathcal{D}$ . We construct a transducer network  $\mathcal{T}$  that is consistent iff  $\varphi$  is *not* finitely satisfiable.

Consider the transducer  $\Pi$  from Example 4.5, that is over transducer schema  $\Upsilon$ . We may assume without loss of generality that the relation names of  $\Upsilon$  do not occur in  $\mathcal{D}$ . We obtain

a new transducer schema  $\Upsilon'$  from  $\Upsilon$  by adding  $\mathcal{D}$  to  $\Upsilon_{\text{in}}$ . We obtain a new transducer  $\Pi'$  from  $\Pi$  by modifying the send rule for relation  $U$  to only send facts when  $\varphi$  is satisfied on the local input over  $\mathcal{D}$ .

Suppose that  $\varphi$  is finitely satisfiable. Consider a network  $\mathcal{N}$  with two (connected) nodes  $x$  and  $y$ . Let  $\mathcal{T}$  denote the transducer network with  $\Pi$  on both nodes. Let  $I$  be a database instance over  $\mathcal{D}$  on which  $\varphi$  is true. Consider the input distributed database instance  $H$  with  $H(x) = I \cup \{R(1), R(2)\}$  and  $H(y) = \emptyset$ . Consider the prefix of a fair run of  $\mathcal{T}$  on  $H$  where first  $x$  does a heartbeat transition: because  $\varphi$  is true on  $I$ , node  $x$  sends messages  $U(1)$  and  $U(2)$  to  $y$ . In the second transition, we can deliver message  $U(1)$  to  $y$  or message  $U(2)$ , or both. Each choice results in a different output at  $y$ , which results in a different global output because  $x$  never outputs anything ( $x$  does not receive messages from  $y$ ). Hence, the transducer network is not consistent.

For the other direction, suppose that  $\varphi$  is not finitely satisfiable. Then there is no input distributed database instance for  $\mathcal{T}$  on which messages will be sent, in which case  $\mathcal{T}$  is consistent because on every input, every run produces the empty output.  $\square$

**PROPOSITION A.2.** *Network independence for FO-transducers is undecidable.*

**PROOF.** We reduce the finite satisfiability problem of FO to deciding network-independence of FO-transducers. Let  $\varphi$  be an FO-sentence over a database schema  $\mathcal{D}$ . We construct a transducer  $\Pi$  that is network-independent iff  $\varphi$  is *not* finitely satisfiable.

Consider the transducer  $\Pi$  in Example 4.6, that is over schema  $\Upsilon$ . We may assume without loss of generality that the relation names of  $\Upsilon$  do not occur in  $\mathcal{D}$ . We obtain a new transducer  $\Pi'$  by modifying  $\Pi$  as follows:

- using the protocol of Lemma 4.8, we let all nodes collect all the input facts of  $\mathcal{D}$  available on the network; and,
- the output query is modified so that output can only be produced if the formula  $\varphi$  is satisfied on the (fully) collected instance over  $\mathcal{D}$  in memory, in addition to detecting a complete network-topology (as before).

It is possible to construct  $\Pi'$  so that its output, message, and memory relations are not in  $\mathcal{D}$ . Note that on any network, the transducer network resulting from  $\Pi'$  is consistent: this is because before the output can be produced at a node, it should have obtained the *entire* input over  $\mathcal{D}$ , and it should have detected that the network topology is complete.

Suppose that  $\varphi$  is finitely satisfiable. Let  $I$  be a database instance over  $\mathcal{D}$  on which  $\varphi$  is true. Denote  $I' = I \cup \{R(1)\}$ . Then, on a complete network-topology, the transducer network resulting from  $\Pi'$  outputs  $T(1)$  on any horizontal partition of  $I'$ . Indeed, the nodes forward all facts of  $I$  to each other and in any run there will be a moment when all nodes have these facts *and* have detected that the network-topology is complete. On any other network-topology, for every input partition of  $I'$ , the resulting transducer network outputs nothing. Therefore  $\Pi'$  is not network-independent.

For the other direction, suppose that  $\varphi$  is not finitely satisfiable. Then every transducer network for  $\Pi'$  computes the empty query and  $\Pi'$  is therefore network-independent.  $\square$

## A.2. Proof of Lemma 4.8

First, we specify the parts of the transducer-schema  $\Upsilon$  that we need:  $\Upsilon_{\text{in}} = \mathcal{D}$ ;  $\Upsilon_{\text{msg}} = \{(R^{\text{msg}}, k+1), (R^{\text{ack}}, k+2) \mid R^{(k)} \in \mathcal{D}\} \cup \{\text{done}^{(2)}\}$ ; and,

$$\begin{aligned} \Upsilon_{\text{mem}} = & \{(R^{\text{mem}}, k), (R^{\text{ackMem}}, k+2) \mid R^{(k)} \in \mathcal{D}\} \\ & \cup \{\text{doneMem}^{(2)}, \text{notDone}^{(1)}, \text{missing}^{(0)}\} \cup \{\text{started}^{(0)}, \text{ready}^{(0)}\}. \end{aligned}$$

We have not specified the database schema  $\Upsilon_{\text{out}}$  because this schema is not used in the transducer construction. The idea is that a node  $x$  will send its local input facts over relation  $R^{(k)} \in \Upsilon_{\text{in}}$  as facts with predicate  $R^{\text{msg}}$ , with as the last component its own node identifier to indicate the origin of the fact (hence the increased arity of  $k + 1$ ). We call this last component the “tag”. Next, when a node  $y$  receives a tagged  $R^{\text{msg}}$ -fact, it removes the tag and stores the fact in its relation  $R^{\text{mem}}$  and it sends an  $R^{\text{ack}}$ -fact to acknowledge the receipt of it. These acknowledgments have the contents of the received  $R^{\text{msg}}$ -fact (including the tag), with an additional last component containing the identifier of  $y$ . The node  $x$  checks whether  $y$  has (eventually) acknowledged all input facts of  $x$ . If yes, then  $x$  sends out  $\text{done}(x, y)$ . From the viewpoint of  $y$ , if  $y$  has received  $\text{done}(x, y)$  from all other nodes  $x$  then it knows that it has accumulated all the input facts on the network.

Some further details are as follows. Let  $x$  be a node. For a relation  $R$  in  $\Upsilon_{\text{in}}$ , node  $x$  uses the relation  $R^{\text{ackMem}}$  to store all received acknowledgments for its local input facts over relation  $R$ . The relation  $\text{notDone}$  is used by  $x$  to remember all the other nodes that have not yet acknowledged all input facts of  $x$ . The relation  $\text{doneMem}$  is used by  $x$  to remember all  $\text{done}$ -messages having as the second component its own identifier. The relation  $\text{missing}$  is nonempty at  $x$  as long as  $x$  has not received a  $\text{done}$ -message from all other nodes. The relation  $\text{started}$  helps  $x$  to differentiate between its first local transition and all the following local transitions. This makes sure that other memory relations have been correctly initialized before they are read.

