

Relational Transducers for Declarative Networking

[Extended Abstract]

Tom J. Ameloot^{*}
Hasselt University
Agoralaan Gebouw D
3590 Diepenbeek, Belgium
tom.ameloot@uhasselt.be

Frank Neven
Hasselt University
Agoralaan Gebouw D
3590 Diepenbeek, Belgium
frank.neven@uhasselt.be

Jan Van den Bussche
Hasselt University
Agoralaan Gebouw D
3590 Diepenbeek, Belgium
jan.vandenbussche@uhasselt.be

ABSTRACT

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.

1. INTRODUCTION

Declarative networking [17] 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 [14, 15], 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 (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

^{*}PhD Fellow of the Fund for Scientific Research, Flanders (FWO). We acknowledge the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under the FET-Open grant agreement FOX, number FP7-ICT-233599.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODS'11, June 13–15, 2011, Athens, Greece.

Copyright 2011 ACM 978-1-4503-0660-7/11/06 ...\$10.00.

have been lacking so far. Moreover, it also requires a formal model of distributed computation.

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 [6, 12, 13, 19]. Relational transducers are firmly grounded in the theory of database queries [4, 5] and also have close connections with Abstract State Machines [10]. 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, nonobvious 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” [15]: 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 in Datalog, for example, then \mathcal{L} can just be the 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 a Datalog 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.

Finally, the present work also lead us to think about the computational power of the language Dedalus [8], the Datalog extension used by Alvaro and Hellerstein. We will show that this language is quite powerful by establishing a monotone simulation of arbitrary Turing machines.

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. Section 6 looks into the CALM conjecture. Section 7 presents some further results. Section 8 compares our results to the language Dedalus. Section 9 is the conclusion.

In this extended abstract, proofs are mainly given on an informal level.

2. PRELIMINARIES

We recall some basic notions from the theory of database queries [1].

A *database schema* is a finite set \mathcal{S} of relation names, each with an associated arity (a natural number). Assume some infinite universe \mathbf{dom} of atomic *data elements*. An *instance* of a database schema \mathcal{S} is an assignment I of finite relations on \mathbf{dom} to the relation names of \mathcal{S} , such that when R has arity k then $I(R)$ is a k -ary relation. Equivalently, we can view an instance as a set of *facts* over \mathcal{S} , where a fact is an expression of the form $R(a_1, \dots, a_k)$ with $a_1, \dots, a_k \in \mathbf{dom}$ and $R \in \mathcal{S}$ of arity k . The *active domain* of an instance I , denoted by $\text{adom}(I)$, is the set of data elements occurring in I .

A *k -ary query over \mathcal{S}* is a partial function Q mapping instances of \mathcal{S} to k -ary relations on \mathbf{dom} such that for each I on which Q is defined, the following two conditions hold: (i) $Q(I)$ is a k -ary relation on $\text{adom}(I)$; and (ii) Q is also defined on the isomorphic instance $h(I)$, for each permutation h of \mathbf{dom} , and $Q(h(I)) = h(Q(I))$. A query Q is *monotone* if for any two instances $I \subseteq J$, if $Q(I)$ is defined then so is $Q(J)$, and $Q(I) \subseteq Q(J)$.

We assume familiarity with *first-order logic (FO)* as a basic database query language. An FO formula $\varphi(x_1, \dots, x_k)$ expresses the k -ary query defined by $\varphi(I) = \{(a_1, \dots, a_k) \in \text{adom}(I)^k \mid (\text{adom}(I), I) \models \varphi[a_1, \dots, a_k]\}$. Note that we evaluate FO formulas on instances under the active-domain semantics. The resulting query language is equivalent in expressive power to the relational algebra, as well as to recursion-free Datalog with negation.

We will also consider the query languages Datalog, stratified Datalog (with negation), and *while*. Datalog and stratified Datalog are well known; *while* is the query language obtained from FO by adding assignment statements and while-loops. Finally, we recall that there exist quite elegant *computationally complete* query languages in which every partial computable query can be expressed.

2.1 Relational transducers

A *transducer schema* is a tuple $(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{sys}}, \mathcal{S}_{\text{msg}}, \mathcal{S}_{\text{mem}}, k)$ consisting of four disjoint database schemas and an *output*

arity k . Here, ‘in’ stands for ‘input’; ‘sys’ stands for ‘system’; ‘msg’ stands for ‘message’; and ‘mem’ stands for ‘memory’.

An *abstract relational transducer* (or just transducer for short) over this schema is a collection of queries $\{Q_{\text{snd}}^R \mid R \in \mathcal{S}_{\text{msg}}\} \cup \{Q_{\text{ins}}^R \mid R \in \mathcal{S}_{\text{mem}}\} \cup \{Q_{\text{del}}^R \mid R \in \mathcal{S}_{\text{mem}}\} \cup \{Q_{\text{out}}\}$, where

- every query is over the combined database schema $\mathcal{S}_{\text{in}} \cup \mathcal{S}_{\text{sys}} \cup \mathcal{S}_{\text{msg}} \cup \mathcal{S}_{\text{mem}}$;
- the arity of each Q_{snd}^R , each Q_{ins}^R and each Q_{del}^R equals the arity of R ; and
- the arity of Q_{out} equals the output arity k .

Here, ‘snd’ stands for ‘send’; ‘ins’ stands for ‘insert’; ‘del’ stands for ‘delete’; and ‘out’ stands for ‘output’.

