DBSP: Automatic Incremental View Maintenance for Rich Query Languages

Mihai Budiu
VMware Research

Tej Chajed
VMware Research

Frank McSherry
Materialize Inc.

Leonid Ryzhyk
VMware Research

Val Tannen
University of Pennsylvania

December 21, 2022

Abstract

Incremental view maintenance (IVM) has been for a long time a central problem in database theory [26]. Many solutions have been proposed for restricted classes of database languages, such as the relational algebra, or Datalog. These techniques do not naturally generalize to richer languages. In this paper we give a general, heuristic-free solution to this problem in 3 steps: (1) we describe a simple but expressive language called DBSP for describing computations over data streams; (2) we give a new mathematical definition of IVM and a general algorithm for solving IVM for arbitrary DBSP programs, and (3) we show how to model many rich database query languages (including the full relational algebra, queries over sets and multisets, arbitrarily nested relations, aggregation, flatmap (unnest), monotonic and non-monotonic recursion, streaming aggregation, and arbitrary compositions of all of these) using DBSP. As a consequence, we obtain efficient incremental view maintenance algorithms for queries over all these languages.

This document is work in progress. It contains a formal specification of the DBSP language, proofs of the theoretical results, and the specification of several query languages in DBSP. A shorter earlier preprint is available at `https://arxiv.org/abs/2203.16684`.

Contents

1 Introduction **4**

2 Related work **9**
- 2.1 Incremental View Maintenance 9
- 2.2 Stream computation models 11
- 2.3 Connection to synchronous circuits 12

I Streaming and incremental computations **12**

3 Streams **13**
3.1 Streams and stream operators 13
3.1.1 Stream operators by lifting 14
3.1.2 Basic stream operator equivalences 14
3.2 Streams over abelian groups 16
3.2.1 Delays and time-invariance 16
3.3 Causal and strict operators 17
3.4 Streams as an abelian group 21
3.5 Differentiation and Integration 23

4 Relational algebra in DBSP **25**
4.0.1 Generalizing box-and-arrow diagrams 25
4.1 $\mathbb{Z}$-sets as an abelian group 25
4.2 Sets, bags, and $\mathbb{Z}$-sets 26
4.3 Streams over $\mathbb{Z}$-sets 28
4.4 Implementing the relational algebra 28
4.4.1 Query composition 28
4.4.2 Set union 30
4.4.3 Projection 30
4.4.4 Selection 30
4.4.5 Filtering 31
4.4.6 Cartesian products 31
4.4.7 Joins 32
4.4.8 Set intersection 32
4.4.9 Set difference 32

5 Incremental computation **32**

6 Incremental relational queries **36**
6.1 Optimizing *distinct* 37
6.1.1 Anti-joins 38
6.2 Parallelization 38
6.3 Incremental relational queries 38
6.3.1 Example 39
6.4 Complexity of incremental circuits 41

7 Nested streams **42**
7.1 Creating and destroying streams 42
7.1.1 Stream introduction 42
7.1.2 Stream elimination 42
7.1.3 The E and X operators 43
7.1.4 Time domains 43
7.2 Streams of streams 44
7.2.1 Defining nested streams 44

7.2.2 Lifting stream operators 44
7.2.3 Strict operators on nested streams 45
7.2.4 Lifted cycles 46
7.3 Fixed-point computations 47

8 Recursive queries in DBSP 48
8.1 Implementing recursive queries 48
8.1.1 Recursive rules 48
8.2 Example: a recursive query in DBSP 51
8.2.1 Mutually recursive rules 52

9 Incremental recursive queries 54
9.1 Incrementalizing a recursive query 55

10 Complexity of recursive incremental circuits 57

11 Additional query languages 58
11.1 Aggregation 58
11.2 Nested relations 59
11.2.1 Indexed partitions 59
11.3 Grouping; indexed relations 60
11.4 `GROUP BY-AGGREGATE` 60
11.5 Streaming joins 61
11.6 Explicit delay 61
11.7 Multisets/bags 62
11.8 Window aggregates 62
11.9 Relational while queries 62

II Implementation 63

12 Well-formed circuits 63
12.1 Primitive nodes 63
12.2 Circuits as graphs 64
12.3 Circuit semantics 64
12.4 Circuit construction rules 64
12.4.1 Single node 65
12.4.2 Delay node 65
12.4.3 Sequential composition 66
12.4.4 Adding a back-edge 67
12.4.5 Lifting a circuit 68
12.4.6 Bracketing 68

13 Implementing WFC as Dataflow Machines 69

14 Implementing Differential Datalog in DBSP **71**
14.1 Differential Datalog syntax and semantics 72
14.2 Compiling Datalog programs to circuits 74
14.2.1 Rule compilation 75
14.2.2 Relation terms in rule bodies 76
14.2.3 Repeated rule heads (set union) 76
14.2.4 Projection 77
14.2.5 Flatmap in DDlog 78
14.2.6 Map in DDlog 79
14.2.7 Filtering 79
14.2.8 Grouping 80
14.2.9 Aggregation 81
14.2.10 Cartesian products 81
14.2.11 Joins 82
14.2.12 Set intersection 83
14.2.13 Negation 83
14.2.14 Set difference 83
14.2.15 Antijoin 84
14.3 Streaming Differential Datalog 84
14.3.1 Streaming Datalog 84
14.3.2 Streaming Differential Datalog 85

15 Implementations **85**
15.1 DBSP Rust library 85
15.2 SQL compiler 86
15.3 Formal verification 86
15.4 Additional Implementation Observations 86
15.4.1 Checkpoint/restore 86
15.4.2 Maintaining a database 86
15.4.3 Materialized views 86
15.4.4 Maintaining input invariants 87

III Appendixes **88**

A Z-transform and stream convolutions **88**

1 Introduction

Incremental view maintenance (IVM) is an important and well-studied problem in databases. The IVM problem can be stated as follows: given a database DB and a view V defined as a function of the database contents (described by a query Q, i.e. $V = Q(DB)$), maintain the contents of V in response to changes of the database, ideally more efficiently than by simply reevaluating $Q(DB)$ from scratch. The goal is to provide an algorithm that can evaluate Q over the

changes ΔDB applied to the database, since in general the size of the changes is small $|\Delta DB| \ll |DB|$.

This paper provides a new perspective by proposing a new definition of IVM based on a streaming model of computation[1]. Whereas previous IVM solutions are based on defining a notion of a (partial) derivative of Q with respect to its inputs, our definition only requires computing *derivatives of streams* as functions of time. Derivatives of streams are always well-defined (assuming that the data computed on has a notion of difference that satisfies some simple mathematical properties — i.e., it forms a commutative group. Fortunately, it has long been known that relational databases can be modeled in such a way, e.g. [24, 36].)

DBSP has several attractive properties:

1. is is **expressive**. (a) It can be used to define precisely multiple concepts: traditional queries, streaming computations, and incremental computations. (b) We have been able to express in DBSP the full relational algebra, computations over sets and bags, nested relations, aggregation, flatmap, monotonic and nonmonotonic recursion, stratified negation, while-relational programs, window queries, streaming queries, streaming aggregation, and incremental versions of all of the above. In fact, we have built a DBSP implementation of the complete SQL language (Section 15).

2. it is **simple**. DBSP is built entirely on elementary concepts such as functions and algebraic groups.

3. mathematically **precise**. All the results in this paper have been formalized and checked using the Lean proof assistant [17].

4. it is **modular**, in the following sense: (a) the incremental version of a complex query can be reduced recursively to incrementalizing its component subqueries. This gives a simple, syntactic, heuristic-free algorithm (Algorithm 6.4) that converts an arbitrary DBSP query to its incremental form. (b) Extending DBSP to support new primitive operators is easy, and they immediately benefit from the rest of the theory of incrementalization. An important consequence of modularity is that the theory can be efficiently implemented, as we briefly discuss in Section 15.

The core concept of DBSP is the *stream*, which is used to model changes over time. We use $\mathcal{S}_A$ to denote the type of infinite streams with values of type A. If $s \in \mathcal{S}_A$ is a stream, then $s[t] \in A, t \in \mathbb{N}$ is the tth element of s, also referred to as the *value of the stream at time* t. A streaming computation is a function that consumes one or more streams and produces another stream. We show streaming computations with diagrams, also called "circuits", where boxes are computations and streams are arrows. The following diagram shows a stream operator $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$, consuming two input streams s_0 and s_1 and producing one output stream s:

[1] Our model is inspired by Digital Signal Processing [50], applied to databases, hence the name DBSP.

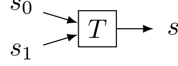

We generally think of streams as sequences of "small" values, such as insertions or deletions in a database. However, we also treat the whole database as a *stream of database snapshots.* We model a database as a stream $DB \in \mathcal{S}_{SCH}$, where SCH is the database schema. Time is not wall-clock time, but counts the transactions applied to the database. (Since transactions are linearizable, they have a total order.) $DB[t]$ is the snapshot of the database contents after t transactions have been applied.

Database transactions also form a stream ΔDB, this time a stream of *changes*, or *deltas* that are applied to the database. The values of this stream are defined by $(\Delta DB)[t] = DB[t] - DB[t-1]$, where "$-$" stands for the difference between two databases, a notion that we will soon make more precise. The ΔDB stream is produced from the DB stream by the *stream differentiation* operator $D : \mathcal{S}_A \to \mathcal{S}_A$; this operator produces as its output the stream of changes from its input stream; we have thus $D(DB) = \Delta DB$.

Conversely, the database snapshot at time t is the cumulative result of applying all transactions up to t: $DB[t] = \sum_{i \leq t} \Delta DB[i]$. The operation of adding up all changes is the inverse of differentiation, and is another basic stream operator, *stream integration*: $I : \mathcal{S}_A \to \mathcal{S}_A$. The following diagram expresses the relationship between the streams ΔDB and DB:

$$\Delta DB \to \boxed{I} \to DB \to \boxed{D} \to \Delta DB$$

Suppose we have a query $Q : SCH \to SCH$ defining a view V. What is a view in a streaming model? It is also a stream! For each snapshot of the database stream we have a snapshot of the view: $V[t] = Q(DB[t])$. In general, given an arbitrary function $f : A \to B$, we define a streaming "version" of f, denoted by $\uparrow\!f$ (read as "f lifted"), which applies f to every element of the input stream independently. We can thus write $V = (\uparrow\!Q)(DB)$.

Armed with these basic definitions, we can now precisely define IVM. What does it mean to maintain a view incrementally? We claim that an efficient maintenance algorithm needs to compute the *changes* to the view given the changes to the database. We thus define the IVM of a query Q by chaining the above three definitions: $\Delta V \stackrel{\text{def}}{=} D(V) = D(\uparrow\!Q(DB)) = D(\uparrow\!Q(I(\Delta DB)))$. This can be shown as the following diagram, which is the central definition of this paper:

$$\Delta DB \to \boxed{I} \xrightarrow{DB} \boxed{\uparrow\!Q} \xrightarrow{V} \boxed{D} \to \Delta V$$

Given a query Q we define its incremental version as $Q^{\Delta} \stackrel{\text{def}}{=} D \circ \uparrow\!Q \circ I$. The incremental version of a query is a *streaming operator* which computes directly on changes and produces changes. The incremental version of a query is thus always well-defined. The above definition shows one way to compute a query

incrementally, but applying it naively will generally produce an inefficient execution plan, since it will reconstruct the database at each step. In Section 5 we show how algebraic properties of the $\cdot^\Delta$ transformation can be used to optimize the implementation of Q^Δ. The first key property is that the composition of queries can be incrementalized by composing the incremental versions of its constituents, that is $(Q_1 \circ Q_2)^\Delta = {Q_1}^\Delta \circ {Q_2}^\Delta$. The second key property is that essentially all primitive database operations have efficient incremental versions.

Armed with this general theory of incremental computation, in Section 4 we show how to model relational queries in DBSP. This immediately gives us a general algorithm to compute the incremental version of any relational query. These results were previously known, but they are cleanly modeled by DBSP. Section 14 shows how recursive Datalog programs with stratified negation can be implemented in DBSP, and Section 7.2 gives *incremental streaming computations for recursive programs*. For example, given an implementation of transitive closure in the natural recursive way, our algorithm produces a program that efficiently maintains the transitive closure of a graph as the graph is changed by adding and deleting edges.

We have formalized the entire DBSP theory in the Lean proof assistant ; our formalization includes machine-checked proofs of correctness for all the theorems stated in this paper.

MIHAI: Need a URL for this

This paper makes the following contributions:

1. DBSP, a **simple** but **expressive** language for streaming computation. DBSP gives an elegant formal foundation unifying the manipulation of streaming and incremental computations.

2. An algorithm for incrementalizing any streaming computation expressed in DBSP.

3. An illustration of how DBSP can be applied to various query classes, such as relational algebra, nested relations, aggregations, flatmap, and stratified-monotonic Datalog.

4. We offer the first fully mechanically-verified theory of IVM.

5. We provide a high-performance open-source implementation of DBSP as a general-purpose streaming query engine in Rust.

The following tables summarize the mathematical notations used in the rest of this paper.

General notations	
$\mathbb{Z}$	The ring of integer numbers
$\mathbb{N}$	The set of natural numbers $0, 1, 2, \ldots$
$\mathbb{B}$	The set of Boolean values
$[n]$	The natural numbers between 0 and $n-1$
id	The identity function over some domain $id : A \to A$, $id(x) = x$
$[\![Q]\!]$	Semantics of query (function) Q
$\langle a, b \rangle$	The pair containing values a and b
$\mathrm{fst}(p)$	The operator that returns the first value of a pair p
$\mathrm{snd}(p)$	The operator that returns the second value of a pair p
$a \mapsto b$	The function that maps a to b and everything else to 0
$\lambda x.M$	An anonymous function with argument x and body M
$\mathrm{fix}\, x.f$	The (unique) solution (fixed point) of the equation $f(x) = x$

Streams	
$\mathcal{S}_A$	The set of streams with elements from a group A; $\mathcal{S}_A = \{f \mid f : \mathbb{N} \to A\}$
$\overline{\mathcal{S}_A}$	Streams with elements from a group A that are 0 almost everywhere
$s[t]$	The t-th element of a stream; $s[t] = s(t)$
$\uparrow f$	An operator applied to a function $f : A \to B$ to produce a function $\uparrow f : \mathcal{S}_A \to \mathcal{S}_B$ operating pointwise
$\text{zpp}(f)$	$\text{zpp}(f)$ iff $f(0) = 0$ for $f : A \to B$ for A, B groups
z^{-1}	The stream delay operator $z^{-1} : \mathcal{S}_A \to \mathcal{S}_A$, that outputs a 0 followed by the input stream
I	The stream integration operator $I : \mathcal{S}_A \to \mathcal{S}_A$
D	The stream differentiation operator $D : \mathcal{S}_A \to \mathcal{S}_A$
Q^Δ	The incremental version of an operator $Q^\Delta = D \circ Q \circ I$
$s\vert_{\leq t}$	A stream that has the same prefix as s up to t, then it is all 0s
$s\vert_{<t}$	A stream that has the same prefix as s up to $t-1$, then it is all 0s
$\cong$	Symbol that indicates that two circuits compute the same function
δ_0	A function that produces a stream from a scalar: scalar, followed by zeros
$\int$	A function that produces a scalar by adding all elements of a stream
E	$E = I \circ \delta_o$
X	$X = \int \circ D$
$\mathbb{Z}$-sets	
$\mathbb{Z}[A]$	$\mathbb{Z}$-sets: finite functions from $A \to \mathbb{Z}$
DB	A database
ΔDB	A change to a database
$\vert s\vert$	Size of $\mathbb{Z}$-set s
isset	A function $\text{isset} : \mathbb{Z}[A] \to \mathbb{B}$ that determines whether its argument is a set
distinct	A function $\mathit{distinct} : \mathbb{Z}[A] \to \mathbb{Z}[A]$ that always returns a set
ispositive	A function $\text{ispositive} : \mathbb{Z}[A] \to \mathbb{B}$ that determines whether all elements of a $\mathbb{Z}$-set have positive weights
toszet	Function converting a set to a $\mathbb{Z}$-set
toset	Function converting a $\mathbb{Z}$-set into a set

2 Related work

2.1 Incremental View Maintenance

Incremental view maintenance [27, 25, 14, 26, 15] is a much studied problem in databases. A survey of results for Datalog queries is present in [46]. The standard approach is as follows: given a query Q, discover a "delta query", a"differential" version ΔQ that satisfies the equation: $Q(d + \Delta d) = Q(d) +$

$\Delta Q(d, \Delta d)$, and which can be used to compute the change for a new input reusing the previous output. DBToaster introduced recursive recursive IVM [6, 36], where the incrementalization process is repeated for the delta query.