Now we specify the transducer  $\Pi$  over  $\Upsilon$ . We describe the queries with the language  $\text{UCQ}^-$ , which is contained in FO. As usual, unions are then expressed by having multiple rules with the same head. As a general remark about message sending, for each message relation, we always have a “forwarding” rule that just resends all received messages, so that eventually all nodes can receive those messages.

First, for each  $R^{(k)} \in \mathcal{D}$  we have the following rules to let all nodes forward their (tagged) input  $R$ -facts to each other, and to store the received facts in memory (including acknowledgments):

$$\begin{aligned}
R_{\text{snd}}^{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}) &\leftarrow R(\mathbf{u}_1, \dots, \mathbf{u}_k), \text{Id}(\mathbf{x}). \\
R_{\text{snd}}^{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}) &\leftarrow R^{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}). \\
R_{\text{ins}}^{\text{mem}}(\mathbf{u}_1, \dots, \mathbf{u}_k) &\leftarrow R(\mathbf{u}_1, \dots, \mathbf{u}_k). \\
R_{\text{ins}}^{\text{mem}}(\mathbf{u}_1, \dots, \mathbf{u}_k) &\leftarrow R^{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}). \\
R_{\text{snd}}^{\text{ack}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}, \mathbf{y}) &\leftarrow R^{\text{msg}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}), \text{Id}(\mathbf{y}). \\
R_{\text{snd}}^{\text{ack}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}, \mathbf{y}) &\leftarrow R^{\text{ack}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}, \mathbf{y}). \\
R_{\text{ins}}^{\text{ackMem}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}, \mathbf{y}) &\leftarrow R^{\text{ack}}(\mathbf{u}_1, \dots, \mathbf{u}_k, \mathbf{x}, \mathbf{y}), \text{Id}(\mathbf{x}).
\end{aligned}$$

We also specify rules for the other relations in  $\Upsilon$ . For convenience, let us denote  $\mathcal{D} = \{R_1^{(k_1)}, R_2^{(k_2)}, \dots, R_n^{(k_n)}\}$ . First, we need to remember that the first local transition has already happened, using the following rule:

$$\text{started}_{\text{ins}}() \leftarrow .$$

Next, on each node  $x$ , the relation  $\text{notDone}$  contains all nodes that have not yet acknowledged the receipt of all local input facts of  $x$ . A node does not have to acknowledge its own input facts, so  $\text{notDone}$  will not contain  $x$  itself. This relation is recomputed during every local transition, using the following rules:

$$\begin{aligned}
\text{notDone}_{\text{ins}}(y) &\leftarrow R_1(\mathbf{u}_1, \dots, \mathbf{u}_{k_1}), \text{Id}(\mathbf{x}), \text{All}(y), \neg \text{Id}(y), \\
&\quad \neg R_1^{\text{ackMem}}(\mathbf{u}_1, \dots, \mathbf{u}_{k_1}, \mathbf{x}, y). \\
&\vdots \\
\text{notDone}_{\text{ins}}(y) &\leftarrow R_n(\mathbf{u}_1, \dots, \mathbf{u}_{k_n}), \text{Id}(\mathbf{x}), \text{All}(y), \neg \text{Id}(y), \\
&\quad \neg R_n^{\text{ackMem}}(\mathbf{u}_1, \dots, \mathbf{u}_{k_n}, \mathbf{x}, y). \\
\text{notDone}_{\text{del}}(y) &\leftarrow \text{notDone}(y).
\end{aligned}$$

Note that if a node was initialized with no local input tuples over some relation  $R_i^{(k_i)} \in \mathcal{D}$ , then the corresponding insertion rule for `notDone` will not fire. In that case, the node will not consider any other nodes “responsible” for acknowledging the receipt of its input facts over relation  $R_i$ . This is the desired behavior. Also, the deletion rule for `notDone` allows for the recomputation of `notDone` during every local step: only the nodes that have not confirmed every local input fact are reinserted again. Thus, after a while, relation `notDone` will become (and remain) empty.

When a node  $x$  notices that another node  $y$  has acknowledged all local input facts of  $x$ , node  $x$  sends out `done(x, y)`. This is accomplished by the following rules:

$$\begin{aligned}
\text{done}_{\text{snd}}(x, y) &\leftarrow \text{started}(), \text{Id}(x), \text{All}(y), \neg \text{Id}(y), \neg \text{notDone}(y). \\
\text{done}_{\text{snd}}(x, y) &\leftarrow \text{done}(x, y).
\end{aligned}$$

These `done`-messages are stored at the addressed node:

$$\text{doneMem}_{\text{ins}}(x, y) \leftarrow \text{done}(x, y), \text{Id}(y).$$

Finally, when a node  $y$  has received `done` from all other nodes, it can output the `ready` flag. This is accomplished by the following rules:

$$\begin{aligned}
\text{missing}_{\text{ins}}() &\leftarrow \text{Id}(y), \text{All}(z), \neg \text{Id}(z), \neg \text{doneMem}(z, y). \\
\text{missing}_{\text{del}}() &\leftarrow \text{missing}(). \\
\text{ready}_{\text{ins}}() &\leftarrow \text{started}(), \neg \text{missing}().
\end{aligned}$$

Note that in a single-node network, the `missing`-fact is never created. In that case, the `ready`-fact is already produced in the second transition (thus after `started()` is created).

The transducer  $\Pi$  above can actually be made inflationary as well. In particular, using the equivalence  $\text{FO} = \text{NrDatalog}^-$ , we can write a  $\text{NrDatalog}^-$ -transducer where relations `notDone` and `missing` do not appear in the memory schema  $\Upsilon_{\text{mem}}$  but are computed locally: one would locally compute `notDone` in the sending query for the `done`-relation, and one would locally compute `missing` in the insertion query for the `ready`-relation.

## B. EXPRESSIVENESS ANALYSIS

### B.1. Proof of Lemma 6.1 (While to FO-transducer)

We show how to simulate a While-program on a single-node FO-transducer network. A While-program can be simulated by iterated heartbeats using well-known techniques [Abiteboul and Vianu 1991a]. Because this is not entirely obvious, we will illustrate the technique. Consider the simple While-program in Algorithm 1 over input schema  $\mathcal{D} = \{R^{(2)}, S^{(1)}\}$  and output schema  $\mathcal{D}' = \{T^{(2)}\}$ . Intuitively, if  $R$  represents a graph then the While program collects all edges that are reachable from the nodes in relation  $S$ . By introducing a temporary relation  $U$ , we can rewrite this program so that  $T$  is only modified at the very end. See

**Algorithm 1** A While-program.

---

```

 $T := \{T(u, v) \mid R(u, v) \wedge S(u)\};$ 
 $T^{\text{prev}} := \emptyset;$ 
while  $(T \setminus T^{\text{prev}} \neq \emptyset)$  do
   $T^{\text{prev}} := T;$ 
   $T := T \cup \{T(v, w) \mid \exists u(T(u, v) \wedge R(v, w))\};$ 
end

```

---

**Algorithm 2** Rewritten version of Algorithm 1.

---

```

 $U := \{U(u, v) \mid R(u, v) \wedge S(u)\};$ 
 $U^{\text{prev}} := \emptyset;$ 
while  $(U \setminus U^{\text{prev}} \neq \emptyset)$  do
   $U^{\text{prev}} := U;$ 
   $U := U \cup \{U(v, w) \mid \exists u(U(u, v) \wedge R(v, w))\};$ 
end
 $T := U;$ 