A *state* of the transducer is an instance of the combined schema $\mathcal{S}_{\text{in}} \cup \mathcal{S}_{\text{sys}} \cup \mathcal{S}_{\text{mem}}$. Intuitively, a state just consists of some input relations, some system relations (we will make these precise in the next section), and some stored relations that constitute the memory of the transducer.

A *message instance* is, plainly, an instance of \mathcal{S}_{msg} . Such a message instance can stand for a set of messages (facts) received by the transducer, but can as well stand for a set of messages sent by the transducer. It will always be clear from the context which of the two meanings we have.

Let Π be a transducer. A *transition* of Π is a five-tuple $(I, I_{\text{rcv}}, J_{\text{snd}}, J_{\text{out}}, J)$, also denoted as $I, I_{\text{rcv}} \xrightarrow{J_{\text{out}}} J, J_{\text{snd}}$, where I and J are states, I_{rcv} and J_{snd} are message instances, and J_{out} is a k -ary relation such that

- every query of Π is defined on $I' = I \cup I_{\text{rcv}}$;
- J agrees with I on \mathcal{S}_{in} and \mathcal{S}_{sys} ;
- $J_{\text{snd}}(R)$, for each $R \in \mathcal{S}_{\text{msg}}$, equals $Q_{\text{snd}}^R(I')$;
- J_{out} equals $Q_{\text{out}}(I')$;
- $J(R)$, for each $R \in \mathcal{S}_{\text{mem}}$, equals

$$\begin{aligned} & (Q_{\text{ins}}^R(I') \setminus Q_{\text{del}}^R(I')) \\ & \cup (Q_{\text{ins}}^R(I') \cap Q_{\text{del}}^R(I') \cap I(R)) \\ & \cup (I(R) \setminus (Q_{\text{ins}}^R(I') \cup Q_{\text{del}}^R(I'))). \end{aligned}$$

The intuition behind the instance I' is that Π sees its input, system and memory relations, plus its received messages. The transducer does not modify the input and system relations. The transducer computes new tuples that can be sent out as messages; this is the instance J_{snd} . The transducer also outputs some tuples; this is the relation J_{out} . These outputs cannot later be retracted! Finally the transducer updates its memory by inserting and deleting some tuples for every memory relation. The intimidating update formula merely expresses that conflicting inserts/deletes are ignored [13, 19]. Note that an assignment $R := Q$ can be expressed by using Q for Q_{ins}^R and R for Q_{del}^R .

Note also that transitions are *deterministic*: if $I, I_{\text{rcv}} \xrightarrow{J_{\text{out}}} J, J_{\text{snd}}$ and $I, I_{\text{rcv}} \xrightarrow{J'_{\text{out}}} J', J'_{\text{snd}}$, then $J'_{\text{snd}} = J_{\text{snd}}$; $J'_{\text{out}} = J_{\text{out}}$; and $J' = J$.

An abstract relational transducer as defined above is just a collection of queries. If we want to write down a transducer then we will of course use some query language to express

these queries. By default, we use FO as the query language. More generally, for any query language \mathcal{L} we can consider the language of \mathcal{L} -transducers consisting of all transducers whose queries are expressed in \mathcal{L} . Because we are going to place transducers in networks, we can think of \mathcal{L} as the language that individual peers use locally. For example, in the language Dedalus [8], the local language is stratified Datalog.

3. TRANSDUCER NETWORKS

Proviso 1. From now on we will only consider transducers where the system schema \mathcal{S}_{sys} consists of the two unary relation names *Id* and *All*.

A *network* is a finite, *connected*, undirected graph over a set of vertices $V \subset \mathbf{dom}$. We refer to the vertices as *nodes*. Note that nodes belong to the universe \mathbf{dom} of atomic data elements; indeed we are going to allow that nodes are stored in relations. We stress again that a network must be connected. This is important to make it possible for flow of information to reach every node.

A *transducer network* is a pair (\mathcal{N}, Π) where \mathcal{N} is a network and Π is a transducer. The general idea is that a copy of Π is running on every node. A database instance is distributed over the input relations of the different nodes. Relation *Id* contains the node identifier where the transducer is running, and relation *All* is the same at all nodes and contains the set of all nodes. When a node v sends a set of facts as messages, these facts are added to the message buffers of v 's neighbors. Nodes receive facts one by one in arbitrary order; messages are not necessarily received in the order they have been sent. A similar situation can happen in the Internet with subsequent TCP/IP connections between the same two nodes, where an earlier connection might be slower than a later one. Moreover, nodes regularly receive a “heartbeat” message which allows them to transition even when no message is read.

We proceed to define the possible runs of a transducer network more formally. A *configuration* of the system is a pair $\gamma = (\text{state}, \text{buf})$ of mappings where

- *state* maps every node v to a state $I = \text{state}(v)$ of Π , so that $I(\text{Id}) = \{v\}$, and $I(\text{All}) = V$ (the set of all nodes of \mathcal{N}).
- *buf* maps every node to a finite multiset of facts over \mathcal{S}_{msg} .

Thus, the system relations *Id* and *All* give the transducer knowledge about the node where it is running and about the other nodes in the network. We will discuss the use and necessity of these relations extensively.

A *general transition* of the system is the transformation of one configuration to another where some node v reads and removes some message instance I_{rcv} from its input buffer, makes a local transition, and sends the resulting message instance J_{snd} to its neighbors. Formally, a general transition is a tuple $\tau = (\gamma_1, v, I_{\text{rcv}}, J_{\text{out}}, \gamma_2)$, also denoted as $\gamma_1 \xrightarrow[v, I_{\text{rcv}}]{J_{\text{out}}} \gamma_2$, where $\gamma_i = (\text{state}_i, \text{buf}_i)$ for $i = 1, 2$ is a configuration, v is a node, and $I_{\text{rcv}} \subseteq \text{buf}_1(v)$ (multiset containment), such that:

- $\text{state}_2(v') = \text{state}_1(v')$ for every node $v' \neq v$.