Many custom algorithms were published for various classes of queries: e.g. [37] handles positive nested relational calculus. DYN [29] and IDYN [30, 31] focus on acyclic conjunctive queries. Instead of keeping the output view materialized they build data structures that allow efficiently querying the output views. PAI maps [4] are specially designed for queries with correlated aggregations. AJU [53] focuses on foreign-key joins. It is a matter of future work to evaluate whether custom DBSP operators can match the efficiency of systems specialized for narrow classes of queries.

DBSP is a bottom-up system, which always produces eagerly the*changes* to the output views. Instead of maintaining the output view entirely, DBSP proposes generating deltas as the output of the computation (similar to the kSQL [32] `EMIT CHANGES` queries). The idea that both inputs and outputs to an IVM system are streams of changes seems trivial, but this is key to the symmetry of our solution: both in our definition of IVM (5.1), and the fundamental reason that the chain rule exists — the chain rule is the one that makes our structural induction IVM algorithm possible.

IVM algorithms for Datalog-like languages frequently use counting based approaches [18, 47] that maintain the number of derivations of each output fact: DRed [27] and its variants [13, 54, 52, 38, 42, 8], the backward-forward algorithm and variants [47, 28, 46]. DBSP is more general than these approaches, and our incrementalization algorithm handles arbitrary recursive queries and generates more efficient plans for recursive queries in the presence of arbitrary updates (especially deletions, where competing approaches may over-delete). Interestingly, the $\mathbb{Z}$-sets multiplicities in DBSP are related to the counting-number-of-derivations approaches, but our use of the *distinct* operator shows that precise counting is not necessary.

Picallo et al. [7] provide a general solution to IVM for rich languages. DBSP requires a group structure on the values operated on; this assumption has two major practical benefits: it simplifies the mathematics considerably (e.g., Picallo uses monoid actions to model changes), and it provides a general, simple algorithm (6.4) for incrementalizing arbitrary programs. The downside of DBSP is that one has to find a suitable group structure (e.g., $\mathbb{Z}$-sets for sets) to "embed" the computation. Picallo's notion of "derivative" is not unique: they need creativity to choose the right derivative definition, we need creativity to find the right group structure.

Finding a suitable group structure has proven easy for relations (both [36] and [23] use $\mathbb{Z}$-sets to uniformly model data and insertions/deletions), but it is not obvious how to do it for other data types, such as sorted collections, or tree-shaped collections (e.g., XML or JSON documents) [19]. An intriguing question is "what other interesting group structures could this be applied to besides $\mathbb{Z}$-sets?" Papers such as [49] explore other possibilities, such as matrix algebra, linear ML models, or conjunctive queries.

[12] implemented a verified IVM algorithm for a particular class of graph

queries called Regular Datalog, with an implementation machine-checked in the Coq proof assistant. Their focus is on a particular algorithm and the approach does not consider other SQL operators, general recursion, or custom operators (although it is modular in the sense that it works on any query by incrementalizing it recursively). Furthermore, for all queries a deletion in the input change stream requires running the non-incremental query to recover.

DBSP does not do anything special for triangle queries [35]. Are there better algorithms for this case?

In Section 11 we have briefly mentioned that DBSP can easily model window and stream database queries [9, 2]; it is an interesting question whether there are CQL queries that cannot be expressed in DBSP (we conjecture that there aren't any).

DBSP is also related to Differential Dataflow (DD) [45, 48] and its theoretical foundations [3] (and recently [44, 16]). DD's computational model is more powerful than DBSP, since it allows past values in a stream to be "updated". In fact, DD is the only other framework which we are aware of which can incrementalize recursive queries as efficiently as DBSP does. In contrast, our model assumes that the inputs of a computation arrive in the time order while allowing for nested time domains via the modular lifting transformer ($\uparrow$). DBSP can express both incremental and non-incremental computations; in essence DBSP is "deconstructing" DD into simple component building blocks. Most importantly, DBSP comes with Algorithm 6.4, a syntax-directed translation that can convert any expressible query into an incremental version — in DD users have to assemble incremental queries manually using incremental operators. (materialize.com offers a product that automates incrementalization, but only for SQL queries. Differential Datalog [51] does it for a Datalog dialect.) Unlike DD, DBSP is a modular theory, which easily accommodates the addition of new operators. In particular, we have given full mechanical proofs of DBSP's correctness.

2.2 Stream computation models

DBSP using non-nested streams is a simplified instance of a Kahn network [34]. Johnson [33] studies a very similar computational model without nested streams and its expressiveness. The implementation of such streaming models of computation and their relationship to dataflow machines has been studied by Lee [40]. Lee [39] also introduced streams of streams and the $\uparrow z^{-1}$ operator.

[20] surveys the connection between synchronous digital circuits and functional programs. Our circuits are nothing but higher order functions computing on streams (functions themselves). The paper's main focus are circuits processing numeric data, whereas, taking advantage of our circuits' ability to compute on arbitrary groups, we use circuits to implement incremental view maintenance for relational databases.

Mamouras [43] gives a formal theory of stream computation.

2.3 Connection to synchronous circuits

There is a vast literature on **synchronous circuits**, which are well-defined models for hardware circuits e.g. [20]. These circuits also compute over infinite streams of values, usually of Booleans $\mathcal{S}_{\mathbb{B}}$. In a **combinational circuit** the output values depend only on the current input values. These are pure lifted streaming computations. A **sequential circuit** can have outputs that depend on past input values. These are always causal circuits. Sequential synchronous circuits use latches or flip-flops to store state; the latches are controlled by a global clock signal. These correspond to the z^{-1} operator. In a well-formed sequential circuit all back-edges must go through some latch — this corresponds to our circuit construction rule that requires a delay element on each back-edge.

Languages such as Verilog or VHDL can be used to specify such circuits. (However, both Verilog and VHDL are strictly more powerful, and can express richer classes of circuits than just synchronous sequential circuits.)

There is a rich literature on synchronous circuits, and some of these results are directly applicable to the circuits we discuss. Here are a few examples.

Retiming [41] is an optimization that "moves" around delay elements while preserving the circuit semantics. Retiming is used traditionally to reduce the clock cycle by minimizing the signal propagation delay between any pair of latches. In our case it could be used for minimizing the amount of internal circuit state.

In a synchronous circuit the *state* is entirely stored in the latches. Saving and restoring the contents of the latches enables such circuits to take a snapshot of their state and resume computation.[2]

Fault tolerance of synchronous circuits is provided by replicating the state elements, to prevent accidental state changes caused by e.g., cosmic rays. We can borrow this idea for building redundant distributed computations.

Pipelining digital circuits is an effective technique for increasing throughput through parallelization, by inserting additional latches and allowing different pipeline stages to compute concurrently on distinct stream values, at the expense of increased latency between the inputs and the corresponding outputs.

Digital circuit latches depend on a special "reset" signal to initialize their state to a pre-established value; this corresponds to the special 0 value in our value domain.

Our nested streams are related to the notion of delta-cycles in the definition of VHDL [10].

[2] In Boolean synchronous circuits this is achieved by connecting all latches into a *scan chain* which can be read and written sequentially after stopping the circuit clock.

Part I
Streaming and incremental computations

3 Streams

3.1 Streams and stream operators

$\mathbb{N}$ is the set of natural numbers, $\mathbb{B}$ is the set of Booleans, and $\mathbb{Z}$ is the set of integers. $[n] = \{0, 1, 2, \ldots n-1\}$ is the set of natural numbers less than n. $\mathbb{R}$ is the set of reals.

Definition 3.1 (stream)**.** Given a set A, a **stream** *of values from* A, or an A-*stream*, is a function $\mathbb{N} \to A$. We denote by $\mathcal{S}_A \stackrel{\text{def}}{=} \{s \,|\, s : \mathbb{N} \to A\}$ the set of all A-streams.

When $s \in \mathcal{S}_A$ and $t \in \mathbb{N}$ we write $s[t]$ for the t-th element of the stream s instead of the usual $s(t)$ to distinguish it from other function applications.

We usually think of the index $t \in \mathbb{N}$ as (discrete) time and of $s[t] \in A$ as the value of the the stream s "at time" t.

For example, the stream of natural numbers given by $id[t] = t$ is the sequence of values $[\ 0 \quad 1 \quad 2 \quad 3 \quad 4 \quad \cdots\ ]$.

Definition 3.2 (stream operator)**.** A (typed) **stream operator** with n inputs is a function $T : \mathcal{S}_{A_0} \times \cdots \times \mathcal{S}_{A_{n-1}} \to \mathcal{S}_B$.

In general we will use "operator" for functions on streams, and "function" for computations on "scalar" values.

DBSP is an extension of the simply-typed lambda calculus – we introduce its elements gradually. However, we find it more readable to also use signal-processing-like circuit diagrams to depict DBSP programs.

Stream operator *composition* (function composition) is shown as chained circuits. The composition of a binary operator $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_A$ with the unary operator $S : \mathcal{S}_A \to \mathcal{S}_B$ into the computation $\lambda s.T(T(s, S(s)), S(s)) : \mathcal{S}_A \to \mathcal{S}_A$ is given by the following circuit:

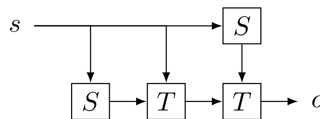

Arrows with a single start and multiple ends denote a stream that is reused multiple times, e.g., s in the above diagram is used 3 times. Diagrams, however, do obscure the ordering of the inputs of an operator; in the above examples we have to indicate which ones are the first and respectively second inputs of T if T is not commutative. Most of our binary operators are commutative.

3.1.1 Stream operators by lifting

One way of building stream operators is by (pointwise) **lifting** functions operating on the stream values. For example, given a (scalar) $f : A \to B$ we can define the stream operator $\uparrow f : \mathcal{S}_A \to \mathcal{S}_B$ by $(\uparrow f)(s) = f \circ s$, or, pointwise, $(\uparrow f)(s)[t] \stackrel{\text{def}}{=} f(s[t])$.

This extends straightforwardly to functions of multiple arguments, e.g., given $T : A \times B \to C$, we can define $\uparrow T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ as $((\uparrow T)(s_0, s_1))[t] \stackrel{\text{def}}{=} T(s_0[t], s_1[t])$.

We call such stream operators **lifted**.

VAL: I hate to be picky about this but we might want to use a different notation for set-theoretical pairing of elements, e.g., in functions of multiple arguments, and category-theoretical pairing of functions. I don't think the latter is captured properly below by $\uparrow\langle ., . \rangle$.

For example, applying the lifted operator $\lambda x.(2x)$ to the stream $id : \mathbb{N} \to \mathbb{N}$ gives as result a stream containing all even values:
$(\uparrow(\lambda x.(2x)))(id) = [\ 0 \quad 2 \quad 4 \quad 6 \quad 8 \quad \cdots\ ]$.

Proposition 3.3 (distributivity)**.** Lifting distributes over function composition: $\uparrow(f \circ g) = (\uparrow f) \circ (\uparrow g)$.

Proof. This is easily proved by using associativity of function composition: $\forall s.(\uparrow(f \circ g))(s) = (f \circ g) \circ s = f \circ (g \circ s) = f \circ (\uparrow g)(s) = (\uparrow f)((\uparrow g)(s)) = (\uparrow f \circ \uparrow g)(s)$. $\square$

We say that two circuits are **equivalent** if they compute the same input-output function on streams. We use the symbol $\cong$ to indicate that two circuits are equivalent. For example, Proposition 3.3 states the following circuit equivalence:

$$s \to \boxed{\uparrow g} \to \boxed{\uparrow f} \to o \quad \cong \quad s \to \boxed{\uparrow(f \circ g)} \to o$$

Two (or more) streams can be combined (**paired**) into a single stream of pairs (tuples) by lifting the scalar pairing operator $\langle \cdot, \cdot \rangle : A \to B \to (A \times B)$, obtaining the stream pair operator: $\uparrow\langle \cdot, \cdot \rangle : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_{A \times B}$, defined as mapping $a \in \mathcal{S}_A$ and $b \in \mathcal{S}_B$ to $\langle a, b \rangle \in \mathcal{S}_{A \times B}$ by pairing elements pointwise $\uparrow\langle a, b \rangle[t] = \langle a[t], b[t] \rangle \in A \times B$.

For example, the stream $\langle id, id \rangle$ is the sequence of pairs
$[\ \langle 0, 0 \rangle \quad \langle 1, 1 \rangle \quad \langle 2, 2 \rangle \quad \langle 3, 3 \rangle \quad \langle 4, 4 \rangle \quad \cdots\ ]$.

Let us also denote by $\text{fst} : A \times B \to A$ the projection that obtains the first element of a pair $\text{fst}(\langle a, b \rangle) = a$, and by $\text{snd} : A \times B \to B$ the projection that obtains the second element of a pair. We obtain useful stream operators by lifting $\uparrow$fst and $\uparrow$snd.

3.1.2 Basic stream operator equivalences

From type theory (or category theory) we recall the standard equalities that pairing and projections satisfy: $\text{fst}(\langle s_0, s_1 \rangle) = s_0$, $\text{snd}(\langle s_0, s_1 \rangle) = s_1$, and $\langle \text{fst}(p), \text{snd}(p) \rangle = p$. By lifting the functions on both left and right we obtain some similar **equivalences of circuits**. In some of the circuits below some inputs are not used.

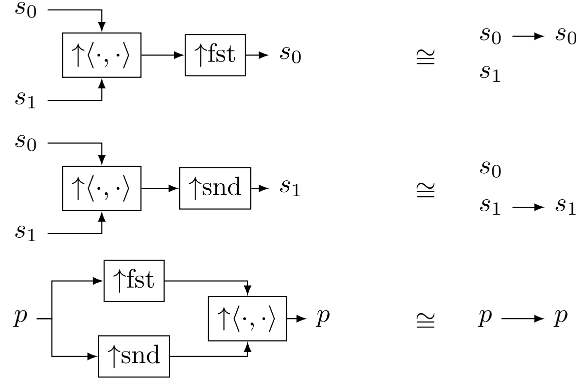

Pairing and projections allow for switching between pairs of streams and streams of pairs, whichever is more convenient. For example, instead of a binary operator $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ we can work with a unary operator $T^u : \mathcal{S}_{A\times B} \to \mathcal{S}_C$ where $T^u(p) = T(\uparrow\text{fst}(p), \uparrow\text{snd}(p))$ and instead of a unary operator $Q : \mathcal{S}_{A\times B} \to \mathcal{S}_C$ we can work with a binary operator $Q^b : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ where $Q^b(s_0, s_1) = Q(\uparrow\langle s_0, s_1\rangle)$.

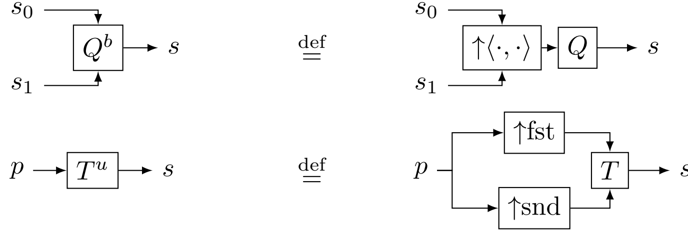

Given two operators $Q : \mathcal{S}_A \to \mathcal{S}_B$ and $R : \mathcal{S}_A \to \mathcal{S}_C$ we define $\uparrow\langle Q, R\rangle : \mathcal{S}_A \to \mathcal{S}_{B\times C}$ by $\uparrow\langle Q, R\rangle(s) = \uparrow\langle Q(s), R(s)\rangle$. In terms of circuit diagrams:

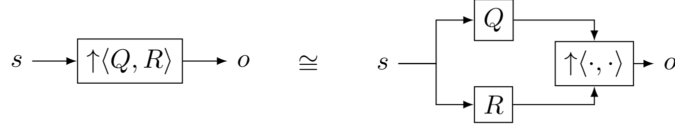

We have standard equalities (from category theory) for this construct such as $\uparrow\text{fst} \circ \uparrow\langle Q, R\rangle = Q$, similarly for snd, and $\uparrow\langle \uparrow\text{fst} \circ W, \uparrow\text{snd} \circ W\rangle = W$. These correspond to equivalences of circuits that follow from the simpler ones above. For example, after substituting the definition of $\langle Q, R\rangle$ we have

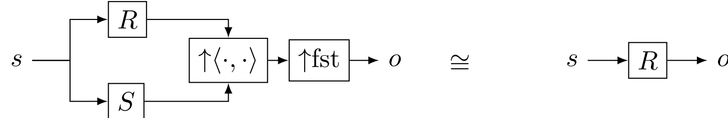

Another useful operator expression notation takes $Q : \mathcal{S}_A \to \mathcal{S}_B$ and $R : \mathcal{S}_D \to \mathcal{S}_C$ and combines them into $Q \times R : \mathcal{S}_{A\times D} \to \mathcal{S}_{B\times C}$ where $(Q \times R)(p) = \langle Q(\uparrow\text{fst}(p)), R(\uparrow\text{snd}(p))\rangle$. This corresponds to the following circuit:

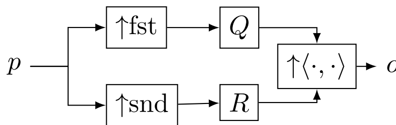

We also have the following circuit equivalence:

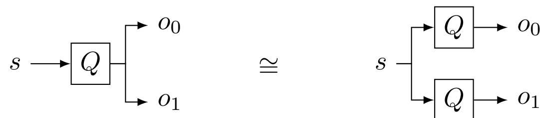

Lifting functions on values and composing stream operators results in a very simple, yet limited, programming language on streams. We next introduce operators that "shift" streams in time. These will be instrumental for enriching the language.