```

---

**Algorithm 3** List program for Algorithm 2.

---

```

 $\text{step}^1: U := \{U(u, v) \mid R(u, v) \wedge S(u)\};$ 
 $\text{step}^2: U^{\text{prev}} := \emptyset;$ 
 $\text{step}^3: V := \{V() \mid U(u, v) \wedge \neg U^{\text{prev}}(u, v)\};$ 
 $\text{step}^4: \text{if } \langle V = \emptyset \rangle \text{ goto } \langle \text{step}^8 \rangle;$ 
 $\text{step}^5: U^{\text{prev}} := U;$ 
 $\text{step}^6: U := U \cup \{U(v, w) \mid \exists u(U(u, v) \wedge R(v, w))\};$ 
 $\text{step}^7: \text{goto } \langle \text{step}^3 \rangle;$ 
 $\text{step}^8: T := U;$ 

```

---

Algorithm 2. Next, any While-program can be translated to a list of statements in which explicit control flow is represented by conditional and unconditional “goto” statements. This also works for nested while-loops. When translating Algorithm 2 to this form, we obtain Algorithm 3. There,  $V$  holds the result of the expression that is tested for non-emptiness by the while-loop condition. Let us refer to this form as a “list program”.

We will now simulate Algorithm 3 with an FO-transducer. First, we define a transducer schema  $\Upsilon$  with  $\Upsilon_{\text{in}} = \mathcal{D}$ ,  $\Upsilon_{\text{out}} = \mathcal{D}'$ ,  $\Upsilon_{\text{msg}} = \emptyset$  and

$$\Upsilon_{\text{mem}} = \{(U, 2), (U^{\text{prev}}, 2), (\text{step}^1, 0), \dots, (\text{step}^8, 0)\}.$$

Relations  $\text{step}^1$  to  $\text{step}^8$  model the program counter of the list program. The idea is that at any moment in time at most one of these relations is active (nonempty) and that they activate each other in the correct way in order to represent the desired control flow. We

now specify a transducer  $\Pi$  over  $\Upsilon$  that simulates Algorithm 3 on a single-node network. As a general remark, some  $\mathbf{step}^i$ -relations with  $i \in \{1, \dots, 8\}$  are read inside the queries that update relations  $U$  and  $U^{\text{prev}}$ , to make sure that  $U$  and  $U^{\text{prev}}$  are updated only at the moment when the original list program updates them. We will also specify the deletion queries for  $U$  and  $U^{\text{prev}}$  because assignment in a While-program is destructive in the sense that previous facts can only stay in the relation if they are on the right hand side of the assignment. The queries are as follows:

$$\mathbf{step}_{\text{ins}}^1() \leftarrow \neg \mathbf{step}^1(), \neg \mathbf{step}^2(), \dots, \neg \mathbf{step}^8().$$

$$U_{\text{ins}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^1(), R(\mathbf{u}, \mathbf{v}), S(\mathbf{u}).$$

$$U_{\text{ins}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^6(), U(\mathbf{u}, \mathbf{v}).$$

$$U_{\text{ins}}(\mathbf{v}, \mathbf{w}) \leftarrow \mathbf{step}^6(), U(\mathbf{u}, \mathbf{v}), R(\mathbf{v}, \mathbf{w}).$$

$$U_{\text{del}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^6(), U(\mathbf{u}, \mathbf{v}).$$

$$\mathbf{step}_{\text{ins}}^2() \leftarrow \mathbf{step}^1().$$

$$\mathbf{step}_{\text{ins}}^3() \leftarrow \mathbf{step}^2().$$

$$\mathbf{step}_{\text{ins}}^3() \leftarrow \mathbf{step}^7().$$

$$V_{\text{ins}}() \leftarrow \mathbf{step}^3(), U(\mathbf{u}, \mathbf{v}), \neg U^{\text{prev}}(\mathbf{u}, \mathbf{v}).$$

$$V_{\text{del}}() \leftarrow \mathbf{step}^4().$$

$$\mathbf{step}_{\text{ins}}^4() \leftarrow \mathbf{step}^3().$$

$$\mathbf{step}_{\text{ins}}^5() \leftarrow \mathbf{step}^4(), V().$$

$$U_{\text{ins}}^{\text{prev}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^5(), U(\mathbf{u}, \mathbf{v}).$$

$$U_{\text{del}}^{\text{prev}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^5(), U^{\text{prev}}(\mathbf{u}, \mathbf{v}).$$

$$\mathbf{step}_{\text{ins}}^6() \leftarrow \mathbf{step}^5().$$

$$\mathbf{step}_{\text{ins}}^7() \leftarrow \mathbf{step}^6().$$

$$\mathbf{step}_{\text{ins}}^8() \leftarrow \mathbf{step}^4(), \neg V().$$

$$T_{\text{out}}(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{step}^8(), U(\mathbf{u}, \mathbf{v}).$$

$$\text{For } i \in \{1, \dots, 7\}: \mathbf{step}_{\text{del}}^i() \leftarrow \mathbf{step}^i().$$

We never delete  $\mathbf{step}^8()$  because we do not want to accidentally restart the simulation of the While program: indeed, when  $\mathbf{step}^i()$  for each  $i \in \{1, \dots, 8\}$  is missing, a new  $\mathbf{step}^1()$

fact is created. Note that the queries for inserting into relations `step`<sup>5</sup> and `step`<sup>8</sup> together simulate the if-goto statement at line 4 of Algorithm 3.

### B.2. Proof of Lemma 6.1 (FO-transducer to While)

Let  $\mathcal{T}$  be a single-node transducer network, running an oblivious FO-transducer  $\Pi$  that does not read message relations. Let  $\mathcal{Q}$  denote the query computed by  $\mathcal{T}$ . We describe a While-program  $P$  to compute  $\mathcal{Q}$ . Intuitively,  $P$  consists of one big loop that during one iteration performs the same state changes as  $\Pi$  during one heartbeat transition. Also, in order to terminate, the While-program must detect repetition of transducer states.

To illustrate, consider the transducer given in Algorithm 4, that computes the query of Section B.1 intentionally in a more complex way. Specifically, the memory relation  $A$  continuously alternates between being empty and nonempty. The insertion queries for relations  $U$  and  $T$  only produce a nonempty output when  $A$  is nonempty.

For two sets  $S_1$  and  $S_2$ , let  $\text{diff}(S_1, S_2)$  abbreviate the expression  $((S_1 \setminus S_2) \cup (S_2 \setminus S_1))$ . Consider now the While-program  $P$  in Algorithm 5 to explicitly simulate the transducer of Algorithm 4. The program  $P$  keeps simulating the updates to the relations  $A$ ,  $U$  and  $T$  until no more changes occur to all of them. Surely, if the transducer state stops changing, no more output facts can be produced because only heartbeat transitions can occur and because local transitions are deterministic. However, because of the alternating behavior of relation  $A$ , the program  $P$  never stops if it uses only this test. For the transducer itself, the alternating behavior of relation  $A$  is no problem because its output is defined on infinite runs anyway. But program  $P$  needs to halt because otherwise its output is undefined. Using the technique of Abiteboul and Simon [Abiteboul and Simon 1991], however,  $P$  can be modified to detect that it is in an infinite loop. This implies that the transducer has repeated a state and will output no new output facts. After detecting the infinite loop, the program  $P$  then breaks the loop and the final contents of relation  $T$  is the output.