- There exists J_{snd} such that $\text{state}_1(v), I_{\text{rcv}} \xrightarrow{J_{\text{out}}} \text{state}_2(v), J_{\text{snd}}$. We call J_{out} the *output* of the transition and denote it also by $\text{out}(\tau)$. Note that, since individual transducer transitions are deterministic, J_{out} and J_{snd} are uniquely determined by $\text{state}_1(v)$ and I_{rcv} .
- $\text{buf}_2(v) = \text{buf}_1(v) \setminus I_{\text{rcv}}$ (multiset difference).
- $\text{buf}_2(v') = \text{buf}_1(v')$ for every node $v' \neq v$ that is not a neighbor of v .
- $\text{buf}_2(v') = \text{buf}_1(v') \cup J_{\text{snd}}$ for every node v' that is a neighbor of v . Note we are using multiset union here.

We will, however, not use transitions in their most general form but only in two special forms:

Heartbeat transition: is of the form $\gamma_1 \xrightarrow[v, \emptyset]{J_{\text{out}}} \gamma_2$. So, some node v transitions without reading any message.

Delivery transition: is of the form $\gamma_1 \xrightarrow[v, \{f\}]{J_{\text{out}}} \gamma_2$. So, some node v reads a single fact f from its received message buffer.

We only consider these two forms because they appear to be the most elementary. We are not sure it is realistic to assume that entire message instances can be read in one transition. Therefore we limit message reading to a single fact. Heartbeat transitions ensure that nodes can transition even if their message buffer is empty.

For any two configurations γ_1 and γ_2 we simply write $\gamma_1 \rightarrow \gamma_2$ to denote that the system can transition from γ_1 to γ_2 either by some heartbeat transition or by some delivery transition. A *run* of the system now is an infinite sequence $(\tau_n)_n$ of heartbeat or delivery transitions such that for each n , if τ_n is $\gamma_n \rightarrow \gamma_{n+1}$ then τ_{n+1} is of the form $\gamma_{n+1} \rightarrow \gamma_{n+2}$. In other words, each transition τ_n with $n > 0$ works on the result configuration of the previous transition.

The *output* of a run ρ is then defined as $\text{out}(\rho) = \bigcup_n \text{out}(\tau_n)$. We note the following:

PROPOSITION 1. *For every run $\rho = (\tau_n)_n$ there exists a natural number m such that $\text{out}(\rho) = \bigcup_{n=0}^m \text{out}(\tau_n)$. The number m is called a quiescence point for ρ .*

Indeed, since the initial configuration contains only a finite number of distinct atomic data elements, and a local query cannot invent new data elements, only a finite number of distinct output tuples are possible.

In the language Dedalus [8], new data elements are invented in the form of increasing timestamps, so the above proposition does not hold for Dedalus. It would be interesting to investigate a version of transducer networks with timestamps, or with object-creating local queries.

4. EXPRESSING QUERIES WITH TRANSDUCER NETWORKS

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-topology independence*.

An instance I of the input part of the transducer schema, \mathcal{S}_{in} , can be distributed over the input relations of the nodes

on the network. Formally, for any instance I of \mathcal{S}_{in} , a *horizontal partition of I on the network \mathcal{N}* is a function H that maps every node v to a subset of I , such that $I = \bigcup_v H(v)$. The *initial configuration for H* is a configuration $(\text{state}, \text{buf})$ such that:

- $\text{buf}(v) = \emptyset$ and $\text{state}(v)(R) = \emptyset$ for each node v and every $R \in \mathcal{S}_{\text{mem}}$. So, each node starts with an empty message buffer and empty memory;
- For each node v , the restriction of $\text{state}(v)$ to the input schema \mathcal{S}_{in} equals $H(v)$.

A *run on H* is a run that starts in the initial configuration for H .

We also need the notion of *fair run*. A run is fair if every node does a heartbeat transition infinitely often, and every fact in every message buffer is eventually taken out by a delivery transition. We omit the obvious formalization.

We now say that a transducer network (\mathcal{N}, Π) is *consistent* if for every instance I of \mathcal{S}_{in} , all fair runs on all possible horizontal partitions of I have the same output. Naturally, a consistent transducer network is said to *compute* a query Q over \mathcal{S}_{in} if for every instance I of \mathcal{S}_{in} on which Q is defined, every fair run on any horizontal partition of I outputs $Q(I)$.

Example 1. Let us see a simple example of a network that is *not* consistent. Consider a network with at least two nodes (indeed if the network consists of a single node the transducer runs all by itself; no messages are delivered and there is only one possible run). The input is a set S of data elements. Each node sends its part of S to its neighbors. Also, each node outputs the first element it receives and outputs no further elements. This is easily programmed on an FO-transducer. When there are at least two nodes and at least two elements in S , different runs may deliver the elements in different orders, so different outputs can be produced, even for the same horizontal partition. \square

Example 2. For a simple example of a consistent network, let the input be a binary relation S . Each node outputs the identical pairs from its part of the input. No messages are sent. This network computes the equality selection $\sigma_{\mathbb{S}_1 = \mathbb{S}_2}(S)$.