VAL: Should this one be here? It is unrelated to products. Let em email you a proposal for organizing these equivalences and definitions

3.2 Streams over abelian groups

For the rest of the technical development we will require the set of values $(A, +, 0, -)$ for any stream $\mathcal{S}_A$ to form a commutative group.

We denote by $0_{\mathcal{S}_A}$ (or simply 0 when the type is clear) the stream that consist of the special value 0_A at each time moment: $0_{\mathcal{S}_A} \in \mathcal{S}_A, \forall t \in \mathbb{N}.0_{\mathcal{S}_A}[t] \stackrel{\text{def}}{=} 0_A$.

3.2.1 Delays and time-invariance

Definition 3.4 (Delay)**.** The **delay operator**[3] emits an output stream that is the input stream delayed by one element: $z_A^{-1} : \mathcal{S}_A \to \mathcal{S}_A$ defined by:

$$z_A^{-1}(s)[t] \stackrel{\text{def}}{=} \begin{cases} s[t-1] & \text{when } t \geq 1 \\ 0_A & \text{when } t = 0 \end{cases}$$

We often omit the type parameter A, and write just z^{-1}. We denote by z^{-k} the composition of z^{-1} with itself k times (delay by k time units).

$$i \to \boxed{z^{-1}} \to o$$

For example, the delay of the id stream is $z^{-1}(id)$, containing the sequence of values $[\ 0 \quad 0 \quad 1 \quad 2 \quad 3 \quad \cdots\ ]$.

The following definition applies to stream operators of any number of arguments but to keep the notation simpler we formulate it only for binary operators.

Definition 3.5 (Time invariance)**.** A stream operator $T : \mathcal{S}_{A_0} \times \mathcal{S}_{A_1} \to \mathcal{S}_B$ is **time-invariant** if it commutes with the delay operator z^{-1}, that is, $T(z^{-1}(s_0), z^{-1}(s_1)) = z^{-1}(T(s_0, s_1))$ for any $s_0 \in \mathcal{S}_{A_0}, s_1 \in \mathcal{S}_{A_1}$.

[3]The name z^{-1} comes form the DSP literature, and it is related to the z-transform in Section A.

In other words, T is time-invariant if and only if the following two circuits are equivalent:

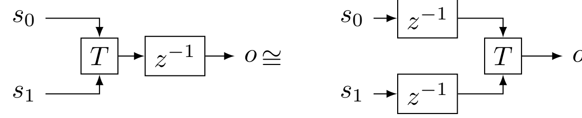

It is straightforward to check that the composition of any number of time-invariant operators of any number of arguments is time invariant. Similarly, the delay operators z^{-k} as well as the pairing and projection operators are time-invariant. In this framework we only deal with time-invariant operators.

Definition 3.6. We say that a function between groups $f : A \to B$ has the **zero-preservation property** iff $f(0_A) = 0_B$. We write $\text{zpp}(f)$. This property generalizes to functions with multiple inputs: e.g., $g : A \times B \to C$ where A, B, C are groups. $\text{zpp}(g)$ iff $g(0_A, 0_B) = 0_C$.

Proposition 3.7. A lifted operator $\uparrow f$ is time-invariant iff $\text{zpp}(f)$.

Notice that it is easy to construct operators that are not time-invariant. Consider the "constant 1" function: $c_1 : \mathbb{N} \to \mathbb{N}$, defined by $c_1(x) = 1, \forall x \in \mathbb{N}$. The stream operator defined by $\uparrow c_1 : \mathcal{S}_N \to \mathcal{S}_N$ produces a stream containing only the constant value 1: $c_1(s)[t] = 1, \forall s \in \mathcal{S}_N, t \in \mathbb{N}$. The stream operator $\uparrow c_1$ is *not* time-invariant, since it does not have the zero preservation property.

3.3 Causal and strict operators

For notation simplicity we again give the next definition only for unary operators; it extends naturally to binary operators through the use of pairing as shown above.

Definition 3.8 (Causality)**.** A stream operator, $S : \mathcal{S}_A \to \mathcal{S}_B$, is **causal** when for any $s, s' \in \mathcal{S}_A$, and all times t we have

$$(\forall i \leq t \; s[i] = s'[i]) \text{ implies } S(s)[t] = S(s')[t]$$

Note that all operators produced by lifting scalar functions are causal. z^{-1} is causal. All DBSP operators are causal.

Definition 3.9 (Cutting)**. Cutting** the stream $s \in \mathcal{S}_A$ at time $t \in \mathbb{N}$ produces a stream

$$(s|_{\leq t})[i] \stackrel{\text{def}}{=} \begin{cases} s[i] & \text{if } i \leq t \\ 0_A & \text{if } i > t \end{cases}$$

For example, cutting the stream *id* at time 2 gives the stream $id|_{\leq 2}$ composed of the sequence $[\; 0 \quad 1 \quad 2 \quad 0 \quad 0 \quad \cdots \;]$.

Note that $s|_{\leq t_1}|_{\leq t_2} = s|_{\leq \min(t_1, t_2)}$. It follows that $s|_{\leq t_1}|_{\leq t_2} = s|_{\leq t_2}|_{\leq t_1}$ (cutting is commutative) and $s|_{\leq t}|_{\leq t} = s|_{\leq t}$, (cutting at time t is idempotent).

Cutting, however, is **not** time-invariant. Cutting is not used as a DBSP operator, it is just a mathematical tool that we will use to reason about the behavior of the circuits we build.

Lemma 3.10. The following are equivalent for a binary stream operator T

(i) T is causal

(ii) $\forall s_1, s_2$ and t we have $T(s_1, s_2)|_{\leq t} \; = \; T(s_1|_{\leq t}, s_2|_{\leq t})|_{\leq t}$.

Using part (ii) it follows immediately that the composition of any number of causal operators of any number of arguments is causal. Moreover, using also the commutativity and idempotence of cutting, it follows that for any t the operator $\lambda s.s|_{\leq t}$ is itself causal.

Definition 3.11 (zero almost everywhere)**.** We say that a stream s is **zero almost-everywhere** if there exists a time $t_0 \in \mathbb{N}$ s.t. $s|_{\leq t_0} = s$.

We denote the set of streams over A that are zero almost everywhere by $\overline{\mathcal{S}_A}$.

Definition 3.12 (Strictness)**.** A stream operator, $F : \mathcal{S}_A \to \mathcal{S}_B$ is **strictly causal** (abbreviated **strict**) if for any $s, s' \in \mathcal{S}_A$ and all times t we have

$$(\forall i < t. s[i] = s'[i]) \;\; \text{implies} \;\; F(s)[t] = F(s')[t]$$

In particular, $F(s)[0] = 0_B$ is the same for all inputs $s \in \mathcal{S}_A$. Strict operators are of course causal. Note that lifted stream operators, while causal, in general are *not* strict.

It can be immediately checked that the operator z^{-1} (in fact, z^{-k} for any positive integer k) is strict. In this text z^{-1} is the only primitive strict operator used.

Definition 3.13 (Strict cutting)**.** **Strictly cutting** the stream $s \in \mathcal{S}_A$ at time $t \in \mathbb{N}$ produces the stream

$$(s|_{<t})[i] \stackrel{\text{def}}{=} \begin{cases} s[i] & \text{if } i < t \\ 0_A & \text{if } i \geq t \end{cases}$$

$s|_{<0}$ is the stream $0_{\mathcal{S}_A}$ that is 0_A at all times. Note also that $s|_{<t+1} = s|_{\leq t}$.

Analogously to Lemma 3.10 an operator $F : \mathcal{S}_A \to \mathcal{S}_B$ is strict iff for any s and t we have

$$F(s)|_{\leq t} \; = \; F(s|_{<t})|_{\leq t}$$

In particular, $F(s)[0] = F(0_{\mathcal{S}_A})[0]$ and $F(s)[t+1] = F(s|_{\leq t})[t+1]$. Note the different zeros in $F(0_{\mathcal{S}_A})[0]$: it features both the stream $0_{\mathcal{S}_A} \in \mathcal{S}_A$, consisting of the group element 0_A at each time moment, and the time moment 0.

The next proposition shows the importance of strict operators.

Proposition 3.14. For any strict operator $F : \mathcal{S}_A \to \mathcal{S}_A$ the equation $\alpha = F(\alpha)$ has a unique solution $\alpha \in \mathcal{S}_A$. In other words, every strict operator has a unique fixed point, which we denote by $\mathrm{fix}\, \alpha. F(\alpha)$.

Proof. Define the solution (the fixed point) by recurrence:

$$\begin{aligned} \alpha[0] &= F(\alpha)[0] = F(0_{\mathcal{S}_A})[0] \\ \alpha[t+1] &= F(\alpha)[t+1] = F(\alpha|_{\leq t})[t+1] \end{aligned}$$

The second equality defines $\alpha[t+1]$ in terms of $\alpha[0], \ldots, \alpha[t]$. Uniqueness follows by strong induction.

□

We will apply the previous proposition to operators obtained by composing strict and causal ones.

Lemma 3.15. Let $k \geq 2$. If F is strict and the k-ary T operator is causal, then for any fixed $s_0, \ldots, s_{k-2}$ the operator $\lambda\alpha.T(s_0, \ldots, s_{k-2}, F(\alpha))$ is strict.

Proof. We show the case $k = 2$. This operator is described by the following diagram with a "feedback loop":

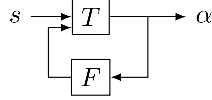

$$\begin{aligned} T(s, F(\alpha))|_{\leq t} &= T(s|_{\leq t}, F(\alpha)|_{\leq t})|_{\leq t} \\ &= T(s|_{\leq t}, F(\alpha|_{<t})|_{\leq t})|_{\leq t} \\ &= T(s, F(\alpha|_{<t}))|_{\leq t} \end{aligned}$$

□

Corollary 3.16. For strict F we have

$$(\operatorname{fix} \alpha.F(\alpha))|_{\leq t} \;=\; \operatorname{fix} \alpha.(F(\alpha)|_{\leq t})$$

Proof. Since cutting itself is a causal operator, it follows from Lemma 3.15 that $\lambda\alpha.(F(\alpha)|_{\leq t})$ is strict so Proposition 3.14 applies and $\operatorname{fix} \alpha.(F(\alpha)|_{\leq t})$ is well-defined.

If we let α be the solution of $\alpha = F(\alpha)$ then, by uniqueness, it suffices to show that $\beta = \alpha|_{\leq t}$ satisfies the equation $\beta = F(\beta)|_{\leq t}$. Indeed, since F is in particular causal

$$\alpha|_{\leq t} = F(\alpha)|_{\leq t} = F(\alpha|_{\leq t})|_{\leq t}$$

□

Corollary 3.17. Let $k \geq 1$. If $F : \mathcal{S}_A \to \mathcal{S}_A$ is strict and $(k+1)$-ary T is causal then the k-ary operator $Q(s_0, \ldots, s_{k-1}) = \operatorname{fix} \alpha.T(s_0, \ldots, s_{k-1}, F(\alpha))$ is well-defined and causal. If, moreover, F and T are time-invariant then so is Q.

Proof. We show the case $k = 1$. The well-definedness of Q follows by applying, for each s, Proposition 3.14 to the operator $\lambda\alpha.T(s, F(\alpha))$ which is strict by Lemma 3.15. For future reference it might be useful to state the defining recurrence for a stream α produced by this operator, that is, $\alpha = Q(s)$:

$$\begin{aligned}
\alpha[0] &= T(s, F(0_{\mathcal{S}_A}))[0] \\
\alpha[t+1] &= T(s, F(\alpha|_{\leq t}))[t+1]
\end{aligned}$$

To prove that Q is causal we could use this recurrence and induction. Instead we use the causality of T and the idempotence of cutting in conjunction with Corollary 3.16 as follows:

$$\begin{aligned}
Q(s)|_{\leq t} &= (\mathrm{fix}\, \alpha.T(s, F(\alpha)))|_{\leq t} \\
&= \mathrm{fix}\, \alpha.(T(s, F(\alpha))|_{\leq t}) && \text{(Corollary 3.16)} \\
&= \mathrm{fix}\, \alpha.(T(s|_{\leq t}, F(\alpha)|_{\leq t})|_{\leq t}) && \text{(Causality of } T\text{)} \\
&= \mathrm{fix}\, \alpha.(T(s|_{\leq t}|_{\leq t}, F(\alpha)|_{\leq t})|_{\leq t}) && \text{(Idempotence of cutting)} \\
&= \mathrm{fix}\, \alpha.(T(s|_{\leq t}, F(\alpha))|_{\leq t}) && \text{(Causality of } T\text{)} \\
&= (\mathrm{fix}\, \alpha.T(s|_{\leq t}, F(\alpha)))|_{\leq t} && \text{(Corollary 3.16)} \\
&= Q(s|_{\leq t})|_{\leq t}
\end{aligned}$$

For time-invariance we observe that if $\alpha = T(s, F(\alpha))$ then $z^{-1}(\alpha) = z^{-1}(T(s, F(\alpha))) = T(z^{-1}(s), z^{-1}(F(\alpha)) = T(z^{-1}(s), F(z^{-1}(\alpha)))$. It follows that $\beta = z^{-1}(\alpha)$ satisfies the equation $\beta = T(z^{-1}(s), F(\beta))$ so, by uniqueness of the fixed point $Q(z^{-1}(s)) = z^{-1}(Q(s))$. □

Ostensibly this covers a form of straightforward recursion but how about mutual recursion? For instance, assuming T_1, T_2 are causal and F_1, F_2 are strict we wish to claim that the following is well defined: $Q_1(s_1, s_2) = \alpha_1$ and $Q_2(s_1, s_2) = \alpha_2$ where:

$$\begin{aligned}
\alpha_1 &= T_1(s_1, F_1(\alpha_2)) \quad \text{and} \\
\alpha_2 &= T_2(s_2, F_2(\alpha_1))
\end{aligned}$$

Here is the corresponding diagram:

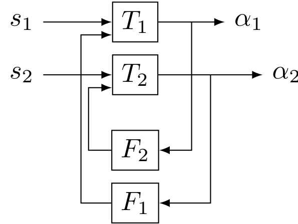

In section 3.1.2 we defined constructions on pairs of streams/streams of pairs. In particular, for binary $T_1 : \mathcal{S}_{A_1} \times \mathcal{S}_{B_1} \to \mathcal{S}_{C_1}$ and binary $T_2 : \mathcal{S}_{A_2} \times \mathcal{S}_{B_2} \to \mathcal{S}_{C_2}$ let binary $T_1 \times T_2 : \mathcal{S}_{A_1 \times A_2} \times \mathcal{S}_{B_1 \times B_2} \to \mathcal{S}_{C_1 \times C_2}$

VAL: For binary operators we should have defined $\times$ this way in section 3.1.2 ... Also, in the def of $T_1 \times T_2$ note that I did not put a lift in front of the stream pairing operator on the RHS.

$$[T_1 \times T_2](p, q) = \langle T_1(\uparrow\text{fst} \circ p, \uparrow\text{fst} \circ q), T_2(\uparrow\text{snd} \circ p, \uparrow\text{snd} \circ q)\rangle$$

In addition, we use $F_1 \times F_2 : \mathcal{S}_{C_1 \times C_2} \to \mathcal{S}_{B_1 \times B_2}$ as defined in section 3.1.2. Let also swap : $C_1 \times C_2 \to C_2 \times C_1$ be the operator that swaps the components of a pairs (obtained by pairing the second projection with the third.

Proposition 3.18. If T_1 and T_2 are causal then $T_1 \circ T_2$ is causal. If F_1 and F_2 are strict then $(F_1 \times F_2) \circ \uparrow\text{swap}$ is strict.