---

#### Algorithm 4 A simple FO-transducer.

---

Schema:  $\Upsilon_{\text{in}} = \{R^{(2)}, S^{(1)}\}$ ;  $\Upsilon_{\text{out}} = \{T^{(2)}\}$ ;  $\Upsilon_{\text{msg}} = \emptyset$ ;  $\Upsilon_{\text{mem}} = \{A^{(0)}, U^{(2)}\}$ .

Queries:

$$A_{\text{ins}}() \leftarrow \neg A().$$

$$A_{\text{del}}() \leftarrow A().$$

$$U_{\text{ins}}(\mathbf{v}, \mathbf{w}) \leftarrow A(), T(\mathbf{u}, \mathbf{v}), R(\mathbf{v}, \mathbf{w}).$$

$$U_{\text{del}}(\mathbf{u}, \mathbf{v}) \leftarrow A(), U(\mathbf{u}, \mathbf{v}).$$

$$T_{\text{out}}(\mathbf{u}, \mathbf{v}) \leftarrow A(), R(\mathbf{u}, \mathbf{v}), S(\mathbf{u}).$$

$$T_{\text{out}}(\mathbf{u}, \mathbf{v}) \leftarrow A(), U(\mathbf{u}, \mathbf{v}).$$


---

### B.3. Proof of Lemma 6.1 (eliminate Id)

Let  $\mathcal{T}$  be a single-node transducer network, with FO-transducer  $\Pi$  over transducer schema  $\Upsilon$ . Let  $x$  denote the single node. Let  $\mathcal{Q}$  denote the query computed by  $\mathcal{T}$ . We assume that  $\Pi$  does not read relation `All` and does not read message relations. Here we will show how to rewrite  $\Pi$  to eliminate the use of relation `Id` as well. We thus obtain an oblivious transducer.

---

**Algorithm 5** While-program based on the transducer of Algorithm 4.

---

[All auxiliary relations start empty]

while ( $\neg \text{started}() \vee (\text{diff}(A, A^{\text{prev}}) \cup \text{diff}(U, U^{\text{prev}}) \cup \text{diff}(T, T^{\text{prev}}) \neq \emptyset)$ ) do

$\text{started} := \{\text{started}()\};$

$A^{\text{prev}} := A;$

$U^{\text{prev}} := U;$

$T^{\text{prev}} := T;$

$A_{\text{ins}} := \{A() \mid \neg A^{\text{prev}}()\};$

$A_{\text{del}} := \{A() \mid A^{\text{prev}}()\};$

$U_{\text{ins}} := \{U(\mathbf{v}, \mathbf{w}) \mid \exists \mathbf{u}(A^{\text{prev}}() \wedge T^{\text{prev}}(\mathbf{u}, \mathbf{v}) \wedge R(\mathbf{v}, \mathbf{w}))\};$

$U_{\text{del}} := \{U(\mathbf{u}, \mathbf{v}) \mid A^{\text{prev}}() \wedge U^{\text{prev}}(\mathbf{u}, \mathbf{v})\};$

[We explicitly simulate the updating of memory relations]

$A := (A^{\text{prev}} \cup (A_{\text{ins}} \setminus A_{\text{del}})) \setminus (A_{\text{del}} \setminus A_{\text{ins}});$

$U := (U^{\text{prev}} \cup (U_{\text{ins}} \setminus U_{\text{del}})) \setminus (U_{\text{del}} \setminus U_{\text{ins}});$

$T := T \cup \{T(\mathbf{u}, \mathbf{v}) \mid A^{\text{prev}}() \wedge R(\mathbf{u}, \mathbf{v}) \wedge S(\mathbf{u})\}$

$\cup \{T(\mathbf{u}, \mathbf{v}) \mid A^{\text{prev}}() \wedge U^{\text{prev}}(\mathbf{u}, \mathbf{v})\};$

end

---

**B.3.1. Definitions and Notations.** Van den Bussche and Cabibbo [Van den Bussche and Cabibbo 1998] have shown how to convert an *untyped* FO-formula to a *typed* one that computes the same query (over a typed relation schema). We will use that result here. We distinguish between two *sorts* of values:

- values in the active domain of an input over  $\Upsilon_{\text{in}}$ ; and,
- the single node  $x$  of  $\mathcal{T}$ .

For these sorts we use the symbols *adom* and *id* respectively. The technique of [Van den Bussche and Cabibbo 1998] requires that each sort has a completely separated domain of values. Hence, we will assume that  $x$  does not occur in the *adom* values. We will see later that this assumption has no undesired consequences.

The definitions below are specifically tailored for the two sorts above and therefore less general as in the paper [Van den Bussche and Cabibbo 1998]. A *type*  $\tau$  is a tuple of sort symbols. An example is  $(\text{adom}, \text{id}, \text{adom}, \text{adom})$ . A *k-type* is a type with arity  $k$ .

A *typed database schema*  $\mathcal{F}$  is a finite set of pairs  $(R, \tau)$  with  $R$  a relation name and  $\tau$  the associated type of  $R$ , such that no relation name occurs twice. This corresponds to an ordinary database schema  $\text{untyped}(\mathcal{F})$  that consists of

- a relation  $R^{(k)}$  for each  $(R, \tau) \in \mathcal{F}$  with  $k$  the arity of  $\tau$ ; and,
- the relations  $\text{Id}^{(1)}$  and  $\text{Adom}^{(1)}$  (assumed not to be in  $\mathcal{F}$  already).

We define a *typed database instance*  $I$  over  $\mathcal{F}$  as a normal database instance over  $\text{untyped}(\mathcal{F})$  such that

- $I|_{\text{Id}} = \{\text{Id}(x)\},$
- $I|_{\text{Adom}} = \{\text{Adom}(a) \mid a \in \text{adom}(I), a \neq x\},$

- for each fact  $R(a_1, \dots, a_k) \in I$ , where  $R$  is not  $\text{Id}$  or  $\text{Adom}$ , and  $\tau = (s_1, \dots, s_k)$  is the type of  $R$  in  $\mathcal{F}$ , we require for each  $i \in \{1, \dots, k\}$  that  $a_i \neq x$  if  $s_i = \text{adom}$  and  $a_i = x$  if  $s_i = \text{id}$ . We say that this fact *has* type  $\tau$ .

Also, we will specify the queries of FO-transducers with *relational calculus* to make a better connection with the previous work [Van den Bussche and Cabibbo 1998].<sup>5</sup> We assume that the active domain semantics is used to evaluate these queries [Abiteboul et al. 1995].