An example of a consistent transducer network that involves communication, computes the transitive closure of S in a distributed fashion in the well-known way [17]. We present here, a naive, unoptimized version. Each node sends its part of the input to its neighbors. Each node also sends all tuples it receives to its neighbors. In this way the input is flooded to all nodes. Each node accumulates the tuples it receives in a memory relation R . Finally, each node maintains a memory relation T in which we repeatedly insert $S \cup R \cup T \cup (T \circ T)$ (here \circ stands for relational composition). This relation T is also output. Thanks to the monotonicity of the transitive closure, this transducer network is consistent. \square

We are mainly interested in the case where the query can be correctly computed by the distributed transducer program, regardless of the network topology. For example, the transitive closure computation from Example 2 is independent of the network topology (as long as the network is connected, but we are requiring that of all networks).

Formally, a transducer Π is *network-topology independent* if for every network \mathcal{N} , the system (\mathcal{N}, Π) is consistent, and regardless of \mathcal{N} computes the same query Q . We say that Q can be *distributedly computed* by Π .

Example 3. For a simple example of a transducer that is consistent for every network topology, but that is not network-topology independent, consider again as input a set S distributed over the nodes. Each node sends its input to its neighbors and also sends the elements it receives to its neighbors. Each node only outputs the elements it receives. On any network with at least two nodes, the identity query is computed, but on the network with a single node, the empty query is computed. \square

In order to state a few first results in Theorem 3, we introduce the following terminology.

Oblivious transducer: does not use the relations *Id* and *All*. Intuitively, the transducer program is unaware of the context in which it is running.

Inflationary transducer: does not do deletions, i.e., each deletion query returns empty on all inputs.

Monotone transducer: uses only monotone local queries.

LEMMA 2. 1. *There is an inflationary FO-transducer such that, on any network, starting on any horizontal partition of any instance I of \mathcal{S}_{in} , any 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 (implemented by a nullary memory relation) is true. Moreover, the flag Ready does not become true at a node before that node has the entire instance in its memory.*

2. *There is an oblivious, inflationary, monotone FO-transducer that accomplishes the same as the previous one, except for the flag Ready.*

3. *A query is expressible in the language ‘while’ if and only if it is computable by an FO-transducer on a single-node network.*

PROOF. For (1), a multicast protocol [9] is implemented. Every node v sends out all the facts in its local input relations, but each fact is tagged with the id of v in an extra coordinate (using relation *Id*). Every node also forwards all input facts it receives, and stores received facts in memory. Moreover, for every input fact received, every node sends an acknowledge fact, additionally tagged with its own identifier. Every node v keeps a record of for which of its local input facts it has received an ack from which node. When v has received an ack from v' for every local input fact, it sends out a message *done*(v, v') meant for v' . When a node has received *done* from every node (which can be checked using the relation *All*), it knows it is ready. No deletions are necessary.

For (2), the program is much simpler. All nodes simply send out their local input facts and forward any message they receive. In any fair run, eventually all nodes will have received all input facts. Relations *Id* and *All* are not needed.

For (3), on a one-node network there are only heartbeat transitions. A *while* program can be simulated by iterated heartbeats using well-known techniques [3]. Conversely, it is

clear that a one-node transducer network can be simulated by a *while* program. The only difficulty is that the transducer keeps running indefinitely whereas the *while* program is supposed to stop. Using the technique described by Abiteboul and Simon [2], however, the program can detect that it is in an infinite loop. This implies that the transducer has repeated a state and will output no more new tuples. \square

- THEOREM 3.**
1. *Every query can be distributedly computed by some abstract transducer. In particular, if \mathcal{L} is computationally complete, every partial computable query can be distributedly computed by an \mathcal{L} -transducer.*
 2. *Every monotone query can be distributedly computed by an oblivious, inflationary, monotone abstract transducer. In particular, if \mathcal{L} is computationally complete, every partial computable monotone query can be distributedly computed by an oblivious \mathcal{L} -transducer.*
 3. *A query is expressible in the language ‘while’ if and only if it can be distributedly computed by an FO-transducer.*
 4. *Every monotone query expressible in while can be distributedly computed by an oblivious FO-transducer.*
 5. *A query is in Datalog if and only if it can be distributedly computed by an oblivious, inflationary, nonrecursive-Datalog-transducer.*

PROOF. For (1), to distributedly compute any query Q , we first run the transducer from Lemma 2(1) to obtain the entire input instance. Then we apply and output Q .

For (2), the idea is the same, but we now use the transducer from Lemma 2(2). We continuously apply Q to the part of the input instance already received, and output the result. Since Q is monotone, no incorrect tuples are output.

For (3) we only still have to argue the only-if implication. We first run the transducer from Lemma 2(1) to obtain the entire input instance. Then every node can act as if it is on its own, ignoring any remaining incoming messages and simulate the *while*-program.

For (4) the idea is the same as in (2). We receive input tuples and store them in memory. We continuously recompute the *while*-program, starting afresh every time a new input fact comes in. We use deletion to start afresh. Since the query is monotone, no incorrect tuples are output.