The circuit above is equivalent to the following (when composed with projections of outputs and pairing of inputs:

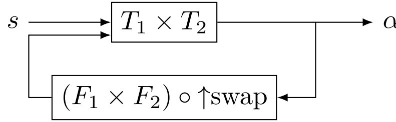

In other words, we can apply Corollary 3.17 to the causal operator $T_1 \times T_2$ and the strict operator $F_1 \times F_2$ and obtain Q_1 and Q_2 from

$$\text{fix}\, \alpha.[T_1 \times T_2](\langle s_1, s_2\rangle, [F_1 \times F_2](\text{swap}(\alpha)))$$

by further projecting, where α is a variable of type pair of streams and $\overline{\alpha}$ swaps the two components.

VAL: We should state this a some kind of corollary to the corollary.

MIHAI: This becomes complicated for n-way mutual recursion, you have a quadratic number of edges. Maybe it's simpler to assume that all of them use all of the alphas, and the projection to a subset is part of T if needed.

3.4 Streams as an abelian group

Remember that we require the elements of a stream to come from an Abelian group: $(A, +, 0, -)$. This structure also lifts to streams:

Proposition 3.19. $(\mathcal{S}_A, +, 0_{\mathcal{S}_A}, -)$ (with the operations lifted pointwise in time) is also an abelian group. Moreover, lifting a group homomorphism produces a stream operator that is itself a group homomorphism. In addition, when A and B are abelian groups there is a standard abelian group structure on $A \times B$, with the zero $0_{A \times B}$ the pair $\langle 0_A, 0_B\rangle$.

Definition 3.20 (linear)**.** If A and B are abelian groups, we call a function $f : A \to B$ **linear** if it is a group homomorphism, that is, $f(a+b) = f(a)+f(b)$ (and therefore $f(0_A) = 0_B$ and $f(-a) = -f(a)$); thus zpp(f).

We use the abbreviation LTI for a stream operator that is linear and time-invariant.

Lifting a linear function $f : A \to B$ produces a stream operator $\uparrow f$ that is causal and LTI. It follows that stream addition and negation are causal and LTI. z^{-1} is LTI, (and so is z^{-k} for all k).

Definition 3.21 (multilinear, bilinear)**.** We define **multilinear** (in particular, **bilinear**) functions as functions (between groups) of multiple arguments that are linear separately in each argument (that is, if we fix all but one argument, the resulting function is linear in that argument. In other words, the function distributes over addition): e.g., for $g : A_0 \times A_1 \to B, \forall a, b \in A_0, c, d \in A_1 . g(a + b, c) = g(a, c) + g(b, c)$, and $g(a, c + d) = g(a, c) + g(c, d)$.

Multiplication over $\mathbb{Z}$ is a bilinear function. For a bilinear function g we have $\text{zpp}(g)$.

This definition extends to stream operators. Lifting any bilinear function $g : A \times B \to C$ produces a bilinear stream operator $\uparrow g$. An example bi-linear operator over $\mathcal{S}_\mathbb{Z}$ is the lifted integer multiplication: $T : \mathcal{S}_\mathbb{Z} \times \mathcal{S}_\mathbb{Z} \to \mathcal{S}_\mathbb{Z}, T(a, b)[t] = a[t] \cdot b[t]$.

The composition of multilinear operators with linear operators is multilinear (since homomorphisms compose). Since linear and bilinear functions have the zero-preservation property, lifted linear and bilinear functions operators are all time-invariant.

Proposition 3.22. The composition of a bilinear operator followed by a linear operator is a bilinear operator.

Proof. Consider $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ a bilinear operator, and $S : \mathcal{S}_C \to \mathcal{S}_C$, a linear operator. Let us compute $S(T(a + b, c)) = S(T(a, c) + T(b, c)) = S(T(a, c)) + S(T(b, c))$. Thus $S \circ T$ is bilinear. $\square$

Lifting a multilinear operator $A_1 \times \cdots \times A_n \to B$ produces a multilinear, time-invariant stream operator. Although combining pairs (tuples) of streams into stream of pairs (tuples) can be useful we must note a distinction (well understood in algebra): we have seen that we can use instead of a binary stream operator $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ a unary version that acts on streams of pairs $T^u : \mathcal{S}_{A \times B} \to \mathcal{S}_C$ where $T^u\langle a, b\rangle = T(a, b)$. However, in contrast to causality, there is, in general, no relation between the linearity.

In traditional signal processing most operators are LTI but in our development we will use some important non-linear ones.

The "feedback-loop" operators defined by recurrence in Corollary 3.17, e.g., $\lambda s.\text{fix}\, \alpha.T(s, F(\alpha))$ are, in general, not (multi)linear. However, multilinearity holds in important particular cases, as shown in the following proposition:

VAL: Counterexample even for bilinear S and $\lambda s_1.\lambda s_2.\text{fix}\, \alpha.S(s_1, s_2 + z^{-1}(\alpha))$? Possibly S is join NOT followed by distinct?

MIHAI: Notice that this is a very nice kind of recursion, a tail-recursion.

Proposition 3.23. Let S be a unary causal, LTI operator. Then, the operator $Q(s) = \text{fix}\, \alpha.S(s + z^{-1}(\alpha))$ is well-defined and LTI.

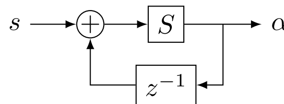

Proof. Since S and the addition operator are causal and z^{-1} is strict, Proposition 3.14 applies, for each s, to the operator $\lambda\alpha.S(s + z^{-1}(\alpha))$ which is strict by Lemma 3.15. Thus Q is well-defined.

Fix streams s_0 and s_1. Let α_0 be the unique solution of $\alpha_0 = S(s_0+z^{-1}(\alpha_0))$ and α_1 be the unique solution of $\alpha_1 = S(s_1 + z^{-1}(\alpha_1))$. Then $\alpha = \alpha_0 + \alpha_1$ is the unique solution of $\alpha = Q(s_0 + s_1 + z^{-1}(\alpha))$. This is justified by adding the equations and using the linearity of S and the linearity of z^{-1}. $\square$

Proposition 3.24. If T is binary causal, G is unary causal, and F is unary strict then:

$$\text{fix}\, \alpha.G(T(s, F(\alpha))) = G(\text{fix}\, \beta.T(s, F(G(\beta))))$$

In terms of diagrams:

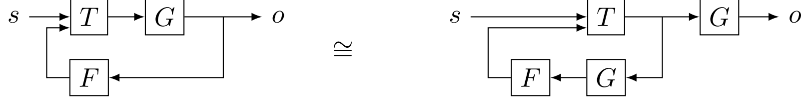

Proof. For the second fixpoint to exist we need to show that $F \circ G$ is strict. Once we do that, by uniqueness of solutions, it remains to show that if β is a solution to $\beta = T(s, F(G(\beta))$ then $\alpha = G(\beta)$ is a solution to $\alpha = G(T(s, F(\alpha))$ which is immediate by applying G to the first equation. So let's prove that $F \circ G$ is strict. For $t > 0$ we have

$$\begin{aligned}
F(G(s))|_{\leq t} &= F(G(s)|_{<t})|_{\leq t} && (F \text{ strict}) \\
&= F(G(s)|_{\leq t-1})|_{\leq t} && (t \geq 1) \\
&= F(G(s|_{\leq t-1}))|_{\leq t} && (G \text{ causal}) \\
&= F(G(s|_{<t}))|_{\leq t}
\end{aligned}$$

For $t = 0$ just observe that $F(G(s))[0]$ is the same for any $G(s)$ therefore for any s. Note that if G is not causal then $F \circ G$ is not strict and the fixed point may not exist. □

3.5 Differentiation and Integration

Definition 3.25 (Differentiation)**.** The **differentiation operator** $D_{\mathcal{S}_A} : \mathcal{S}_A \to \mathcal{S}_A$ is defined by:

$$D_{\mathcal{S}_A}(s) \stackrel{\text{def}}{=} s - z^{-1}(s)$$

We generally omit the type, and write just D when the type can be inferred from the context.

The value of $D(s)$ (at time t) is the difference between the current (time t) value of s and the previous (time $t - 1$) value of s.

As an example, applying D to the stream id gives a result a stream $D(id)$ containing the values $[\ 0\ \ 1\ \ 1\ \ 1\ \ 1\ \ \cdots\ ]$.

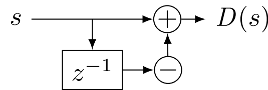

Proposition 3.26. The differentiation operator D is causal and LTI.

Proof. Follow from definition using the properties of subtraction and delay. □

The integration operator "reconstitutes" a stream from its changes:

Definition 3.27 (Integration)**.** The **integration operator** $I_{\mathcal{S}_A} : \mathcal{S}_A \to \mathcal{S}_A$ is defined by $I_{\mathcal{S}_A}(s) \stackrel{\text{def}}{=} \lambda s.\text{fix}\, \alpha.(s + z^{-1}(\alpha))$.

We also generally omit the type, and write just I. This is the construction from Proposition 3.23 using the identity function for S.

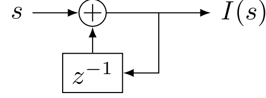

As an example, applying I to the id stream gives as result a stream $I(id)$ composed of the values $[\ 0 \quad 1 \quad 3 \quad 6 \quad 10 \quad \cdots\ ]$.

Proposition 3.28. $I(s)$ is the discrete (indefinite) integral applied to the stream s: $I(s)[t] = \sum_{i \leq t} s[i]$.

Proof. The recurrence from Corollary 3.17 specializes to

$$\begin{aligned} \alpha[0] &= s[0] \\ \alpha[t+1] &= \alpha[t] + s[t+1] \end{aligned}$$

and it's straightforward to check that $\alpha[t] = \sum_{i \leq t} s[i]$ satisfies it. □

Proposition 3.29 (Properties of I)**.** The integration operator I is causal and LTI.

Proof. By Proposition 3.28 these properties follow from Corollary 3.17 and Proposition 3.23. They also be checked directly using the definition by summation. □

Theorem 3.30 (Inversion)**.** The integration and differentiation operators are inverse to each other. Equivalently, for any streams α and s we have $\alpha = I(s)$ iff $D(\alpha) = s$.

Proof. This can be shown directly from the definitions, for example

$$\begin{aligned} D(I(s))[t] &= (I(s) - z^{-1}(I(s)))[t] && \text{definition of } D \\ &= \sum_{k \leq t} s[k] - z^{-1}(\sum_{k \leq t} s[k])[t] && \text{Property 3.29} \\ &= \sum_{k \leq t} s[k] - \sum_{k \leq t-1} s[k] && \text{definition of } z^{-1} \\ &= s[t] \end{aligned}$$

(and similarly we can show that $I(D(s'))[t] = s'[t]$).

Alternatively, the equivalent form of the theorem follows from Proposition 3.28 by observing that $D(\alpha) = s$ iff $\alpha = s + z^{-1}(\alpha)$ which is the equation on streams that defines $\alpha = I(s)$. □

So we have the following circuit equivalence:

$$s \to \boxed{I} \to \boxed{D} \to o \quad \cong \quad s \to o \quad \cong \quad s \to \boxed{D} \to \boxed{I} \to o$$

Since I and D are inverse to each other they are both bijections on streams.

If we define addition of pairs as adding the pair elements pointwise, we also have the following identities: $I(\langle s, t \rangle) = \langle I(s), I(t) \rangle$ and $D(\langle s, t \rangle) = \langle D(s), D(t) \rangle$.

Observation It is a standard algebraic fact that the inverse of a homomorphism is also a homomorphism. Thus, I is linear iff D is linear. Another immediate consequence of this theorem is that I is time-invariant iff D is time-invariant. Moreover I is causal iff D is causal. Therefore by the Inversion Theorem we could have stated and proved only one of Proposition 3.29 or Proposition 3.26.

Observation In digital signal processing, I is a IIR, an infinite-impulse response filter: given a cut stream it can produce an unbounded stream. D is a FIR, a finite-impulse response filter: from a cut stream it always produces a cut stream.

4 Relational algebra in DBSP

4.0.1 Generalizing box-and-arrow diagrams

From now on our circuits will mix computations on scalars and streams. We will use the same graphical representation for functions that compute on scalars: boxes with input and output arrows. The values on the the connecting arrows will be scalars instead of streams; otherwise the interpretation of boxes as function application is unchanged.

When connecting boxes the types of the arrows must match. E.g., the output of a box producing a stream cannot be connected to the input of a box consuming a scalar.

Results in Section 3 apply to streams of arbitrary group values. In this section we turn our attention to using these results in the context of relational databases.

However, we face a technical problem: the I and D operators were defined on abelian groups, and relational databases in general are not abelian groups, since they operate on sets. Fortunately, there is a well-known tool in the database literature which converts set operations into group operations by using $\mathbb{Z}$-sets (also called z-relations [24]) instead of sets.

We start by defining the $\mathbb{Z}$-sets group, and then we explain how relational queries are converted into DBSP circuits over $\mathbb{Z}$-sets.

4.1 $\mathbb{Z}$-sets as an abelian group

Given a set A we define **$\mathbb{Z}$-sets**[4] over A as functions with *finite support* from A to $\mathbb{Z}$ (i.e., which are 0 almost everywhere). These are functions $f : A \to \mathbb{Z}$

[4] Also called $\mathbb{Z}$-relations elsewhere [23].

where $f(x) \neq 0$ for at most a finite number of values $x \in A$. We also write $\mathbb{Z}[A]$ for the type of $\mathbb{Z}$-sets with elements from A. The values in $\mathbb{Z}[A]$ can also be thought as being key-value maps with keys in A and values in $\mathbb{Z}$, justifying the array indexing notation.

Since $\mathbb{Z}$ is an abelian group, $\mathbb{Z}[A]$ is also an abelian group. This group $(\mathbb{Z}[A], +_{\mathbb{Z}[A]}, 0_{\mathbb{Z}[A]}, -_{\mathbb{Z}A})$ has addition and subtraction defined pointwise:

$$(f +_{\mathbb{Z}[A]} g)(x) = f(x) + g(x). \forall x \in A.$$

The 0 element of $\mathbb{Z}[A]$ is the function $0_{\mathbb{Z}[A]}$ defined by $0_{\mathbb{Z}[A]}(x) = 0. \forall x \in A$. (In fact, since $\mathbb{Z}$ is a ring, $\mathbb{Z}[A]$ is also ring, endowed with a multiplication operation, also defined pointwise.)

A particular $\mathbb{Z}$-set $m \in \mathbb{Z}[A]$ can be denoted by enumerating the inputs that map to non-zero values and their multiplicities: $m = \{x_1 \mapsto w_1, \ldots, x_n \mapsto w_n\}$. We call $w_i \in Z$ the **multiplicity** (or weight) of $x_i \in A$. Multiplicities can be negative. We write that $x \in m$ for $x \in A$, iff $m[x] \neq 0$.

For example, let's consider a concrete $\mathbb{Z}$-set $R \in \mathbb{Z}[\texttt{string}]$, defined by $R = \{\texttt{joe} \mapsto 1, \texttt{anne} \mapsto -1\}$. R has two elements in its domain, `joe` with a multiplicity of 1 (so $R[\texttt{joe}] = 1$), and `anne` with a multiplicity of -1. We say $\texttt{joe} \in R$ and $\texttt{anne} \in R$.

Given a $\mathbb{Z}$-set $m \in \mathbb{Z}[A]$ and a value $v \in A$, we overload the array index notation $m[v]$ to denote the multiplicity of the element v in m. Thus we write $R[\texttt{anne}] = -1$. When $c \in \mathbb{Z}$, and $v \in A$ we also write $c \cdot v$ for the **singleton** $\mathbb{Z}$-set $\{v \mapsto c\}$. In other words, $3 \cdot \texttt{frank} = \{\texttt{frank} \mapsto 3\}$. We extend scalar multiplication to operate on $\mathbb{Z}$-sets: for $c \in Z, m \in \mathbb{Z}[A]$, $c \cdot m \stackrel{\text{def}}{=} \sum_{x \in m} (c \cdot m[x]) \cdot x$. We then have $2 \cdot R = \{\texttt{joe} \mapsto 2, \texttt{anne} \mapsto -2\}$: multiplying each row weight by 2.

We define the **size** of a $\mathbb{Z}$-set as the size of its support set, and we use the modulus symbol to represent the size: $|m| \stackrel{\text{def}}{=} \sum_{x \in m} 1$. So $|R| = 2$.