**B.3.2. Split tuples by type.** We now construct an intermediate transducer  $\Pi^2$  to compute on the single-node network  $x$  the same query  $\mathcal{Q}$  as  $\Pi$ , when we restrict attention to input instances  $I$  over  $\Upsilon_{\text{in}}$  with  $x \notin \text{adom}(I)$ . We define the schema  $\Upsilon^2$  of  $\Pi^2$  as follows. First,  $\Upsilon_{\text{in}}^2 = \Upsilon_{\text{in}}$ ;  $\Upsilon_{\text{out}}^2 = \Upsilon_{\text{out}}$ ;  $\Upsilon_{\text{msg}}^2 = \emptyset$ ; and  $\Upsilon_{\text{mem}}^2$  consists of

- relation  $\text{Adom}^{(1)}$ ;
- relation  $\text{started}^{(0)}$ ; and,
- the relations  $R_\tau^{(k)}$  for each  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$  and *each*  $k$ -type  $\tau$ .

Concerning  $\Upsilon_{\text{mem}}^2$ , the idea is that in  $\Pi^2$  the facts over relation  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$  of  $\Pi$  are represented by all the disjoint relations  $\{R_\tau^{(k)} \mid \tau \text{ is a } k\text{-type}\}$ .

We now describe the FO-queries of  $\Pi^2$ . To start, in the first transition, we initialize relation  $\text{Adom}$  to contain the active domain of the input over  $\Upsilon_{\text{in}}$ , and we also compute the fact  $\text{started}()$  so that other queries can know that relation  $\text{Adom}$  has been initialized. We omit the details of these relatively simple queries.

Next, we define the queries for the other memory relations of  $\Pi^2$ . Consider  $R_\tau^{(k)} \in \Upsilon_{\text{mem}}^2$ , with  $R^{(k)} \in \Upsilon_{\text{mem}}$ . Let the insertion query for  $R$  in  $\Pi$  be the following relational calculus query:

$$\{(\mathbf{u}_1, \dots, \mathbf{u}_k) \mid \varphi(\mathbf{u}_1, \dots, \mathbf{u}_k)\}.$$

The FO-formula  $\varphi$  is over  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}} \cup \{\text{Id}^{(1)}\}$ . We write  $\varphi^{\text{split}}$  to denote the modification of  $\varphi$  that is obtained by replacing for each  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$ , each occurrence of an atomic subformula  $R(\mathbf{u}_1, \dots, \mathbf{u}_k)$  by the non-atomic formula  $(\bigvee_{\tau \in \alpha} R_\tau(\mathbf{u}_1, \dots, \mathbf{u}_k))$  where  $\alpha$  is the set of all  $k$ -types. Now, denoting  $\tau = (s_1, \dots, s_k)$ , we define the insertion query for relation  $R_\tau$  in  $\Pi^2$  to be the following relational calculus query:

$$\{(\mathbf{u}_1, \dots, \mathbf{u}_k) \mid \varphi^{\text{split}}(\mathbf{u}_1, \dots, \mathbf{u}_k) \wedge S_1(\mathbf{u}_1) \wedge \dots \wedge S_k(\mathbf{u}_k) \wedge \text{started}()\}$$

where for each  $i \in \{1, \dots, k\}$  we define  $S_i = \text{Adom}$  if  $s_i = \text{adom}$  and  $S_i = \text{Id}$  if  $s_i = \text{id}$ . The deletion query for  $R_\tau$  in  $\Pi^2$  can be defined in a similar way, based on the deletion query of  $R$  in  $\Pi$ . Again, this is similar for a relation  $R_\tau^{(k)}$  with  $R^{(k)} \in \Upsilon_{\text{out}}$ , but with the difference that the deletion query in  $\Pi^2$  will always return empty (because output only accumulates).

Finally, for  $R^{(k)} \in \Upsilon_{\text{out}}^2 = \Upsilon_{\text{out}}$  we define the output query of  $R$  in  $\Pi^2$  to copy the contents of memory relation  $R_\tau$  where  $\tau$  is the  $k$ -type  $(\text{adom}, \dots, \text{adom})$ . This way, the value  $x$  can not be output on input instances that do not contain  $x$  in their active domain (see our earlier assumption). This completes the description of transducer  $\Pi^2$ .

As for notation, for a fact  $\mathbf{f} = R_\tau(a_1, \dots, a_k)$  we write  $\widehat{\mathbf{f}}$  to denote  $R(a_1, \dots, a_k)$ , i.e., to drop the type subscript. The following lemma can now be shown by induction on the structure of FO-formulas.

**LEMMA B.1.** *Let  $I$  be an input instance over  $\Upsilon_{\text{in}} = \Upsilon_{\text{in}}^2$  such that  $x \notin \text{adom}(I)$ . On input  $I$ , on the single-node network  $x$ , let  $J_1$  and  $J_2$  be transducer states for  $\Pi$  and  $\Pi^2$  respectively.*

<sup>5</sup>For easier technical presentation, we assume that relational calculus queries produce tuples instead of facts. These tuples can be easily turned into facts by prepending the correct predicate.

This implies  $\text{Id}(x) \in J_1|_{(\text{sys})}$  and  $\text{Id}(x) \in J_2|_{(\text{sys})}$ . Suppose for each  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$  that

$$J_1|_R = \bigcup_{\tau \in \alpha} \{\widehat{\mathbf{f}} \mid \mathbf{f} \in J_2|_{R_\tau}\}$$

where  $\alpha$  is the set of all  $k$ -types. In words: the relation  $R$  in  $J_1$  is represented exactly by the split  $R$ -relations in  $J_2$ .

Consider some relational calculus query over input schema  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{out}} \cup \Upsilon_{\text{mem}} \cup \{\text{Id}^{(1)}\}$ :

$$\mathcal{C}_1 = \{(\mathbf{u}_1, \dots, \mathbf{u}_k) \mid \varphi(\mathbf{u}_1, \dots, \mathbf{u}_k)\}.$$

Consider the modified relational calculus query:

$$\mathcal{C}_2 = \{(\mathbf{u}_1, \dots, \mathbf{u}_k) \mid \varphi^{\text{split}}(\mathbf{u}_1, \dots, \mathbf{u}_k)\}.$$

We have  $\mathcal{C}(J_1) = \mathcal{C}_2(J_2)$ . □

Let  $\mathcal{T}^2$  denote the transducer network where we run  $\Pi^2$  on  $x$ . Now we have the following property:

LEMMA B.2.  $\mathcal{T}^2$  computes the query  $\mathcal{Q}$  when restricted to inputs not containing  $x$ .

PROOF. Let  $I$  be a database instance over  $\Upsilon_{\text{in}}$  with  $x \notin \text{adom}(I)$ . Let  $\mathcal{R}$  and  $\mathcal{R}^2$  denote the unique runs of  $\mathcal{T}$  and  $\mathcal{T}^2$  on input  $I$  respectively.<sup>6</sup> Let  $\rho_1, \rho_2, \dots$ , and  $\rho'_1, \rho'_2, \dots$ , denote the sequences of configurations of  $\mathcal{R}$  and  $\mathcal{R}^2$  respectively. For  $i \geq 1$  we denote  $\rho_i = (s_i, b_i)$  and  $\rho'_i = (s'_i, b'_i)$ . Using Lemma B.1, it can be shown by induction on  $i \geq 1$  that for each relation  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$  we have

$$s_i(x)|_R = \bigcup_{\tau \in \alpha} \{\widehat{\mathbf{f}} \mid \mathbf{f} \in s'_{i+1}(x)|_{R_\tau}\}.$$

In words: although in  $\Pi^2$  the tuples are divided by their type, they still represent exactly the same tuples as in  $\Pi$ . In this expression, the configuration-index for run  $\mathcal{R}^2$  is offset by 1 because the queries of  $\Pi^2$  first have to wait until relation **started** becomes nonempty.