For (5) the idea for the only-if implication is again the same as in (2). We receive input tuples and apply continuously the T_P -operator of the Datalog program. By the monotone nature of Datalog evaluation, deletions are not needed, so the transducer is inflationary. The Datalog program for the if-implication is obtained by taking together the rules of all update queries Q_{ins}^R and the output query Q_{out} . The answer predicate of Q_{out} is the global answer predicate. The answer predicate of each Q_{ins}^R is replaced by R , so that we obtain a recursive program. \square

Without proof we note that the transducer from Lemma 2(1) can actually be implemented to use only unions of conjunctive queries with negation (UCQ \neg). By simulating FO queries by fixed compositions of UCQ \neg , we obtain: (proof omitted)

PROPOSITION 4. *Every (monotone) query that can be distributedly computed by an FO-transducer can be distributedly computed by an (oblivious) UCQ \neg -transducer.*

To conclude, we remark that in a transducer network of at least two nodes, each node can establish a linear order on the active domain, by first collecting all input tuples, then sending out all elements of the active domain, forwarding messages and storing the elements that are received back in the order they are received. But such a transducer is not truly network-topology independent, as it does not work in the same way on a one-node network. At any rate, by the well-known characterization of PSPACE [1], we obtain:

COROLLARY 5. *On any network with at least two nodes, every PSPACE query can be computed by an FO-transducer.*

5. COORDINATION

The CALM conjecture hinges on an intuitive notion of “coordination-freeness” of certain distributed computations. For some tasks, coordination is required to reach consistency across the network. Two-phase commit is a classical example. The multicast protocol used in Lemma 2(1) also requires heavy coordination. Since coordination typically blocks distributed computations, it is good to understand precisely when it can be avoided. This is what the CALM conjecture is all about.

It seems hard to give a definitive formalization of coordination. 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 CALM conjecture will turn out to hold.

Formally, consider a network-topology independent transducer Π and a network \mathcal{N} . We call Π *coordination-free* on \mathcal{N} if for every instance I of \mathcal{S}_{in} , there exists a horizontal partition H of I on \mathcal{N} and a run ρ of (\mathcal{N}, Π) on H , in which a quiescence point (Proposition 1) is already reached by only performing heartbeat transitions. Intuitively, if the horizontal partition is “right”, then no communication is required to correctly compute the query. Finally we call Π *coordination-free* if it is coordination-free on any network.

Example 4. Consider again the transitive closure computation from Example 2. When every node already has the full input, they can each individually compute the transitive closure. Hence this transducer is coordination-free.

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. A (contrived) example of a coordination-free transducer that needs communication even if each node has the entire input is the following. The input, distributed over the nodes as usual, consists of two sets A and B , and the query is to determine if at least one of A and B are nonempty. If the network has only one node (which can be tested by looking at the *All* relation), the transducer simply outputs the answer to the query. Otherwise, it first tests if its local input fragments A and B are both nonempty. If yes, nothing is output, but the value ‘true’ (encoded by the empty tuple) is sent out. Any node that receives the message ‘true’ will output it. When A or B is empty locally, the transducer simply outputs the desired output directly. The transducer is coordination-free, since if we take care to

have at least one node that knows A , and another node that knows B , but no node knows both, then the query can be computed without communication. When A and B are both nonempty, however, a run on the horizontal partition where every node has the entire input will not reach quiescence without communication. \square

Example 5. A simple example of a transducer that is not coordination-free, i.e., requires communication, is the one that computes the emptiness query on an input set S . Since every node can have a part of the input, the nodes must coordinate with each other to be certain that S is empty at every node. Every node sends out its identifier (using the relation Id) on condition that its local relation S is empty. Received messages are forwarded, so that if S is globally empty, eventually all nodes will have received the identifiers of *all* nodes, which can be checked using the relation All . When this has happened the transducer at each node outputs ‘true’. \square

Coordination-freeness is undecidable for FO-transducers, but 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. Specifically, recall that an *oblivious* transducer is one that does not use the system relations Id and All . For now we observe:

PROPOSITION 6. *Every network-topology independent, oblivious transducer is coordination-free.*

PROOF. Let Π be a network-topology independent, oblivious transducer. Let Q be the query distributedly computed by Π . On a one-node network and given any input instance I , transducer Π reaches quiescence and outputs $Q(I)$ by doing only heartbeat transitions. Now consider any other network, any instance I over S_{in} , and the horizontal partition H that places the entire I at every node. Since Π is oblivious, nodes cannot detect that they are on a network with multiple nodes unless they communicate. So, by doing only heartbeat transitions initially, every node will act the same as if in a one-node network and will already output the entire $Q(I)$. \square

6. THE CALM CONJECTURE

CONJECTURE 1. *(CALM conjecture [15]) A program has an eventually consistent, coordination-free execution strategy if and only if it is expressible in (monotonic) Datalog.*

It is not specified what is meant by “program” or “strategy”; here, we will take these terms to mean “query” and “distributedly computable by a transducer”, respectively. The term “eventually consistent” is then formalized by our notions of consistency and network-topology independence. Under this interpretation, the conjecture becomes “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 the conjecture out of the way. It surely holds, and many versions of it are already contained in Theorem 3. That theorem talks about oblivious transducers, but we have seen in Proposition 6 that these are coordination-free.

As to the only-if side, the explicit mention of Datalog is a bit of a nuisance. Datalog is limited to polynomial

time whereas there certainly are monotone queries outside PTIME. This continues to hold for queries expressible in the language Dedalus that Alvaro and Hellerstein use; we will show this in Section 8. We also mention the celebrated paper [7] where Afrati, Cosmadakis and Yannakakis show that even within PTIME there exist queries that are monotone but not expressible in Datalog.