4.2 Sets, bags, and $\mathbb{Z}$-sets

$\mathbb{Z}$-sets generalize sets and bags. Given a set with elements from A, it can be represented as a $\mathbb{Z}$-set $\mathbb{Z}[A]$ by associating a weight of 1 with each set element. The function tozset : $2^A \to \mathbb{Z}[A]$, defined as $\text{tozset}(s) = \sum_{x \in s} 1 \cdot x$, converts a set to a $\mathbb{Z}$-set by associating a multiplicity of 1 with each set element. Thus $\text{tozset}(\{\texttt{joe}, \texttt{anne}\}) = \{\texttt{joe} \mapsto 1, \texttt{anne} \mapsto 1\}$.

Definition 4.1. We say that a $\mathbb{Z}$-set represents a **set** if the multiplicity of every element is one. We define a function to check this property isset : $\mathbb{Z}[A] \to \mathbb{B}$, given by:

$$\text{isset}(m) \stackrel{\text{def}}{=} \begin{cases} \text{true} & \text{if } m[x] = 1, \forall x \in m \\ \text{false} & \text{otherwise} \end{cases}$$

For our example $\text{isset}(R) = \text{false}$, since $R[\texttt{anne}] = -1$. $\text{isset}(\text{tozset}(m)) = \text{true}$ for any set $m \in 2^A$.

Definition 4.2. We say that a $\mathbb{Z}$-set is **positive** (or a **bag**) if the multiplicity of every element is positive. We define a function to check this property ispositive : $\mathbb{Z}[A] \to \mathbb{B}$, given by

$$\text{ispositive}(m) \stackrel{\text{def}}{=} \begin{cases} \text{true} & \text{if } m[x] \geq 0, \forall x \in A \\ \text{false} & \text{otherwise} \end{cases}$$

For our example ispositive(R) = false, since $R[\texttt{anne}] = -1$, but isset$(m) \Rightarrow$ ispositive$(m).\forall m \in \mathbb{Z}[A]$.

We also write $m \geq 0$ when m is positive. For positive m, n we write $m \geq n$ for $m, n \in \mathbb{Z}[A]$ iff $m - n \geq 0$. The relation $\geq$ is a partial order.

Definition 4.3. The function *distinct* : $\mathbb{Z}[A] \to \mathbb{Z}[A]$ projects a $\mathbb{Z}$-set into an underlying set (but *the result is still a* $\mathbb{Z}$*-set*). The definition is $\forall x \in A$

$$\mathit{distinct}(m)[x] \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } m[x] > 0 \\ 0 & \text{otherwise} \end{cases}$$

$\mathit{distinct}(R) = \{\texttt{joe} \mapsto 1\}$.

distinct "removes" elements with negative multiplicities. zpp(*distinct*).

Circuits derived from relational program will only operate with positive $\mathbb{Z}$-sets; non-positive values will be only used to represent *changes* to $\mathbb{Z}$-sets (a change with negative weights will remove elements from a $\mathbb{Z}$-set).

Proposition 4.4. *distinct* is idempotent: $\mathit{distinct} = \mathit{distinct} \circ \mathit{distinct}$.

Proposition 4.5. For any $m \in \mathbb{Z}[A]$ we have: isset$(\mathit{distinct}(m))$ and ispositive$(\mathit{distinct}(m))$.

We call a function $f : \mathbb{Z}[I] \to \mathbb{Z}[O]$ **positive** if $\forall x \in \mathbb{Z}[I], x \geq 0_{\mathbb{Z}[I]} \Rightarrow f(x) \geq 0_{\mathbb{Z}[0]}$. We extend the notation used for $\mathbb{Z}$-sets for functions as well: ispositive(f).

Correctness of the DBSP implementations The function toset : $\mathbb{Z}[A] \to 2^A$, defined as toset$(m) = \cup_{x \in \mathit{distinct}(m)} \{x\}$, converts a $\mathbb{Z}$-set into a set.

A relational query f that transforms a set V into a set U will be implemented by a DBSP computation f' on $\mathbb{Z}$-sets. The correctness of the implementation requires that the following diagram commutes:

$$\begin{array}{ccc} V & \xrightarrow{\quad f \quad} & U \\ {\scriptstyle \text{tozset}}\downarrow & & \uparrow{\scriptstyle \text{toset}} \\ VZ & \xrightarrow{\quad f' \quad} & UZ \end{array}$$

Remark: We can generalize the notion of $\mathbb{Z}$-sets to functions $m : A \to \mathbf{G}$ for a ring $\mathbf{G}$ other than $\mathbb{Z}$. The properties we need from the ring structure are the following: the ring must be a commutative group (needed for defining I, D, and z^{-1}), the multiplication operation must distribute over addition (needed to define Cartesian products), and there must be a notion of positive values, needed to define the *distinct* function. Rings such as $\mathbb{Q}$ or $\mathbb{R}$ would work perfectly.

4.3 Streams over $\mathbb{Z}$-sets

Since all the results from Section 3 are true for streams over an arbitrary abelian group, they extend to streams where the elements are $\mathbb{Z}$-sets. In the rest of this text we only consider streams of the form $\mathcal{S}_{\mathbb{Z}[A]}$, for some element type A.

An example of a stream of $\mathbb{Z}$-sets is
$s = [\ 0 \quad R \quad -1 \cdot R \quad 2 \cdot R \quad -2 \cdot R \quad \cdots\ ]$. We have $s[2] = -R = \{\texttt{joe} \mapsto -1, \texttt{anne} \mapsto 1\}$.

Definition 4.6. A stream $s \in \mathcal{S}_{\mathbb{Z}[A]}$ is **positive** if every value of the stream is positive: $s[t] \geq 0 . \forall t \in \mathbb{N}$.

Definition 4.7. A stream $s \in \mathcal{S}_{\mathbb{Z}[A]}$ is **monotone** if $s[t] \geq s[t-1], \forall t \in \mathbb{N}$.

Lemma 4.8. Given a positive stream $s \in \mathcal{S}_{\mathbb{Z}[A]}$ the stream $I(s)$ is monotone.

Proof. Let us compute $I(s)[t+1] - I(s)[t] = \sum_{i \leq t+1} s[i] - \sum_{i \leq t} s[i] = s[t+1] \geq 0$, by commutativity and positivity of s. □

Lemma 4.9. Given a monotone stream $s \in \mathcal{S}_{\mathbb{Z}[A]}$, the elements of the stream $D(s)$ are positive.

Proof. By the definition of monotonicity $s[t+1] \geq s[t]$. By definition of D we have $D(s)[t+1] = s[t+1] - s[t] \geq 0$. □

4.4 Implementing the relational algebra

The fact that the relational algebra can be implemented by computations on $\mathbb{Z}$-sets has been shown before, e.g. [24]. The translation of the core relational operators is summarized in Table 1 and discussed below.

The translation is fairly straightforward, but many operators require the application of a *distinct* to produce sets. The correctness of this implementation is predicated on the global circuit inputs being sets as well.

4.4.1 Query composition

A composite query is translated by compiling each sub-query separately into a circuit and composing the respective circuits.

For example, consider the following SQL query:

```
SELECT ... FROM (SELECT ... FROM ...)
```

given circuits C_O implementing the outer query and C_I implementing the inner query, the translation of the composite query is:

$$\texttt{I} \rightarrow \boxed{C_I} \rightarrow \boxed{C_O} \rightarrow \texttt{O}$$

We have $\text{ispositive}(C_I) \wedge \text{ispositive}(C_0) \Rightarrow \text{ispositive}(C_O \circ C_I)$ and $\text{zpp}(C_I) \wedge \text{zpp}(C_O) \Rightarrow \text{zpp}(C_O \circ C_I)$.

Operation	SQL example	DBSP circuit
Composition	`SELECT DISTINCT ... FROM` `(SELECT ... FROM ...)`	I → C_I → C_O → O

Union	`(SELECT * FROM I1)`
	`UNION`
	`(SELECT * FROM I2)`
		I1, I2 → $+$ → *distinct* → O

Projection	`SELECT DISTINCT I.c` `FROM I`	I → π → *distinct* → O
Filtering	`SELECT * FROM I` `WHERE p(I.c)`	I → σ_P → *distinct* → O
Selection	`SELECT DISTINCT f(I.c, ...)` `FROM I`	I → map(f) → *distinct* → O
Cartesian product	`SELECT I1.*, I2.*` `FROM I1, I2`	I1, I2 → $\times$ → O

Join	`SELECT I1.*, I2.*`
	`FROM I1 JOIN I2`
	`ON I1.c1 = I2.c2`
		I1, I2 → $\bowtie$ → O

Intersection	`(SELECT * FROM I1)`
	`INTERSECT`
	`(SELECT * FROM I2)`
		I1, I2 → $\bowtie$ → O

Difference	`SELECT * FROM I1`
	`EXCEPT`
	`SELECT * FROM I2`
		I1, (I2 → $-$) → $+$ → *distinct* → O

Table 1: Implementation of SQL relational set operators in DBSP. Each query assumes that inputs `I`, `I1`, `I2`, are sets and it produces output sets.

4.4.2 Set union

Consider the following SQL query:

```
(SELECT * FROM I1) UNION (SELECT * FROM I2)
```

The following circuit implements the union program:

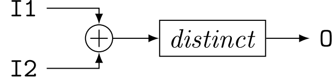

Given $\mathbb{Z}$-sets $a, b \in \mathbb{Z}[I]$ s.t. isset(a) and isset(b), their *set union* can be computed as: $\cup : \mathbb{Z}[I] \times \mathbb{Z}[I] \to \mathbb{Z}[I]$,

$$a \cup b \stackrel{\text{def}}{=} \mathit{distinct}(a +_{\mathbb{Z}[I]} b).$$

The *distinct* application is necessary to provide the set semantics.

4.4.3 Projection

Consider a query such as:

```
SELECT I.c FROM I
```

We can assume without loss of generality that table I has two columns, and that a single column is preserved in the projection. Hence the type of I is $\mathbb{Z}[A_0 \times A_1]$ while the result has type is $\mathbb{Z}[A_0]$. In terms of $\mathbb{Z}$-sets, the projection of a $\mathbb{Z}$-set i on A_0 is defined as:

$$\pi(i)[y] = \sum_{x \in i, x|_0 = y} i[x]$$

where $x|_0$ is first component of the tuple x. In other words, the multiplicity of a tuple in the result is the sum of the multiplicities of all input tuples that project to it.

The circuit for a projection query is:

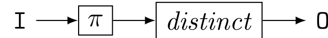

The *distinct* is necessary to convert the result to a set.
π is linear; ispositive(π), zpp(π).

4.4.4 Selection

We generalize the SQL selection operator to allow it to apply an arbitrary function to each row of the selected set. Given a function $f : A \to B$, the mathematical **map** operator "lifts" the function f to operate on $\mathbb{Z}$-sets: map(f) : $\mathbb{Z}[A] \to \mathbb{Z}[B]$. A map operator appears in SQL due to the use of expressions in the SELECT clause, as in the following example:

```
SELECT f(I.c) FROM I
```

The circuit implementation of this query is:

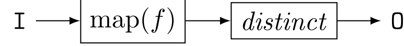

For any function f we have the following properties: $\text{map}(f)$ is linear, $\text{ispositive}(\text{map}(f)), and \text{zpp}(\text{map}(f))$.

4.4.5 Filtering

Filtering occurs in SQL through a WHERE clause, as in the following example:

```
SELECT * FROM I WHERE p(I.c)
```

Let us assume that we are filtering with a predicate $P : A \rightarrow \mathbb{B}$. We define the following function $\sigma_P : \mathbb{Z}[A] \rightarrow \mathbb{Z}[A]$ as:

$$\sigma_P(m)[t] = \begin{cases} m[t] \cdot t & \text{if } P(t) \\ 0 & \text{otherwise} \end{cases}$$

The circuit for filtering with a predicate P is:

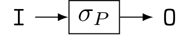

For any predicate P we have $\text{isset}(i) \Rightarrow \text{isset}(\sigma_P(i))$ and $\text{ispositive}(\sigma_P)$. Thus a *distinct* is not needed. σ_P is linear and $\text{zpp}(\sigma_P)$.

4.4.6 Cartesian products

Consider this SQL query performing a Cartesian product between sets `I1` and `I2`:

```
SELECT I1.*, I2.* FROM I1, I2
```

We first define a product operation on $\mathbb{Z}$-sets. For $a \in \mathbb{Z}[A]$ and $b \in \mathbb{Z}[B]$ we define $a \times b \in \mathbb{Z}[A \times B]$ by

$$(a \times b)((x, y)) \stackrel{\text{def}}{=} a[x] \times b[y]. \forall x \in a, y \in b.$$

The weight of a pair in the result is the product of the weights of the elements in the sources. The circuit for the query is:

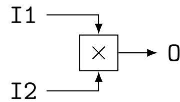

$\text{isset}(x) \wedge \text{isset}(y) \Rightarrow \text{isset}(x \times y)$. $\times$ is bilinear, $\text{ispositive}(\times), \text{zpp}(\times)$.

4.4.7 Joins

As is well-known, joins can be modeled as Cartesian products followed by filtering. Since a join is a composition of a bilinear and a linear operator, it is also a bilinear operator. $\text{ispositive}(\bowtie), \text{zpp}(\bowtie)$.

In practice joins are very important computationally, and they are implemented by a built-in scalar function on $\mathbb{Z}$-sets:

$$(a \bowtie b)((x,y)) \stackrel{\text{def}}{=} a[x] \times b[y] \text{ if } x|_{c1} = y|_{c2}.$$

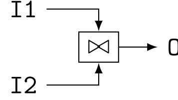

4.4.8 Set intersection

Set intersection is a special case of join, where both relations have the same schema. It follows that set intersection is bilinear, and has the zero-preservation property.

4.4.9 Set difference

Consider the following query:

```
SELECT * FROM I1 EXCEPT SELECT * FROM I2
```

We define the set difference on $\mathbb{Z}$-sets as follows: $\setminus : \mathbb{Z}[I] \times \mathbb{Z}[I] \to \mathbb{Z}[I]$, where

$$i_1 \setminus i_2 = \mathit{distinct}(i_1 - i_2).$$

Note that we have $\forall i_1, i_2, \text{ispositive}(i_1 \setminus i_2)$ due to the *distinct* operator. The circuit computing the above query is:

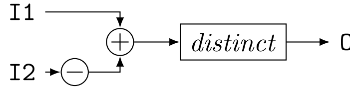

5 Incremental computation

In this section we formally define incremental computations over streams and analyze their properties.

Definition 5.1. Given a unary stream operator $Q : \mathcal{S}_A \to \mathcal{S}_B$ we define the **incremental version** of Q as $Q^\Delta \stackrel{\text{def}}{=} D \circ Q \circ I$. Q^Δ has the same "type" as Q: $Q^\Delta : \mathcal{S}_A \to \mathcal{S}_B$. For an operator with multiple inputs we define the incremental version by applying I to each input independently: e.g., if $T : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ then $T^\Delta : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C$ and $T^\Delta(a,b) \stackrel{\text{def}}{=} D(T(I(a), I(b)))$.

The following diagram illustrates the intuition behind this definition:

$$\Delta s \rightarrow \boxed{I} \xrightarrow{s} \boxed{Q} \xrightarrow{o} \boxed{D} \rightarrow \Delta o$$

If $Q(s) = o$ is a computation, then Q^Δ performs the "same" computation as Q, but between streams of changes Δs and Δo. This is the diagram from the introduction, substituting Δs for the transaction stream T, and o for the stream of view versions V.

Notice that our definition of incremental computation is meaningful only for *streaming* computations; this is in contrast to classic definitions, e.g. [26] which consider only one change. Generalizing the definition to operate on streams gives us additional power, especially when operating with recursive queries.

The following proposition is one of our central results.

Proposition 5.2. (Properties of the incremental version):
For computations of appropriate types, the following hold:

inversion: $Q \mapsto Q^\Delta$ is bijective; its inverse is $Q \mapsto I \circ Q \circ D$.

invariance: $+^\Delta = +, (z^{-1})^\Delta = z^{-1}, -^\Delta = -, I^\Delta = I, D^\Delta = D$

push/pull: $Q \circ I = I \circ Q^\Delta$; $D \circ Q = Q^\Delta \circ D$

chain: $(Q_1 \circ Q_2)^\Delta = {Q_1}^\Delta \circ {Q_2}^\Delta$ (This generalizes to operators with multiple inputs.)

add: $(Q_1 + Q_2)^\Delta = {Q_1}^\Delta + {Q_2}^\Delta$

cycle: $(\lambda s. \text{fix}\, \alpha. T(s, z^{-1}(\alpha)))^\Delta = \lambda s. \text{fix}\, \alpha. T^\Delta(s, z^{-1}(\alpha))$

Proof. The inversion and push-pull properties follow straightforwardly from the fact that I and D are inverses of each other.