We are left to show that  $\Pi^2$  outputs  $\mathcal{Q}(I)$ . Let  $R^{(k)} \in \Upsilon_{\text{out}}$ . Let  $\mathbf{f} = R(a_1, \dots, a_k)$  be an output fact produced in run  $\mathcal{R}$ . Because by assumption  $x \notin \text{adom}(I)$  and because the query  $\mathcal{Q}$  is generic, we must have that  $\mathbf{f}$  has  $k$ -type  $\tau = (\text{adom}, \dots, \text{adom})$ . Using the above property, the memory fact  $R_\tau(a_1, \dots, a_k)$  is produced in run  $\mathcal{R}^2$  (and is never deleted). By specification of the output queries in  $\Pi^2$ , the fact  $R(a_1, \dots, a_k)$  is output in  $\mathcal{R}^2$  as well. The other direction is similar. □

**B.3.3. Well-typed formulas.** Consider the *typed* database schema  $\mathcal{E}$  that consists of:

- the relation  $(R, \tau)$  for each  $R^{(k)} \in \Upsilon_{\text{in}}$  where  $\tau$  is  $k$ -type  $(\text{adom}, \dots, \text{adom})$ ; and,
- the relation  $(R_\tau, \tau)$  for each  $R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$  and *each*  $k$ -type  $\tau$ .

Let  $\varphi$  be an FO-formula used in a query of  $\Pi^2$ . Formula  $\varphi$  is over the schema  $\Upsilon_{\text{in}}^2 \cup \Upsilon_{\text{mem}} \cup \{\text{Id}^{(1)}\}$ . When we would ignore the (simple) usage of relation **started** in  $\varphi$ , formula  $\varphi$  is over the schema *untyped*( $\mathcal{E}$ ). Now we can apply Proposition 1 of [Van den Bussche and Cabibbo 1998] to  $\varphi$ , to obtain  $\varphi^{\text{well}}$ , which is a *well-typed* formula over  $\mathcal{E}$ . Formula  $\varphi^{\text{well}}$  computes the same query as  $\varphi$  when applied to a typed database instance over  $\mathcal{E}$ , but importantly,  $\varphi^{\text{well}}$  does not read the relations **Adom** and **Id** directly. Instead, it has variables with the sort *adom* or *id* that range over the active domain of the input instance and  $\{x\}$ , respectively. Since  $\varphi^{\text{well}}$  does not read relation **Adom** anymore, we can also safely remove the occurrence of relation **started** from it.

<sup>6</sup>The runs are unique because there are only heartbeat transitions.

Because formula  $\varphi^{\text{well}}$  uses two sorts of variables, it is not directly usable for a (normal) FO-transducer. We now explain how to remove the *id*-variables from  $\varphi^{\text{well}}$  so that only the *adom*-variables remain, giving us again a (normal) formula with a single sort of variable. So, let  $\varphi$  be an FO-formula used for a memory relation in  $\Pi^2$ , either for insertion or deletion. Abbreviate  $\psi = \varphi^{\text{well}}$ . Let us also define the following sentences: *true* :=  $\forall \mathbf{u}(\mathbf{u} = \mathbf{u})$  and *false* :=  $\exists \mathbf{u}(\mathbf{u} \neq \mathbf{u})$ , where  $\mathbf{u}$  is an *adom*-variable not yet occurring in  $\psi$ . By structural induction, we now convert  $\psi$  to a normal FO-formula  $\psi^\downarrow$ , by keeping only the *adom* variables as follows:

- Suppose  $\psi$  is  $(\mathbf{u} = \mathbf{v})$  with  $\mathbf{u}$  an *adom*-variable and  $\mathbf{v}$  an *id*-variable. We define  $\psi^\downarrow$  as  $(\mathbf{u} = \mathbf{u}) \wedge \textit{false}$ , because *adom*- and *id*-variables can never point to the same value (using our assumption that inputs do not contain value  $x$ ).
- Suppose  $\psi$  is  $(\mathbf{v}_1 = \mathbf{v}_2)$  with  $\mathbf{v}_1$  and  $\mathbf{v}_2$  being *id*-variables. We define  $\psi^\downarrow$  as *true*, because *id*-variables can only point to the same value  $x$ .
- Suppose  $\psi$  is  $(\mathbf{u}_1 = \mathbf{u}_2)$  with  $\mathbf{u}_1$  and  $\mathbf{u}_2$  being *adom*-variables. We define  $\psi^\downarrow$  as  $(\mathbf{u}_1 = \mathbf{u}_2)$ .
- Suppose  $\psi$  is  $R(\mathbf{w}_1, \dots, \mathbf{w}_k)$  with  $R^{(k)} \in \Upsilon_{\text{in}}$ . Then  $R$  has type  $(\textit{adom}, \dots, \textit{adom})$  in  $\mathcal{E}$ . Hence, each variable  $\mathbf{w}_i$  is an *adom*-variable. We define  $\psi^\downarrow$  as  $R(\mathbf{w}_1, \dots, \mathbf{w}_k)$ .
- Suppose  $\psi$  is  $R_\tau(\mathbf{w}_1, \dots, \mathbf{w}_k)$  with  $R_\tau^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}$ . Denote  $\tau = (s_1, \dots, s_k)$ . Let  $\mathbf{u}_1, \dots, \mathbf{u}_n$  be the *adom*-variables of  $\mathbf{w}_1, \dots, \mathbf{w}_k$  in order. By construction of  $\psi$ , for each  $i \in \{1, \dots, k\}$ , the sort of  $\mathbf{w}_i$  is  $s_i$ . We define  $\psi^\downarrow$  as  $R_\tau(\mathbf{u}_1, \dots, \mathbf{u}_n)$ .
- Suppose  $\psi$  is  $\psi_1 \vee \psi_2$ . Let  $\psi_1^\downarrow$  and  $\psi_2^\downarrow$  denote the conversions of  $\psi_1$  and  $\psi_2$  respectively. We define  $\psi^\downarrow$  as  $\psi_1^\downarrow \vee \psi_2^\downarrow$ .
- Suppose  $\psi$  is  $\exists \mathbf{w}(\psi_1)$ . If  $\mathbf{w}$  is an *adom*-variable then we define  $\psi^\downarrow$  as  $\exists \mathbf{w}(\psi_1^\downarrow)$  and otherwise we define  $\psi^\downarrow$  as  $\psi_1^\downarrow$ .
- Suppose  $\psi = \neg \psi_1$ . We define  $\psi^\downarrow$  as  $\neg(\psi_1^\downarrow)$ .