We will see in Corollary 9(3) how a Datalog version of the CALM conjecture can be obtained. Datalog aside, however, the true emphasis of the CALM conjecture clearly lies in the monotonicity aspect. Indeed we confirm it in this sense:

THEOREM 7. *Every query that is distributedly computed by a coordination-free transducer is monotone.*

PROOF. Let Q be the query distributedly computed by the coordination-free transducer Π . Let $I \subseteq J$ be two input instances and let $t \in Q(I)$. We must show $t \in Q(J)$. Consider a network \mathcal{N} with at least two nodes. Since Π is coordination-free, there exists a horizontal partition H of I such that $Q(I)$ is already output distributedly over the nodes, by letting the nodes do only heartbeat transitions. Let v be a node where t is output. Let v' be a node different from v and consider the horizontal partition H' of J where $H'(v) = H(v)$ and $H'(v') = H(v') \cup (J \setminus I)$. The partial run on H where v first does only heartbeat transitions until t is output, is also a partial run on H' . This partial run can be extended to a fair run, so t is output by some fair run of (\mathcal{N}, Π) on H' . Since (\mathcal{N}, Π) is consistent, t will also be output in any other fair run on any horizontal partition of J . Hence, t belongs to the query computed by (\mathcal{N}, Π) applied to J . Moreover, Π is network-topology independent, so t belongs to $Q(J)$. \square

COROLLARY 8 (CALM PROPERTY). *The following are equivalent for any query Q :*

1. Q can be distributedly computed by a transducer that is coordination-free.
2. Q can be distributedly computed by a transducer that is oblivious.
3. Q is monotone.

PROOF. Theorem 3 yields (3) \Rightarrow (2); Proposition 6 yields (2) \Rightarrow (1); Theorem 7 yields (1) \Rightarrow (3). \square

Similarly we obtain the following versions of the CALM property:

COROLLARY 9. *The following groups of statements are equivalent for any query Q :*

1. Let \mathcal{L} be a computationally complete query language.
 - (a) Q can be distributedly computed by a coordination-free \mathcal{L} -transducer.
 - (b) Q can be distributedly computed by an oblivious \mathcal{L} -transducer.
 - (c) Q is monotone and partial computable.
2.
 - (a) Q can be distributedly computed by a coordination-free FO-transducer.
 - (b) Q can be distributedly computed by an oblivious FO-transducer.

- (c) Q is monotone and expressible in the language ‘while’.
3. (a) Q can be distributedly computed by a coordination-free nonrecursive-Datalog-transducer.
 - (b) Q can be distributedly computed by a nonrecursive-Datalog-transducer that is oblivious.
 - (c) Q is expressible in Datalog.

7. 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 have only Id or only All ? As to coordination-freeness, it is readily verified that the argument given in the proof of Proposition 6 still works in the presence of Id . It does not work in the presence of All , and indeed we have the following counterexample.

Example 6. We describe a transducer that is network-topology independent, does not use Id , but that is not coordination-free. The query expressed is simply the identity query on a set S . The transducer can detect whether he is alone in the network by looking at the relation All . If so, he simply outputs the result. If he is not alone, he sends out a ping message. Only upon receiving a ping message he outputs the result. Regardless of the horizontal partition, on a multiple-node network, communication is required for the transducer network to produce the required output. \square

So, coordination-freeness is not guaranteed when we use the relation All , but yet, monotonicity is not lost:

THEOREM 10. *Every query distributedly computed by a transducer that does not use the system relation Id , is monotone.*

PROOF. Let Π be a network-topology independent transducer and let Q be the query distributedly computed by Π . Let $I \subseteq J$ be two input instances and let $t \in Q(I)$. We must show $t \in Q(J)$. Consider the network \mathcal{R}_4 with four nodes 1–2–3–4–1 in the form of a ring. Let H be the horizontal partition of I that places the entire I at every node. Consider now the following, fair, run ρ of (\mathcal{R}_4, Π) on H . This particular run has a fifo behavior of the message buffers. We go around the network in rounds. The construction is such (proof omitted) that after each round, all nodes have the same state and the same fifo message buffer queue. In each round, we first let each node do a heartbeat transition. Then, if some (hence every) input buffer is nonempty, let each node do a delivery transition, receiving the first tuple in its message buffer. If the buffers are empty, we let each node do a second heartbeat transition. Since $t \in Q(I)$, we know that t is output during run ρ . Without loss of generality, assume node 1 outputs tuple t in round m during run ρ .

We now consider the modified network \mathcal{R}' on the same four nodes, obtained by adding the shortcut 2–4 to \mathcal{R}_4 . Consider the horizontal partition H' of J defined by $H'(1) = H'(2) = H'(4) = I$ and $H'(3) = J \setminus I$. Consider now the following prefix ρ' of a possible run of (\mathcal{R}', Π) on H' . The idea is that run ρ is mimicked until round m , but we ignore node 3 completely. The construction is such (proof omitted) that

after each round, nodes 1, 2 and 4 have the same state and the same fifo message buffer queue as after the same round in ρ . In each round i , we first let each node 1, 2 and 4 do a heartbeat transition. Then, if in the same round in ρ we made delivery transitions, then we make the same delivery transitions in ρ' but not for node 3. If in round i we did a series of second heartbeat transitions, we do the same in ρ' but again not for node 3.

The result is that t is also output by node 1 during any fair run that has ρ' as a prefix. Since Π is network-topology independent, we have $t \in Q(J)$ as desired. \square

As a corollary we can add two more statements to the three equivalent statements of the CALM Property (Corollary 8):

COROLLARY 11. *The following are equivalent for any query Q :*

1. Q can be distributedly computed by an oblivious transducer.
2. Q can be distributedly computed by a transducer that does not use the Id relation.
3. Q can be distributedly computed by a transducer that does not use the All relation.