For proving invariance we have $+^\Delta(a, b) \stackrel{\text{def}}{=} D(I(a) + I(b)) = a + b$, due to linearity of I. $-^\Delta(a) = D(-I(a)) = D(0 - I(a)) = D(I(0) - I(a)) = D(I(0 - a)) = -a$, also due to linearity of I.

The chain rule follows from push-pull. Indeed,

$$I \circ {Q_1}^\Delta \circ {Q_2}^\Delta = Q_1 \circ I \circ {Q_2}^\Delta = Q_1 \circ Q_2 \circ I$$

I.e., we have the following sequence of equivalent circuits:

$$i \rightarrow \boxed{I} \rightarrow \boxed{Q_1} \rightarrow \boxed{Q_2} \rightarrow \boxed{D} \rightarrow o \quad \cong$$

$$i \rightarrow \boxed{I} \rightarrow \boxed{Q_1} \rightarrow \boxed{D} \rightarrow \boxed{I} \rightarrow \boxed{Q_2} \rightarrow \boxed{D} \rightarrow o \quad \cong$$

$$i \rightarrow \boxed{{Q_1}^\Delta} \rightarrow \boxed{{Q_2}^\Delta} \rightarrow o$$

Here is a version of the chain rule with a binary operator:

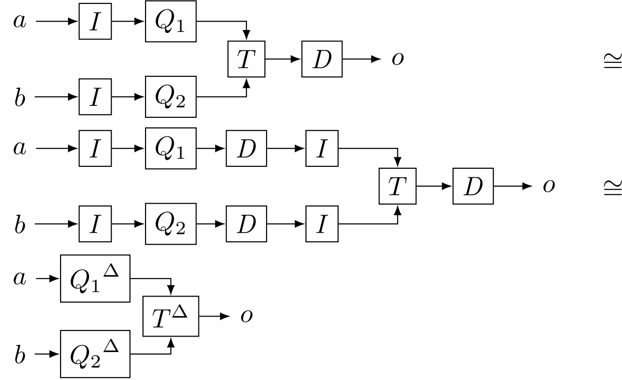

The add rule follows from push/pull and the linearity of I (or D). Indeed,

$$I \circ ({Q_1}^{\Delta} + {Q_2}^{\Delta}) = I \circ {Q_1}^{\Delta} + I \circ {Q_2}^{\Delta} = Q_1 \circ I + Q_2 \circ I = (Q_1 + Q_2) \circ I$$

I.e., the following diagrams are equivalent:

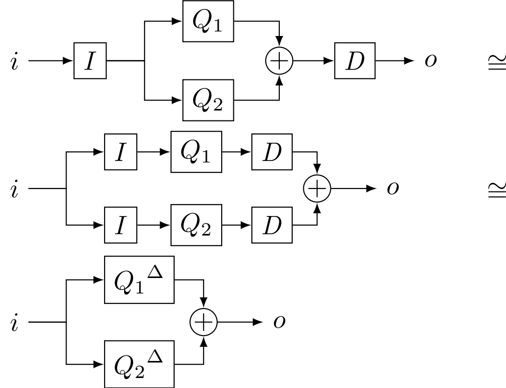

The cycle rule is most interesting. First, observe that if T is causal then so is T^{Δ} thus both sides of the equality are well-defined. Next, we can use again push/pull to show the equality if we can check that

$$I \circ (\lambda s.\text{fix}\, \alpha.T^{\Delta}(s, z^{-1}(\alpha)) = (\lambda s.\text{fix}\, \alpha.T(s, z^{-1}(\alpha)) \circ I$$

that is, for any s,

$$I(\text{fix}\, \alpha.T^{\Delta}(s, z^{-1}(\alpha)) = \text{fix}\, \alpha.T(I(s), z^{-1}(\alpha))$$

This follows from the following lemma. □

Lemma 5.3. If the parameters a and b are related by $b = D(a)$ (equivalently $a = I(b)$) then the unique solutions of the fixed point equations

$$\alpha \;=\; T(a, z^{-1}(\alpha)) \quad \text{and} \quad \beta \;=\; T^{\Delta}(b, z^{-1}(\beta))$$

are related by $\alpha = I(\beta)$ (equivalently $\beta = D(\alpha)$).

Proof. (Of Lemma 5.3) Let β be the unique solution of $\beta = T^{\Delta}(D(a), z^{-1}(\beta))$. We verify that $\alpha = I(\beta)$ satisfies the equation $\alpha = T(a, z^{-1}(\alpha))$. Indeed, using the fact that I and D are inverses as well as the time-invariance of I we have

$$\begin{aligned} I(\beta) &= I(T^{\Delta}(D(a), z^{-1}(\beta))) \\ &= I(D(T(I(D(a)), I(z^{-1}(\beta))))) \\ &= T(a, I(z^{-1}(\beta)) \\ &= T(a, z^{-1}(I(\beta)) \end{aligned}$$

I.e., starting from this diagram we apply a sequence of term-rewriting semantics-preserving transformations:

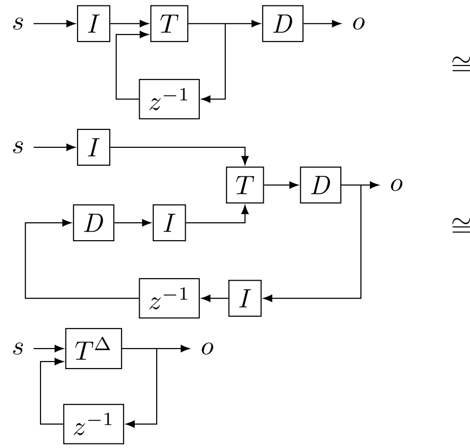

∎

If we specialize the above formula for the case $T(a, b) = Q(a + b)$ (for some time-invariant operator Q), by us the linearity of I we get that:

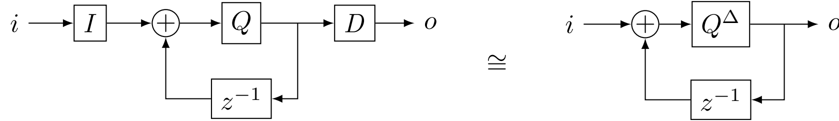

Theorem 5.4 (Linear)**.** For any LTI operator Q we have $Q^{\Delta} = Q$.

Proof. By the push/pull rule from Proposition 5.2 it suffices to show that Q commutes with differentiation:

$$\begin{aligned} D(Q(s)) &= Q(s) - z^{-1}(Q(s)) && \text{by definition of } D \\ &= Q(s) - Q(z^{-1}(s)) && \text{by time-invariance of } Q \\ &= Q(s - z^{-1}(s)) && \text{by linearity of } Q \\ &= Q(D(s)) && \text{by definition of } D. \end{aligned}$$

∎

As we have shown, the incremental version of a linear unary operator equals the operator itself. However, this is not true, in general, for multilinear operators. Nonetheless, there is a useful relationship between the incremental version of a multilinear operator and the operator itself. We illustrate with bilinear operators.

Theorem 5.5 (Bilinear)**.** For any bilinear time-invariant operator $\times$ we have $(a \times b)^{\Delta} \; = \; a \times b \; + \; I(z^{-1}(a)) \times b \; + \; a \times I(z^{-1}(b))$.

By rewriting this statement using Δa for the stream of changes to a we get the familiar formula for incremental joins: $\Delta(a \times b) = \Delta a \times \Delta b + a \times (\Delta b) + (\Delta a) \times b$.

In other words, the following three diagrams are equivalent:

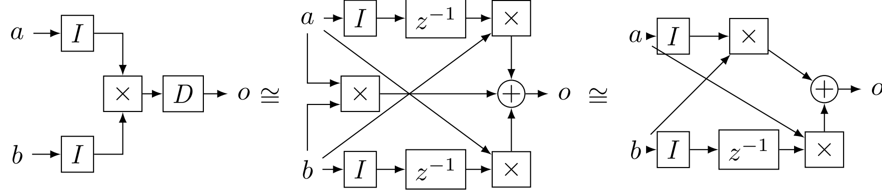

Proof.

$$
\begin{aligned}
(a \times b)^{\Delta} &= D(I(a) \times I(b)) && \text{def of } \cdot^{\Delta} \\
&= (I(a) \times I(b)) \; - \; z^{-1}(I(a) \times I(b)) && \text{def of } D \\
&= I(a) \times I(b) \; - \; z^{-1}(I(a)) \times z^{-1}(I(b)) && \times \text{ time inv.} \\
&= (a + z^{-1}(I(a))) \times (b + z^{-1}(I(b))) \; - \; z^{-1}(I(a)) \times z^{-1}(I(b)) && I \text{ fixpoint equation} \\
&= a \times b + z^{-1}(I(a)) \times b + a \times z^{-1}(I(b)) + z^{-1}(I(a)) \times z^{-1}(I(b)) \\
&\qquad - \; z^{-1}(I(a)) \times z^{-1}(I(b)) && \text{bilinearity} \\
&= a \times b \; + \; z^{-1}(I(a)) \times b \; + \; a \times z^{-1}(I(b)) && \text{cancel} \\
&= a \times b \; + \; I(z^{-1}(a)) \times b \; + \; a \times I(z^{-1}(b)) && I \text{ time inv.} \\
&= (a + I(z^{-1}(a))) \times b \; + \; a \times I(z^{-1}(b)) && \text{commutativity} \\
&= I(a) \times b \; + \; a \times I(z^{-1}(b)) && \text{def of } I
\end{aligned}
$$

□

6 Incremental relational queries

We start by giving a few rules that can be used to optimize relational DBSP circuits. Later we show how DBSP relational circuits can be converted to incremental versions.

6.1 Optimizing *distinct*

All standard algebraic properties of the relational operators can be used to optimize circuits. In addition, a few optimizations are related to the *distinct* operator, which is not linear, and thus expensive to incrementalize:

Proposition 6.1. Let Q be one of the following $\mathbb{Z}$-sets operators: filtering σ, join $\bowtie$, or Cartesian product $\times$. Then we have $\forall i \in \mathbb{Z}[I], \text{ispositive}(i) \Rightarrow Q(\mathit{distinct}(i)) = \mathit{distinct}(Q(i))$.

$$i \to \boxed{\mathit{distinct}} \to \boxed{Q} \to o \quad \cong \quad i \to \boxed{Q} \to \boxed{\mathit{distinct}} \to o$$

This rule allows us to delay the application of *distinct*.

Proposition 6.2. Let Q be one of the following $\mathbb{Z}$-sets operators: filtering σ, projection π, selection (map(f)), addition $+$, join $\bowtie$, or Cartesian product $\times$. Then we have $\forall i \in \mathbb{Z}[I], \text{ispositive}(i) \Rightarrow \mathit{distinct}(Q(\mathit{distinct}(i))) = \mathit{distinct}(Q(i))$.

This is Proposition 6.13 in [23].

$$i \to \boxed{\mathit{distinct}} \to \boxed{Q} \to \boxed{\mathit{distinct}} \to o \quad \cong$$
$$i \to \boxed{Q} \to \boxed{\mathit{distinct}} \to o$$

These properties allow us to "consolidate" distinct operators by performing one distinct at the end of a chain of computations.

Proposition 6.3. The following circuit implements $(\uparrow\!\mathit{distinct})^\Delta$:

$$d \to \boxed{(\uparrow\!\mathit{distinct})^\Delta} \to o \quad \cong \quad d \to \boxed{I} \to \boxed{z^{-1}} \xrightarrow{i} \boxed{\uparrow\!H} \to o \;\text{(with } d \text{ also fed into } \uparrow\!H\text{)}$$

where $H : \mathbb{Z}[A] \times \mathbb{Z}[A] \to \mathbb{Z}[A]$ is defined as:

$$H(i, d)[x] \stackrel{\text{def}}{=} \begin{cases} -1 & \text{if } i[x] > 0 \text{ and } (i + d)[x] \leq 0 \\ 1 & \text{if } i[x] \leq 0 \text{ and } (i + d)[x] > 0 \\ 0 & \text{otherwise} \end{cases}$$

The function H detects whether the multiplicity of an element in the input set i is changing from negative to positive or vice-versa.

6.1.1 Anti-joins

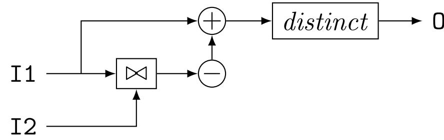

This can be optimized as follows:

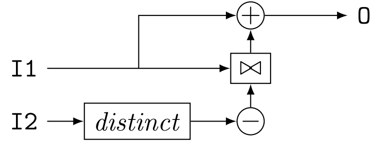

6.2 Parallelization

DBSP is rich enough to express operators such as the Volcano exchange operator [21], which can be used to parallelize DBSP circuits. The following circuit shows a parallel implementation of a join operator.

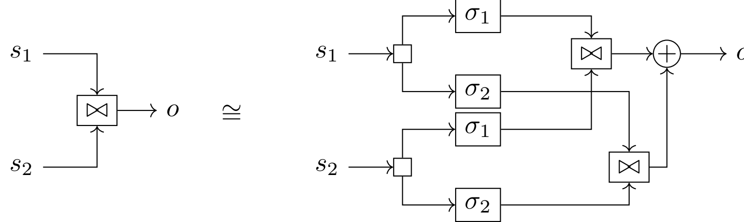

Here is an example of an exchange operator which repartitions the data in a collection from two partitions to three partitions, where σ_i for $i \in [3]$ are disjoint selection functions that partition the space of tuples.

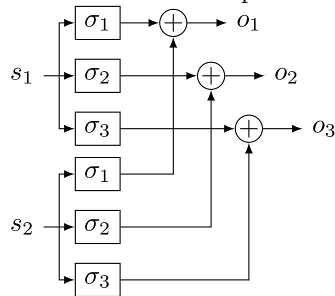

6.3 Incremental relational queries

Let us consider a relational query Q defining a view. To create a circuit that maintains incrementally the view defined by Q we apply the following mechanical steps:

Algorithm 6.4 (incremental view maintenance)**.**

1. Translate Q into a circuit using the rules in Table 1.
2. Apply optimization rules, including *distinct* consolidation.
3. Lift the whole circuit, by applying Proposition 3.3, converting it to a circuit operating on streams.
4. Incrementalize the whole circuit "surrounding" it with I and D.
5. Apply the chain rule and other properties of the $\cdot^{\Delta}$ operator from Proposition 5.2 to optimize the incremental implementation.

It is known that a query can be implemented by multiple plans, with varying data-dependent costs. The input provided to this algorithm is a standard relational query plan, and this algorithm produces an incremental plan that is "similar" to the input plan[5]. Step (2) generates an equivalent circuit, with possibly fewer *distinct* operators (the result is deterministic no matter the order of elimination). Step (3) yields a circuit that consumes a *stream* of complete database snapshots and outputs a stream of complete view snapshots. Step (4) yields a circuit that consumes a stream of changes to the database and outputs a stream of view changes; however, the internal operation of the circuit is non-incremental, as it rebuilds the complete database using integration operators. Step (5) incrementalizes the circuit by rewriting all operators to compute directly on changes.

6.3.1 Example

In this section we apply the incremental view maintenance algorithm to a concrete query. Let us consider the following query:

```
CREATE VIEW v AS
SELECT DISTINCT t1.x, t2.y FROM (
     SELECT t1.x, t1.id
     FROM t
     WHERE t.a > 2
) t1
JOIN (
     SELECT t2.id, t2.y
     FROM r
     WHERE r.s > 5
) t2 ON t1.id = t2.id
```

Step 1: First we create a DBSP circuit to represent this query using the translation rules from Table 1:

[5] Query planners generally use cost-based heuristics to optimize plans, but IVM planning in general does not have this luxury, since the plan must be generated *before* the data has been fed to the database. Nevertheless, standard query optimization techniques, perhaps based on historical statistics, can be applied to the query plan before generating the incremental plan.