In  $\psi^\downarrow$ , there are no *id*-variables and all *adom*-variables have been preserved. In conclusion, to remove relation *Id* from an FO-formula  $\varphi$  of  $\Pi^2$ , we use the transformation  $(\varphi^{\text{well}})^\downarrow$ . We will use this below.

**B.3.4. Construct new transducer.** We construct a third and last FO-transducer  $\Pi^3$  that computes the query  $\mathcal{Q}$ .

For a type  $\tau$ , we write  $\#\tau$  to denote the number of *adom*-components. We now define the schema  $\Upsilon^3$  of  $\Pi^3$  as

- $\Upsilon_{\text{in}}^3 = \Upsilon_{\text{in}}$ ;  $\Upsilon_{\text{out}}^3 = \Upsilon_{\text{out}}$ ;  $\Upsilon_{\text{msg}}^3 = \emptyset$ ; and,
- $\Upsilon_{\text{mem}}^3 = \{R_\tau^{(\#\tau)} \mid R^{(k)} \in \Upsilon_{\text{mem}} \cup \Upsilon_{\text{out}}, \tau \text{ is a } k\text{-type}\}$ .

Note that  $\Upsilon^3$  is very similar to schema  $\Upsilon^2$ , with the difference that (i) the memory relations in  $\Upsilon^3$  in general have a lower arity to store just the *adom*-values, and (ii) relation **started** is omitted because the relation **Adom** does not need to be computed anymore.

Now we define the queries of  $\Pi^3$ . First, let  $R_\tau^{(l)} \in \Upsilon_{\text{mem}}^3$ . By definition,  $l = \#\tau$ . Let  $k$  be the arity of  $\tau$ . We have  $R_\tau^{(k)} \in \Upsilon_{\text{mem}}^2$ . Let the insertion query for  $R_\tau^{(k)}$  in  $\Pi^2$  be

$$\{(\mathbf{w}_1, \dots, \mathbf{w}_k) \mid \varphi(\mathbf{w}_1, \dots, \mathbf{w}_k)\}.$$

We define the insertion query for  $R_\tau^{(l)}$  in  $\Pi^3$  to be:

$$\{(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid (\varphi^{\text{well}})^\downarrow(\mathbf{u}_1, \dots, \mathbf{u}_n)\}$$

where  $\mathbf{u}_1, \dots, \mathbf{u}_n$  are the *adom*-variables of  $\mathbf{w}_1, \dots, \mathbf{w}_k$  (in order). The deletion query for  $R_\tau^{(l)}$  in  $\Pi^3$  is defined similarly. For  $R^{(k)} \in \Upsilon_{\text{out}}^3$ , we define the output query in  $\Pi^3$  as the one that copies the memory relation  $R_\tau$  to  $R$  with  $\tau$  the  $k$ -type  $(\textit{adom}, \dots, \textit{adom})$ .

Before we can look at the properties of  $\Pi^3$ , we need some additional notation. Consider  $R_\tau^{(l)} \in \Upsilon_{\text{mem}}^3$ . Let  $k$  be the arity of  $\tau$ . We have  $R_\tau^{(k)} \in \Upsilon_{\text{mem}}^2$ . Let  $A \subseteq \{1, \dots, k\}$  be the component indices of  $\tau$  corresponding to an *adom*-variable. Let  $f : A \rightarrow \mathbb{N}$  be the strictly increasing function that maps  $A$  to contiguous integers starting at 1. For example, if  $A = \{3, 6, 7\}$  then  $f = \{3 \mapsto 1, 6 \mapsto 2, 7 \mapsto 3\}$ . Let  $\bar{t}$  be a tuple  $(a_1, \dots, a_l)$ . We write  $\bar{t}^{\uparrow(\tau)}$  to denote the tuple  $(b_1, \dots, b_k)$  where for each  $i \in \{1, \dots, k\}$  we have  $b_i = a_{f(i)}$  if  $i \in A$  and  $b_i = x$  otherwise. Intuitively, we insert the single *id*-value  $x$  back to obtain a  $k$ -tuple. We use this notation for facts as well.

Let  $\mathcal{T}^3$  denote the transducer network obtained by putting  $\Pi^3$  at node  $x$ . The following lemma can again be shown by structural induction on the (well-typed) FO-formulas.

LEMMA B.3. *Let  $I$  be a database instance over  $\Upsilon_{\text{in}} = \Upsilon_{\text{in}}^2 = \Upsilon_{\text{in}}^3$  with  $x \notin \text{adom}(I)$ . On input  $I$ , on the single-node network  $x$ , let  $J_2$  and  $J_3$  be transducer states for  $\Pi^2$  and  $\Pi^3$  respectively. Suppose for each  $R_\tau^{(k)} \in \Upsilon_{\text{mem}}^2$  that*

$$J_2|_{R_\tau} = \{\mathbf{f}^{\uparrow(\tau)} \mid \mathbf{f} \in J_3|_{R_\tau}\}.$$

*Let  $\sigma$  be a type. Consider some relational calculus query over input schema  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{mem}}^2 \cup \{\text{Id}^{(1)}\}$  that produces only tuples of type  $\sigma$ :*

$$\mathcal{C}_2 = \{(\mathbf{w}_1, \dots, \mathbf{w}_k) \mid \varphi(\mathbf{w}_1, \dots, \mathbf{w}_k)\}.$$

*Consider the following transformed query, which is over input schema  $\Upsilon_{\text{in}} \cup \Upsilon_{\text{mem}}^3$ :*

$$\mathcal{C}_3 = \{(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid (\varphi^{\text{well}})^\downarrow(\mathbf{u}_1, \dots, \mathbf{u}_n)\}$$

*where  $\mathbf{u}_1, \dots, \mathbf{u}_n$  are the free *adom*-variables of  $(\varphi^{\text{well}})^\downarrow$ , in the same relative order as they occur in  $\mathbf{w}_1, \dots, \mathbf{w}_k$ . We have  $\mathcal{C}_2(J_2) = \{\bar{t}^{\uparrow(\sigma)} \mid \bar{t} \in \mathcal{C}_3(J_3)\}$ .  $\square$*

We now have the following property:

LEMMA B.4.  $\mathcal{T}^3$  computes the query  $\mathcal{Q}$ .

PROOF. Similarly to the proof of Lemma B.2, we can show by using Lemma B.3 that  $\mathcal{T}^3$  computes the same query as  $\mathcal{T}^2$  when restricted to inputs not containing  $x$ , which (by Lemma B.2) is the original query  $\mathcal{Q}$  restricted to those inputs. At this point, we cannot say yet that  $\mathcal{T}^3$  computes the full query  $\mathcal{Q}$ , i.e., for all inputs on which  $\mathcal{Q}$  is defined. We will show now that this is actually the case.