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 [9]. This justifies our modeling of this knowledge in the form of the system relations Id and All . Importantly, we have shown that these relations are only necessary if one wants to compute a nonmonotone query in a distributed fashion.

8. DEDALUS

Dedalus [8] is the declarative language used by Alvaro et al. to model and program network protocols. The precise expressive power of Dedalus needs to be better understood. Here, we compare Dedalus to our setting and we also show that Dedalus can at least simulate arbitrary Turing machines in an eventually consistent manner.¹ By the time hierarchy theorem [18], it follows that eventually-consistent Dedalus programs are not contained in PTIME, let alone in Datalog.

Dedalus is a temporal version of Datalog with negation where the last position of each predicate carries a timestamp. All subgoals of any rule must be joined on this timestamp. The timestamp of the head of the rule can either be the timestamp of the body (a “deductive rule”), or it can be the successor timestamp (an “inductive rule”). The deductive rules must be stratifiable, thus guaranteeing modular stratification and a deterministic semantics through a unique minimal model. Note how this corresponds well to transducers using stratified Datalog as local query language.

Furthermore, Dedalus has a non-deterministic construct by which facts can be derived with a random timestamp, used to model asynchronous communication. In our transducer networks, the same effect is achieved by the semantics we have given, by which one node may send a message in its

¹This is not the same as saying that all computable queries can be expressed in Dedalus; we conjecture this is not the case.

n th local step, whereas another node may receive the message in its m th local step where m can be smaller as well as larger than n . As long as a Dedalus program is monotone in the relations derived by asynchronous rules, the program remains deterministic, but there is no longer a simple syntactic guarantee for this.

The feature that makes Dedalus quite powerful is that timestamp values can also occur as data values, i.e., in other predicate positions than the last one. This feature, called “entanglement”, is intriguing and makes Dedalus go beyond languages such as temporal Datalog [11]. Note that entanglement does not involve arithmetic on timestamps; it merely allows them to be copied in relations in a safe, Datalog-like manner.

Turing machine simulations in database query languages are well known [1, 3, 16], but the Dedalus setting is new, so we describe the Turing machine simulation in some detail. For any database schema \mathcal{S} we can consider the database schema $\mathcal{S}^{\text{time}}$ with the same relation names as \mathcal{S} , but in $\mathcal{S}^{\text{time}}$ each relation name has an arity higher by one than in \mathcal{S} , in order to accommodate timestamps. Dedalus works with temporal database instances; these are instances over schemas of the form $\mathcal{S}^{\text{time}}$ in which the last coordinate of every fact is a natural number acting as timestamp. For any instance I over $\mathcal{S}^{\text{time}}$ and any timestamp value n , let $I|_n$ be the instance over \mathcal{S} obtained from the facts in I that have timestamp n , and let \hat{I} equal $\bigcup_n I|_n$.

Now let Σ be an arbitrary but fixed finite alphabet, and consider the database schema \mathcal{S}_Σ consisting of relation names *Tape* of arity two, *Begin* and *End* of arity one, and a of arity one for each $a \in \Sigma$. Recall that any string $s = a_1 \dots a_p$ over Σ can be presented as an instance I_s over \mathcal{S}_Σ . We consider only strings of length at least two. Then I_s consists of the facts *Tape*(1, 2), ..., *Tape*($p-1$, p), *Begin*(1), *End*(p), $a_1(1)$, ..., $a_p(p)$. Such instances, and isomorphic instances, are known as *word structures* [20].

For any Turing machine M , we define the boolean (0-ary) query Q_M over the class of temporal instances over $\mathcal{S}_\Sigma^{\text{time}}$ as follows.

- If \hat{I} is a word structure representing string s , and M accepts s , then $Q_M(I)$ equals true (encoded by the 0-ary relation containing the empty tuple). If M does not terminate on s , then $Q_M(I)$ is undefined.
- If \hat{I} contains a word structure, but is not a word structure (due to spurious facts), then $Q_M(I)$ also equals true.
- In all other cases $Q_M(I)$ equals false (encoded by the empty 0-ary relation).

The second item in the definition is there to ensure that Q_M is monotone; nevertheless, when we give Q_M a proper word structure as input, a faithful simulation of M is required. Hence, the computational complexity of Q_M is as high as that of the language accepted by M .

We say that a (deterministic) Dedalus program Π expresses a boolean query Q over temporal instances, if for every I such that $Q(I)$ is defined, $\Pi(I)$ contains a fact *Accept*(n) for some n if and only if $Q(I)$ is true. Moreover, Π expresses Q in an eventually consistent way if for every I such that $Q(I)$ is defined, there exists n such that $\Pi(I)|_m = \Pi(I)|_n$ for all $m \geq n$.

THEOREM 12. *For every Turing machine M , the query Q_M is expressible in an eventually consistent way by a Dedalus program.*

PROOF. We only sketch the proof and assume some familiarity with Dedalus. The main difficulties to overcome are the following.

1. Detection of a word structure. Since input facts can arrive at any timestamp, they are persisted, e.g.,

$$a(x, n+1) \leftarrow a(x, n) \quad \text{for each } a \in \Sigma$$

(Officially, this should be done using “pos-predicates” [8].) A word structure is detected at time n if there is a path in the *Tape* relation, beginning in an element in *Begin*, and ending in an element in *End*, where all elements on the path are *labeled*, i.e., belong to some a relation. This is readily expressed in Datalog.