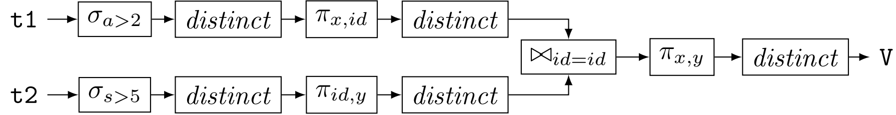

Step 2: we apply the *distinct* optimization rules; first the rule from 6.2 gives us the following equivalent circuit:

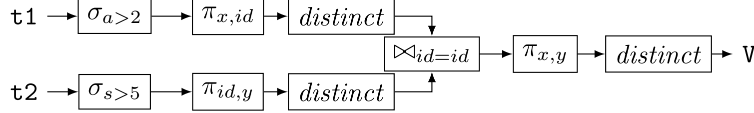

Applying the rule from 6.1 we get:

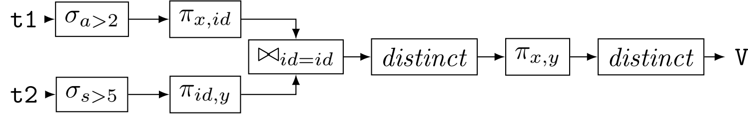

And applying again 6.2 we get:

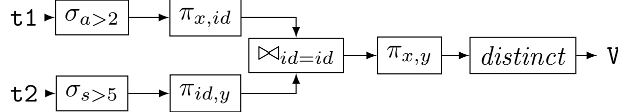

Step 3: we lift the circuit using distributivity of composition over lifting; we obtain a circuit that computes over streams, i.e., for each new input pair of relations `t1` and `t2` it will produce an output view `V`:

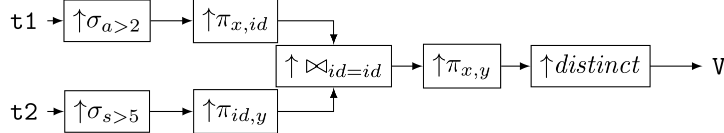

Step 4: incrementalize circuit, obtaining a circuit that computes over changes; this circuit receives changes to relations `t1` and `t2` and for each such change it produces the corresponding change in the output view `V`:

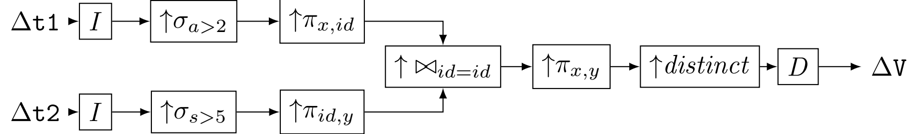

Step 5: apply the chain rule to rewrite the circuit as a composition of incremental operators;

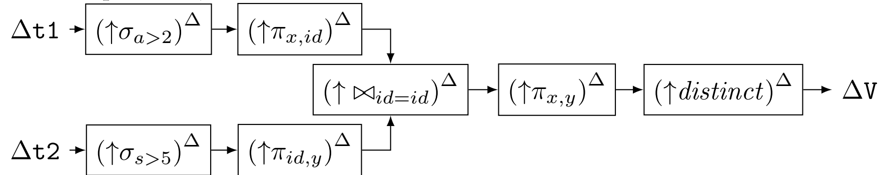

Use the linearity of σ and π to simplify this circuit:

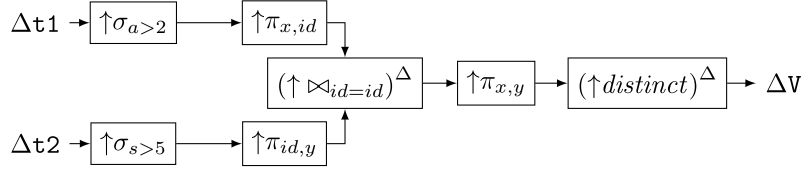

Finally, replace the incremental join using the formula for bilinear operators (Theorem 5.5), and the incremental *distinct* (Proposition 6.3), obtaining the circuit below:

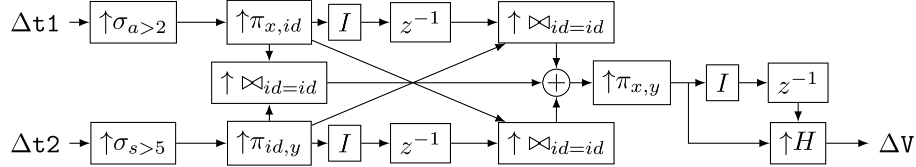

Notice that the resulting circuit contains three integration operations: two from the join, and one from the *distinct*. It also contains three join operators. However, the work performed by each operator for each new input is proportional to the size of change, as we argue in the following section.

6.4 Complexity of incremental circuits

Incremental circuits are efficient. We evaluate the cost of a circuit while processing the t-th input change from two points of view: the work performed, and the total memory used. Even if Q is a pure function, Q^Δ is actually a streaming system, with internal state. This state is stored entirely in the delay operators z^{-1}, some of which appear in I and D operators. The result produced by Q^Δ on the t-th input depends in general not only on the new t-th input, but also on all prior inputs it has received.

We argue that each operator in the incremental version of a circuit is efficient in terms of work and space. We make the standard IVM assumption that the input changes *of each operator* are small[6]: $|\Delta DB[t]| \ll |DB[t]| = |(I(\Delta DB))[t]|$.

An unoptimized incremental operator $Q^\Delta = D \circ Q \circ I$ evaluates query Q on the whole database DB, the integral of the input stream: $DB = I(\Delta DB)$; hence its time complexity is the same as that of the non-incremental evaluation of Q. In addition, each of the I and D operators uses $O(|DB[t]|)$ memory.

Step (5) of the incrementalization algorithm applies the optimizations described in Section 5; these reduce the time complexity of each operator to be a function of $O(|\Delta DB[t]|$. For example, Theorem 5.4, allows evaluating S^Δ, where S is a linear operator, in time $O(|\Delta DB[t]|)$. The I operator can also be evaluated in $O(|\Delta DB[t]|)$ time, because all values that appear in the output of $I(\Delta DB)[t]$ must be present in current input change $\Delta DB[t]$. Similarly, while the *distinct* operator is not linear, $(\uparrow distinct)^\Delta$ can also be evaluated in $O(|\Delta DB[t]|)$ according to Proposition 6.3. Bilinear operators, including join, can be evalu-

[6]In the worst case this may not hold for all operators in a composite query plan because outputs of operators can be large even if inputs are small.

ated in time $O(|DB[t]| \times |\Delta DB[t]|)$, which is a factor of $|DB[t]/\Delta DB[t]|$ better than full re-evaluation.

The space complexity of linear operators is 0 (zero), since they store no data persistently. The space complexity of operators such as $(\uparrow distinct)^\Delta$, $(\uparrow \bowtie)^\Delta$, I, and D is $O(|DB[t]|)$ (the first two because they contain one or more integrals I in their expansion).

7 Nested streams

7.1 Creating and destroying streams

We introduce two new stream operators that are instrumental in expressing recursive query evaluation. These operators allow us to build circuits implementing looping constructs, which are used to iterate computations until a fixed-point is reached.

7.1.1 Stream introduction

Definition 7.1 (Dirac delta)**.** The delta function (named from the Dirac delta function) $\delta_0 : A \rightarrow \mathcal{S}_A$ produces a stream from a scalar value. The output stream is produced as follows from the input scalar:

$$\delta_0(v)[t] \stackrel{\text{def}}{=} \begin{cases} v & \text{if } t = 0 \\ 0_A & \text{otherwise} \end{cases}$$

Here is a diagram showing a δ_0 operator; note that the input is a scalar value, while the output is a stream:

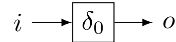

For example, $\delta_0(5)$ is the stream $[\ 5 \quad 0 \quad 0 \quad 0 \quad 0 \quad \cdots\ ]$.

7.1.2 Stream elimination

Recall that $\overline{\mathcal{S}_A}$ was defined in Definition 3.11 to be the set of A-streams over a group A that are zero almost everywhere.

Definition 7.2 (indefinite integral)**.** We define a function $\int : \overline{\mathcal{S}_A} \rightarrow A$ as $\int(s) \stackrel{\text{def}}{=} \sum_{t \geq 0} s[t]$.

$\int$ is closely related to I; if I is the indefinite integral, $\int$ is the definite integral on the interval $0 - \infty$. Unlike I $\int$ produces a scalar value, the "last" distinct value that would appear in the stream produced by I. For example $\int(id|_{\leq 4}) = 0 + 1 + 2 + 3 = 6$, because $I(id|_{\leq 4}) = [\ 0 \quad 1 \quad 3 \quad 6 \quad 6 \quad \cdots\ ]$.

An alternative definition for $\int$ for all streams $\mathcal{S}_A$ would extend the set A with an "infinite": $\overline{A} \stackrel{\text{def}}{=} A \cup \{\top\}$, and define $\int s \stackrel{\text{def}}{=} \top$ for streams that are not zero a.e., $s \in \mathcal{S}_A \setminus \overline{\mathcal{S}_A}$.

Here is a diagrams showing the $\int$ operator; note that the result it produces is a scalar, and not a stream:

$$i \longrightarrow \boxed{\int} \longrightarrow o$$

δ_0 is the left inverse of $\int$, i.e., the following equation holds: $\int \circ \delta_0 = id_A$.

7.1.3 The E and X operators

The composition $I \circ \delta_0$ is frequently used, so we will give it a name, denoting it by $E : A \to \mathcal{S}_A$, $E \stackrel{\text{def}}{=} I \circ \delta_0$.

$$\longrightarrow \boxed{E} \longrightarrow \quad \stackrel{\text{def}}{=} \quad \longrightarrow \boxed{\delta_0} \longrightarrow \boxed{I} \longrightarrow$$

Notice that the output of the E operator is a constant infinite stream, consisting the scalar value at the input. $E(5) = [\ 5 \quad 5 \quad 5 \quad 5 \quad 5 \quad \cdots\ ]$.

Similarly, we denote by $X : \mathcal{S}_A \to A$ the combination $X \stackrel{\text{def}}{=} \int \circ D$.

$$\longrightarrow \boxed{X} \longrightarrow \quad \stackrel{\text{def}}{=} \quad \longrightarrow \boxed{D} \longrightarrow \boxed{\int} \longrightarrow$$

Proposition 7.3. For a monotone stream $o \in \mathcal{S}_A$ we have $X(o) = \lim_{n \to \infty} o[n]$, if the limit exists.

Proof. $X(o|_{\leq n}) = (\int \circ D)(o|_{\leq n}) = o[0] + (o[1] - o[0]) + \ldots + (o[n] - o[n-1]) = o(n)$. The result follows by taking limits on both sides. □

MIHAI: This looks almost right, but it is not.

Clearly, E is the left-inverse of X.

Proposition 7.4. δ_0, $\int$, E, and X are LTI.

Proof. The proof is easy using simple algebraic manipulation of the definitions of these operators. □

7.1.4 Time domains

So far we had the tacit assumption that "time" is common for all streams in a program. For example, when we add two streams, we assume that they use the same "clock" for the time dimension. However, the δ_0 operator creates a streams with a "new", independent time dimension. In Section 12 we will define some well-formed circuit construction rules that will ensure that such time domains are always "insulated", by requiring each diagram that starts with a δ_0 operator to end with a corresponding $\int$ operator:

$$i \longrightarrow \boxed{\delta_0} \longrightarrow \boxed{Q} \longrightarrow \boxed{\int} \longrightarrow o$$

Proposition 7.5. If Q is time-invariant, the circuit above has the zero-preservation property: $\text{zpp}(\int \circ Q \circ \delta_o)$.

Proof. This follows from the fact that all three operators preserve zeros, and thus so does their composition. □

7.2 Streams of streams

7.2.1 Defining nested streams

Since all streams we work with are defined over abelian groups and streams themselves form an abelian group, as pointed in Section 3.4, it follows that we can naturally define streams of streams. $\mathcal{S}_{\mathcal{S}_A} = \mathbb{N} \to (\mathbb{N} \to A)$. This construction can be iterated, but our applications do not require more than two-level nesting. Box-and-arrow diagrams can be used equally to depict functions computing on nested streams; in this case an arrow represent a stream where each value is another stream.

Equivalently, a nested stream in $\mathcal{S}_{\mathcal{S}_A}$ is a value in $\mathbb{N} \times \mathbb{N} \to A$, i.e., a "matrix" with an infinite number of rows, where each row is a stream. For example, we can depict the nested stream $i \in \mathcal{S}_{\mathcal{S}_\mathbb{N}}$ defined by $i[t_0][t_1] = t_0 + 2t_1$ as:

$$i = \begin{bmatrix} [& 0 & 1 & 2 & 3 & \cdots &] \\ [& 2 & 3 & 4 & 5 & \cdots &] \\ [& 4 & 5 & 6 & 7 & \cdots &] \\ [& 6 & 7 & 8 & 9 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

(t_0 is the column index, and t_1 is the row index).

7.2.2 Lifting stream operators

We have originally defined lifting (Section 3.1.1) for scalar functions. We can generalize lifting to apply to stream operators as well. Consider a stream operator $S : \mathcal{S}_A \to \mathcal{S}_B$. We define $\uparrow\!S : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_B}$ as: $(\uparrow\!S(s))[t_0][t_1] \stackrel{\text{def}}{=} S(s[t_0])[t_1], \forall t_0, t_1 \in \mathbb{N}$. Alternatively, we can write $(\uparrow\!S)(s) = S \circ s$.

In particular, a scalar function $f : A \to B$ can be can lifted twice to produce an operator between streams of streams: $\uparrow\uparrow\!f : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_B}$.

Lifting twice a scalar function computes on elements of the matrix pointwise:

$$(\uparrow\uparrow(x \mapsto x \bmod 2))(i) = \begin{bmatrix} [& 0 & 1 & 0 & 1 & \cdots &] \\ [& 0 & 1 & 0 & 1 & \cdots &] \\ [& 0 & 1 & 0 & 1 & \cdots &] \\ [& 0 & 1 & 0 & 1 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

z^{-1} delays the rows of the matrix:

$$z^{-1}(i) = \begin{bmatrix} [& 0 & 0 & 0 & 0 & \cdots &] \\ [& 0 & 1 & 2 & 3 & \cdots &] \\ [& 2 & 3 & 4 & 5 & \cdots &] \\ [& 4 & 5 & 6 & 7 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

while its lifted counterpart delays each column of the matrix:

$$(\uparrow z^{-1})(i) = \begin{bmatrix} [& 0 & 0 & 1 & 2 & \cdots &] \\ [& 0 & 2 & 3 & 4 & \cdots &] \\ [& 0 & 4 & 5 & 6 & \cdots &] \\ [& 0 & 6 & 7 & 8 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

We can also apply both operators, and they commute:

$$(\uparrow z^{-1})(z^{-1}(i)) = z^{-1}((\uparrow z^{-1})(i)) = \begin{bmatrix} [& 0 & 0 & 0 & 0 & \cdots &] \\ [& 0 & 0 & 1 & 2 & \cdots &] \\ [& 0 & 2 & 3 & 4 & \cdots &] \\ [& 0 & 4 & 5 & 6 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

Similarly, we can apply D to nested streams $D : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_A}$, computing on rows of the matrix:

$$D(i) = \begin{bmatrix} [& 0 & 1 & 2 & 3 & \cdots &] \\ [& 2 & 2 & 2 & 2 & \cdots &] \\ [& 2 & 2 & 2 & 2 & \cdots &] \\ [& 2 & 2 & 2 & 2 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

while $\uparrow D : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_A}$ computes on the columns:

$$(\uparrow D)(i) = \begin{bmatrix} [& 0 & 1 & 1 & 1 & \cdots &] \\ [& 2 & 1 & 1 & 1 & \cdots &] \\ [& 4 & 1 & 1 & 1 & \cdots &] \\ [& 6 & 1 & 1 & 1 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

Similarly, we can apply both differentiation operators in sequence:

$$(D(\uparrow D))(i) = \begin{bmatrix} [& 0 & 1 & 1 & 1 & \cdots &] \\ [& 2 & 0 & 0 & 0 & \cdots &] \\ [& 2 & 0 & 0 & 0 & \cdots &] \\ [& 2 & 0 & 0 & 0 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

7.2.3 Strict operators on nested streams

In order to show that operators defined using feedback are well-defined on nested streams we need to extend the notion of strict operators from Section 3.3.

We define a partial order over timestamps: $(i_0, i_1) \leq (t_0, t_1)$ iff $i_0 \leq t_0$ and $i_1 \leq t_1$. We extend the definition of strictness for operators over nested

streams: a stream operator $F : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_B}$ is strict if for any $s, s' \in \mathcal{S}_{\mathcal{S}_A}$ and all times $t, i \in \mathbb{N} \times \mathbb{N}$ we have $\forall i < t, s[i] = s'[i]$ implies $F(s)[t] = F(s')[t]$. Proposition 3.14 holds for this notion of strictness, i.e., the fixed point operator $\text{fix}\, \alpha.F(\alpha)$ is well defined for a strict operator F.

MIHAI: Should write down this proof.

Proposition 7.6. The operator $\uparrow z^{-1} : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_A}$ is strict.