First, observe that transducer  $\Pi^3$  is oblivious. Indeed,  $\Pi$  does not use **All**, and we have further eliminated the use of **Id** from transducer  $\Pi$  to obtain  $\Pi^3$ . Then, Lemma 4.2 tells us that  $\mathcal{T}^3$  computes a *generic* query  $\mathcal{Q}'$ . Let  $I$  be an instance over  $\Upsilon_{\text{in}}$  with possibly  $x \in \text{adom}(I)$ . There is another instance  $J$  over  $\Upsilon_{\text{in}}$  and a permutation  $h$  of **dom** such that  $h(J) = I$  and  $x \notin \text{adom}(J)$ . As seen above, we have  $\mathcal{Q}'(J) = \mathcal{Q}(J)$  and thus  $h(\mathcal{Q}'(J)) = h(\mathcal{Q}(J))$ . By genericity of both  $\mathcal{Q}'$  and  $\mathcal{Q}$  we then have  $\mathcal{Q}'(h(J)) = \mathcal{Q}(h(J))$  and thus  $\mathcal{Q}'(I) = \mathcal{Q}(I)$ . In conclusion, the transducer network  $\mathcal{T}^3$  computes the same query  $\mathcal{Q}$  as the original transducer network  $\mathcal{T}$  but without reading the relations **Id** and **All**.  $\square$

## References

- ABITEBOUL, S., BIENVENU, M., GALLAND, A., ET AL. 2011. A rule-based language for Web data management. In *Proceedings 30th ACM Symposium on Principles of Database Systems*. ACM Press, 293–304.
- ABITEBOUL, S., HULL, R., AND VIANU, V. 1995. *Foundations of Databases*. Addison-Wesley.
- ABITEBOUL, S. AND SIMON, E. 1991. Fundamental properties of deterministic and nondeterministic extensions of Datalog. *Theor. Comput. Sci.* 78, 137–158.
- ABITEBOUL, S. AND VIANU, V. 1991a. Datalog extensions for database queries and updates. *J. Comput. Syst. Sci.* 43, 1, 62–124.

- ABITEBOUL, S. AND VIANU, V. 1991b. Generic computation and its complexity. In *Proceedings 23rd ACM Symposium on the Theory of Computing*. 209–219.
- ABITEBOUL, S. AND VIANU, V. 1995. Computing with first-order logic. *J. Comput. Syst. Sci.* 50, 2, 309–335.
- ABITEBOUL, S., VIANU, V., ET AL. 2000. Relational transducers for electronic commerce. *J. Comput. Syst. Sci.* 61, 2, 236–269.
- AFRATI, F., COSMADAKIS, S., AND YANNAKAKIS, M. 1995. On Datalog vs polynomial time. *J. Comput. Syst. Sci.* 51, 2, 177–196.
- ALVARO, P., CONWAY, N., HELLERSTEIN, J., AND MARCZAK, W. 2011a. Consistency analysis in Bloom: A CALM and collected approach. In *Proceedings 5th Biennial Conference on Innovative Data Systems Research*. www.cidrdb.org, 249–260.
- ALVARO, P., MARCZAK, W., ET AL. 2011b. Dedalus: Datalog in time and space. In *Datalog Reloaded: First International Workshop, Datalog 2010*, O. de Moor, G. Gottlob, T. Furche, and A. Sellers, Eds. Lecture Notes in Computer Science Series, vol. 6702. 262–281.
- AMELOOT, T., NEVEN, F., AND VAN DEN BUSSCHE, J. 2011. Relational transducers for declarative networking. In *Proceedings 30th ACM Symposium on Principles of Database Systems*. ACM Press, 283–292.
- AMELOOT, T. AND VAN DEN BUSSCHE, J. 2012a. Deciding eventual consistency for a simple class of relational transducers. In *Proceedings of the 15th International Conference on Database Theory*. ACM Press, 86–98.
- AMELOOT, T. AND VAN DEN BUSSCHE, J. 2012b. On the cron conjecture. In *Proceedings of the Second International Workshop, Datalog 2.0, Vienna, Austria*. Springer. To appear.
- APT, K., FRANCEZ, N., AND KATZ, S. 1988. Appraising fairness in languages for distributed programming. *Distributed Computing* 2, 226–241.
- ATTIYA, H. AND WELCH, J. 2004. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. Wiley.
- BLASS, A., GUREVICH, Y., AND VAN DEN BUSSCHE, J. 2002. Abstract state machines and computationally complete query languages. *Information and Computation* 174, 1, 20–36.
- DEUTSCH, A., HULL, R., PATRIZI, F., AND VIANU, V. 2009. Automatic verification of data-centric business processes. In *Proceedings 12th International Conference on Database Theory*.
- DEUTSCH, A., SUI, L., AND VIANU, V. 2007. Specification and verification of data-driven Web applications. *J. Comput. Syst. Sci.* 73, 3, 442–474.
- DEUTSCH, A., SUI, L., VIANU, V., AND ZHOU, D. 2006. Verification of communicating data-driven Web services. In *Proceedings 25th ACM Symposium on Principles of Database Systems*. ACM Press, 90–99.
- EUGSTER, P., GUERRAoui, R., KERMARREC, A., AND MASSOULIÉ, L. 2004. From epidemics to distributed computing. *IEEE Computer* 37, 60–67.
- FRANCEZ, N. 1986. *Fairness*. Springer-Verlag New York, Inc., New York, NY, USA.
- GRAY, J. 1978. Notes on data base operating systems. In *Operating Systems—An Advanced Course*, M. Flynn et al., Eds. Lecture Notes in Computer Science Series, vol. 60. Springer-Verlag, 393–481.
- GRUMBACH, S. AND WANG, F. 2010. Netlog, a rule-based language for distributed programming. In *Proceedings 12th International Symposium on Practical Aspects of Declarative Languages*, M. Carro and R. Peña, Eds. Lecture Notes in Computer Science Series, vol. 5937. 88–103.
- HELLERSTEIN, J. 2010a. Datalog redux: experience and conjecture. Video available (under the title “The Declarative Imperative”) from <http://db.cs.berkeley.edu/jmh/>. PODS 2010 keynote.
- HELLERSTEIN, J. 2010b. The declarative imperative: experiences and conjectures in distributed logic. *SIGMOD Record* 39, 1, 5–19.
- LAMPORT, L. 2000. Fairness and hyperfairness. *Distributed Computing* 13, 239–245.
- LOO, B. ET AL. 2009. Declarative networking. *Commun. ACM* 52, 11, 87–95.
- LOO, B., CONDIE, T., ET AL. 2006. Declarative networking: language, execution and optimization. In *SIGMOD Conference*, S. Chaudhuri, V. Hristidis, and N. Polyzotis, Eds. ACM, 97–108.
- NIGAM, V., JIA, L., LOO, B., AND SCEDROV, A. 2012. Maintaining distributed logic programs incrementally. *Computer Languages, Systems & Structures* 38, 2, 158–180.
- SPIELMANN, M. 2003. Verification of relational transducers for electronic commerce. *J. Comput. Syst. Sci.* 66, 1, 40–65.
- VAN DEN BUSSCHE, J. AND CABIBBO, L. 1998. Converting untyped formulas into typed ones. *Acta Informatica* 35, 8, 637–643.
- ZINN, D., GREEN, T., AND LUDAESCHER, B. 2012. Win-move is coordination-free. In *Proceedings of the 15th International Conference on Database Theory*. ACM Press, 99–113.