2. Detection of spurious tuples. When a word structure is already detected, we can detect spurious tuples by checking for one of the following conditions, which can be expressed in stratified Datalog:

- (a) *Begin* and *End* contain more than a single element.
- (b) An element in the active domain is labeled by two different alphabet letters.
- (c) *Tape* is more than a successor relation from its begin to its end point, i.e., there is an element on the tape with out-degree or in-degree more than one, or there is an element on the tape that is not reachable from *Begin*.
- (d) There exists a phantom element, i.e., an element in the active domain that is not labeled, or that is not on the tape.

3. Turing machine simulation. When a proper word structure is discovered, without spurious tuples, the simulation of M is started. We copy the a predicates to a^{simul} predicates. This is necessary because a is persisted, which would cause the simulation to be overwritten. And we need to continue persisting a because new (spurious) a facts may still arrive after the simulation has already started. Each transition of M goes to the next timestamp. For each state q of M we use a predicate $q(x, n)$ that holds if M at time n is in state q with its head on position x . Timestamp values (entanglement) is used to extend the finite tape to the right. Care must be taken to do this only when necessary, to ensure eventual consistency. Moreover, we must avoid confusing timestamp values that may already occur as input tape cells, with timestamp values that are used to build the tape extension. Thereto we use a separate predicate *TapeExt* to represent the tape extension. For example, the first time that M extends the input tape, in some state q and seeing letter a at the last input position, is expressed by the following rules:

$$\begin{aligned} \text{ExtNext}(x, n) &\leftarrow \text{TapeExt}(x, y, n) \\ \text{TapeExt}(x, n, n+1) &\leftarrow q(x, n), a(x, n), \\ &\quad \text{End}(x, n), \neg \text{ExtNext}(x, n) \end{aligned}$$

For positions on the extension tape, we use predicates q_{ext} instead of q and $a_{\text{ext}}^{\text{simul}}$ instead of a^{simul} .

□

Distribution is not built in Dedalus but can be simulated using precisely one extra data value in each fact that serves as a location specifier. For a rule all subgoals must be joined on the location specifier and one can differentiate between rules depending on how the location specifier in the head relates to the location specifier in the body [8]. The above theorem can be extended to a distributed setting where different peers send around their input data to their peers. The receiving peer treats these messages as EDB facts. This works without coordination since the program is monotone in the EDB relations. More generally, it seems one can define a syntactic class of “oblivious” Dedalus programs in analogy to our notion of oblivious transducers: for each Dedalus rule disallow the projection of location variables in the body into positions of the head fact that are not used as the location specifier. Roughly speaking, that prevents a Dedalus program to reason about what nodes send what facts.

9. CONCLUSION

Encouraged by Hellerstein [14, 15], 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.
- Investigate the expressiveness of variations or extensions of the basic distributed computation model presented here.
- Understand the exact expressive power of the Dedalus language, as well as the automated verification of Dedalus programs.
- Identify special cases where essential semantic notions such as monotonicity, consistency, network-topology independence, coordination-freeness, are decidable.

10. REFERENCES

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] S. Abiteboul and E. Simon. Fundamental properties of deterministic and nondeterministic extensions of Datalog. *Theoretical Computer Science*, 78:137–158, 1991.
- [3] S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences*, 43(1):62–124, 1991.
- [4] S. Abiteboul and V. Vianu. Generic computation and its complexity. In *Proceedings 23rd ACM Symposium on the Theory of Computing*, pages 209–219, 1991.
- [5] S. Abiteboul and V. Vianu. Computing with first-order logic. *Journal of Computer and System Sciences*, 50(2):309–335, 1995.
- [6] S. Abiteboul, V. Vianu, et al. Relational transducers for electronic commerce. *Journal of Computer and System Sciences*, 61(2):236–269, 2000.
- [7] F.N. Afrati, S.C. Cosmadakis, and M. Yannakakis. On Datalog vs polynomial time. *Journal of Computer and System Sciences*, 51(2):177–196, 1995.
- [8] P. Alvaro, W. Marczak, et al. Dedalus: Datalog in time and space. Technical Report EECS-2009-173, University of California, Berkeley, 2009.
- [9] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. Wiley, 2004.
- [10] A. Blass, Y. Gurevich, and J. Van den Bussche. Abstract state machines and computationally complete query languages. *Information and Computation*, 174(1):20–36, 2002.
- [11] J. Chomicki and T. Imieliński. Finite representation of infinite query answers. *ACM Transactions on Database Systems*, 18(2):181–223, June 1993.
- [12] A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic verification of data-centric business processes. In *Proceedings 12th International Conference on Database Theory*, 2009.
- [13] A. Deutsch, L. Sui, and V. Vianu. Specification and verification of data-driven Web applications. *Journal of Computer and System Sciences*, 73(3):442–474, 2007.
- [14] J.M. Hellerstein. Datalog redux: experience and conjecture. In *Proceedings 29th ACM Symposium on Principles of Database Systems*, pages 1–2. ACM Press, 2010.
- [15] J.M. Hellerstein. The declarative imperative: experiences and conjectures in distributed logic. *SIGMOD Record*, 39(1):5–19, 2010.
- [16] N. Immerman. *Descriptive Complexity*. Springer, 1999.
- [17] B.T. Loo et al. Declarative networking. *Communications of the ACM*, 52(11):87–95, 2009.
- [18] C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [19] M. Spielmann. Verification of relational transducers for electronic commerce. *Journal of Computer and System Sciences*, 66(1):40–65, 2003.
- [20] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, chapter 7. Springer, 1997.