The I operator on $\mathcal{S}_{\mathcal{S}_A}$ is well-defined: it operates on rows of the matrix, treating each row as a single value:

$$I(i) = \begin{bmatrix} [& 0 & 1 & 2 & 3 & \cdots &] \\ [& 2 & 4 & 6 & 8 & \cdots &] \\ [& 6 & 9 & 12 & 15 & \cdots &] \\ [& 12 & 16 & 20 & 24 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

With this extended notion of strictness we have that the lifted integration operator is also well-defined: $\uparrow I : \mathcal{S}_{\mathcal{S}_A} \to \mathcal{S}_{\mathcal{S}_A}$. This operator integrates each column of the stream matrix:

$$(\uparrow I)(i) = \begin{bmatrix} [& 0 & 1 & 3 & 6 & \cdots &] \\ [& 2 & 5 & 9 & 14 & \cdots &] \\ [& 4 & 9 & 15 & 22 & \cdots &] \\ [& 6 & 13 & 21 & 30 & \cdots &] \\ & & & \vdots & & & \end{bmatrix}$$

Notice the following commutativity properties for integration and differentiation on nested streams: $I \circ (\uparrow I) = (\uparrow I) \circ I$ and $D \circ (\uparrow D) = (\uparrow D) \circ D$.

7.2.4 Lifted cycles

Proposition 7.7 (Lifting cycles)**.** For a binary, causal T we have:

$$\uparrow(\lambda s. \text{fix}\, \alpha.T(s, z^{-1}(\alpha))) = \lambda s. \text{fix}\, \alpha.(\uparrow T)(s, (\uparrow z^{-1})(\alpha))$$

i.e., lifting a circuit containing a "cycle" can be accomplished by lifting all operators independently.

Proof. Consider a stream of streams $a = [a_0, a_1, a_2, \cdots] \in \mathcal{S}_{\mathcal{S}_A}$ (where each $a_i \in \mathcal{S}_A$). The statement to prove becomes:

$$\uparrow(\lambda s. \text{fix}\, \alpha.T(s, z^{-1}(\alpha)))(a) = \text{fix}\, \alpha.(\uparrow T)(a, (\uparrow z^{-1})(\alpha))$$

This follows if we show that the value defined as:

$$\beta = \uparrow(\lambda s. \text{fix}\, \alpha.T(s, z^{-1}(\alpha)))(a)$$

satisfies

$$\beta = (\uparrow T)(a, (\uparrow z^{-1})(\beta))$$

Now, expanding the definition of lifting a function:

$$\begin{aligned}\uparrow(\lambda s.\text{fix}\,\alpha.T(s,z^{-1}(\alpha)))(a) &\stackrel{\text{def}}{=} \uparrow(\lambda s.\text{fix}\,\alpha.T(x,z^{-1}(\alpha)))([a_0,a_1,\cdots])\\ &\stackrel{\text{def}}{=} [\text{fix}\,\alpha.T(a_0,z^{-1}(\alpha)),\text{fix}\,\alpha.T(a_1,z^{-1}(\alpha)),\cdots]\\ &= [\alpha_0,\alpha_1,\alpha_2,\cdots]\end{aligned}$$

where, $\forall i.\alpha_i$ is the unique solution of the equation $\alpha_i = T(a_i, z^{-1}(\alpha_i))$. Finally, for $\beta = [\alpha_0, \alpha_1, \cdots]$ we have

$$(\uparrow T)([a_0,a_1,\ldots],(\uparrow z^{-1})(\beta)) = [T(a_0,z^{-1}(\alpha_0)),T(a_1,z^{-1}(\alpha_1)),\ldots]$$

which finishes the proof. □

Proposition 7.7 gives us the tool to lift whole circuits. For example, we have:

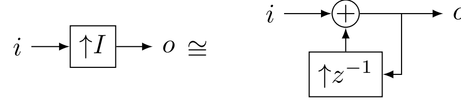

As another example, consider the following circuit $T : A \to B$ that represents a *scalar* function:

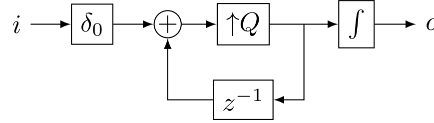

Since this circuit represents a scalar function, it can be lifted like any other scalar function to create a stream computation[7]:

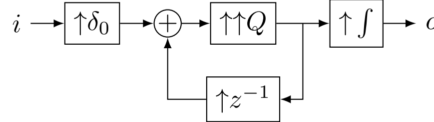

7.3 Fixed-point computations

Theorem 7.8. Given a scalar operator $Q : A \to A$, the output stream o computed by the following diagram (using the lifted version of Q) is given by $\forall t \in \mathbb{N}.o[t] = Q^{t+1}(i)$, where k is the composition of Q with itself k times.

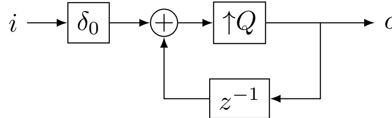

[7]Notice that + is not shown lifted in this circuit, since there is in fact a + operator for any type, and we have that $\uparrow(+_A) = +_{\mathcal{S}_A}$

Proof. Let us compute $o[t]$. We have that $o[0] = Q(\delta_0(i)[0] + 0) = Q(i)$. $o[1] = Q(o[0] + \delta_0(i)[1]) = Q(Q(i) + 0) = Q^2(i)$. We can prove by induction over t that $o[t] = Q^{t+1}(i)$. □

Observation: in this circuit the "plus" operator behaves as an `if`, selecting between the base case and the inductive case in a recursive definition. This is because the left input contains a single non-zero element (i), in the first position, while the bottom input starts with a 0.

8 Recursive queries in DBSP

Recursive queries are very useful in a many applications. For example, many graph algorithms (such as graph reachability or transitive closure) are naturally expressed using recursive queries.

We illustrate the implementation of recursive queries in DBSP for stratified Datalog.

8.1 Implementing recursive queries

For warm-up we start with a single recursive queries, and then we discuss the case of mutually recursive queries.

8.1.1 Recursive rules

In Datalog a recursive rule appears when a relation that appears in the head of a rule is also used in a positive term in the rule's body. (Stratification disallows the use of the same relation negated in the rule's body). The Datalog semantics of recursive rules is to compute a fixedpoint.

Consider a Datalog program of the form:

```
O(v) :- C(v).  // base case
O(v) :- ..., O(x), I(z), ... .  // recursive case
```

Note that relation `O` is recursively defined. Let us assume wlog that the `O` relation depends on two other relations (i.e., in the rule bodies defining `O` the two other relations appear) : a "base case" relation `C` (which appears in a non-recursive rule), and a relation `I` which appears in the recursive rule, but does not itself depend on `O`.

To implement the computation of `O` as a circuit we perform the following algorithm:

Algorithm 8.1 (recursive queries)**.**

1. Implement the non-recursive relational query R as described in Section 4 and Table 1; this produces an acyclic circuit whose inputs and outputs are a $\mathbb{Z}$-sets:

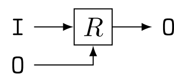

2. Lift this circuit to operate on streams:

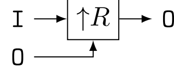

We construct $\uparrow R$ by lifting each operator of the circuit individually according to Proposition 3.3.

3. Build a cycle, connecting the output to the corresponding input via a delay:

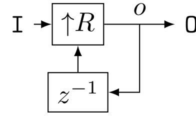

4. "Bracket" the circuit in I and D nodes, and then in δ_0 and $\int$:

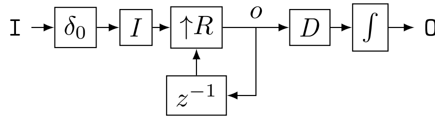

This circuit as drawn is not a well-formed circuit. It can, however, be modified into an equivalent well-formed circuit by adding two constant zero value streams:

VAL: Why not well-formed? As far as I can see its semantics follows from Corollary 3.17.

MIHAI: The WFC rules require any circuit bracketed by $\delta_0 - \int$ to have no other input or output edges. It also requires back-edges to go through a plus only. It is more strict than the stream computation rules.

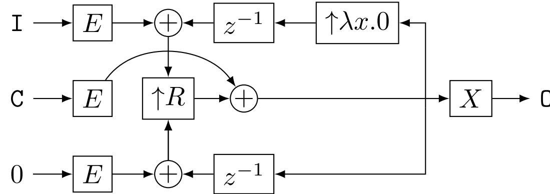

Theorem 8.2. If isset(I) and isset(C), the output of the circuit above is the relation O as defined by the Datalog semantics of recursive relations as a function of the input relations I and C.

Proof. The proof is by structural induction on the structure of the circuit. As a basis for induction we assume that the circuit R correctly implements the semantics of the recursive rule body when treating O as an independent input. We need to prove that the output of circuit encompassing R produces the correct value of the O relation, as defined by the recursive Datalog equation.

Let us compute the contents of the o stream, produced at the output of the *distinct* operator. We will show that this stream is composed of increasing approximations of the value of O, and in fact $\texttt{O} = \lim_{t \to \infty} o[t]$ if the limit exists.

We define the following one-argument function: $R'(x) = \lambda x.R(\texttt{I}, x)$. Notice that the left input of the $\uparrow R$ block is a constant stream with the value I. Due to the stratified nature of the language, we must have ispositive(R'), so $\forall x.R'(x) \geq x$. Also $\uparrow R'$ is time-invariant, so $R'(0) = 0$.

With this notation for R' the previous circuit has the output as the following simpler circuit:

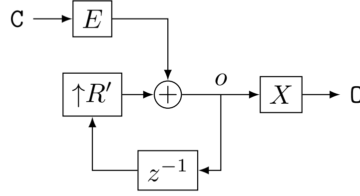

We use the following notation: $x \cup y = \mathit{distinct}(x + y)$. As discussed in Section 14.2.3 the $\cup$ operation computes the same result as set union when x and y are sets. With this notation let us compute the values of the o stream:

$$\begin{aligned} o[0] =& \mathtt{C} + R'(0) = \mathtt{C} \cup R'(0) = \mathtt{C} \\ o[1] =& \mathtt{C} + R'(o[0]) = \mathtt{C} \cup R'(\mathtt{C}) \\ o[t] =& \mathtt{C} + R(o[t-1]) = \mathtt{C} \cup R'(o[t-1]) \end{aligned}$$

Defining a new helper function $S(x) = \mathtt{C} \cup R'(x)$, the previous system of equations becomes:

$$\begin{aligned} o[0] =& S(0) \\ o[1] =& S(S(0)) \\ o[t] =& S(o[t-1]) \end{aligned}$$

So, by induction $o[t] = S^t(0)$, where by S^t we mean $\underbrace{S \circ S \circ \ldots \circ S}_{t}$. S is monotone because R' is monotone; thus, if there is a time k such that $S^k(0) = S^{k+1}(0)$, we have $\forall j \in \mathbb{N}. S^{k+j}(0) = S^k(0)$.

$\mathtt{O}$ is computed by the X operator as the limit of stream o: $\mathtt{O} = X(o) = \lim_{n \to \infty} o[n]$. If this limit exists (i.e., a fixed-point is reached), the circuit computes the fixed point $\text{fix}\, x.S(x)$. This is exactly the definition of the Datalog semantics of a recursive relation definition: $\mathtt{O} = \text{fix}\, x.\mathtt{C} \cup R(\mathtt{I}, x)$. $\square$

Note that the use of unbounded domains (like integers with arithmetic operations) does not guarantee convergence for all programs.

Our circuit implementation is in fact computing the value of relation $\mathtt{O}$ using the standard **naïve evaluation** algorithm (e.g., see Algorithm 1 from [22]).

Observe that the "inner" part of the circuit is the incremental form of another circuit, since is "sandwiched" between I and D operators. According to Proposition 5.2, part 7, the circuit can be rewritten as:

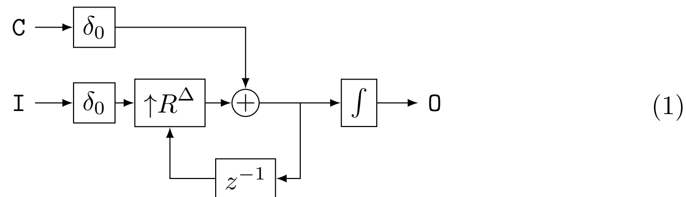

This form of the circuit is effectively implementing the **semi-naïve evaluation** of the same relation (Algorithm 2 from [22]). So the correctness of semi-naïve evaluation is an immediate consequence of the cycle rule from Proposition 5.2.

8.2 Example: a recursive query in DBSP

Here we apply the algorithm 8.1, that converts a recursive query into a DBSP circuit, to a concrete Datalog program. The program computes the transitive closure of a directed graph:

```
// Edge relation with head and tail
input relation E(h: Node, t: Node)
// Reach relation with source s and sink t
output relation R(s: Node, t: Node)
R(x, y) :- E(x, y).
R(x, y) :- E(x, z), R(z, y).
```

Step 1: we ignore the fact that R is both an input and an output and we implement the DBSP circuit corresponding to the body of the query. This query could be expressed in SQL as:

```
( SELECT * FROM E)
UNION
( SELECT E.h , R.t
  FROM E JOIN R
  ON E.t = R.s )
```

This query generates a DBSP circuit with inputs `E` and `R`:

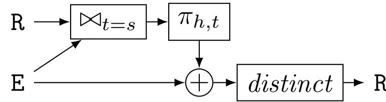

Step 2: Lift the circuit by lifting each operator pointwise:

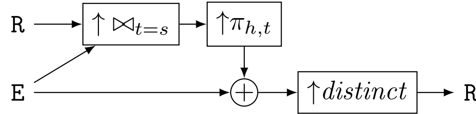

Step 3: Connect the feedback loop implied by relation `R`:

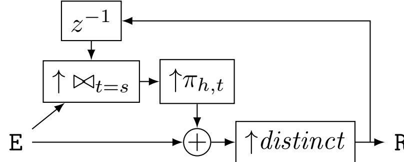

Step 4: "bracket" the circuit, once with I-D, then with δ_0-$\int$:

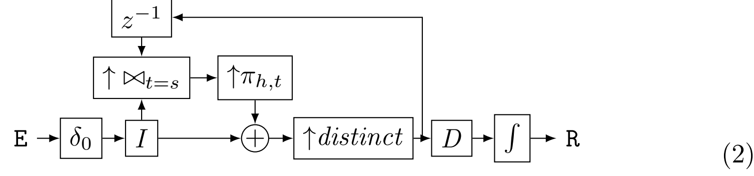

(2)

The above circuit is a complete implementation of the non-streaming recursive query; given an input relation E it will produce its transitive closure R at the output.

Now we use seminaïve evaluation 1 to rewrite the circuit (to save space in the figures we will omit the indices from π and σ in the subsequent figures):

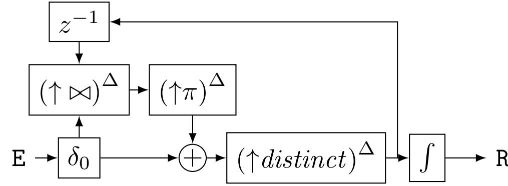

Using the linearity of $\uparrow\pi$, this can be rewritten as:

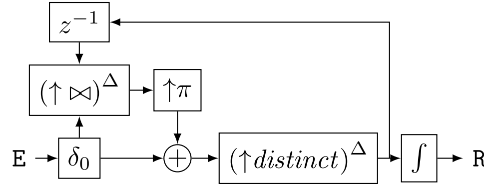

8.2.1 Mutually recursive rules

Given a stratified Datalog program we can compute a graph where relations are nodes and dependences between relations are edges. We then compute the strongly connected components of this graph. All relations from a strongly-connected component are mutually recursive.

Let us consider the implementation of a single strongly-connected component defining n relations $O_i, i \in [n]$. We can assume wlog that the definition of O_i has the following structure:

$$\begin{array}{lll} O_i(v) & \texttt{:-} & C_i(v)\texttt{.} \\ O_i(v) & \texttt{:-} & I_i(x), O_0(v_0), O_1(v_1), \ldots, O_{n-1}(v_{n-1}), v = \ldots. \end{array}$$

There are exactly n base cases, one defining each O_i. Also, we assume that each O_i relation depends on an external relation I_i, which does not itself depend on any O_k.

To compile these into circuits we generalize the algorithm from Section 8.1.1:

1. For each recursive rule for O_i implement a circuit R_i that treats all $O_j.j \in [n]$ and I_i in the rule body as inputs. Here is the circuit R_0 for relation O_0:

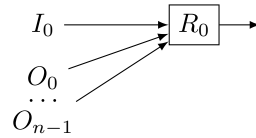

2. Embed each circuit R_i as part of a "widget" as follows:

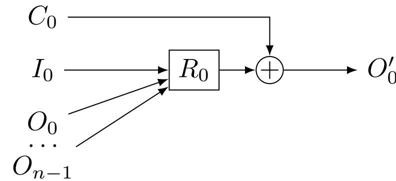

3. Finally, lift each such widget and connect them to each other via z^{-1} operators to the corresponding recursive inputs. The following is the shape of the circuit computing O_0; the O'_j sources correspond to the widget outputs of the other recursive circuits:

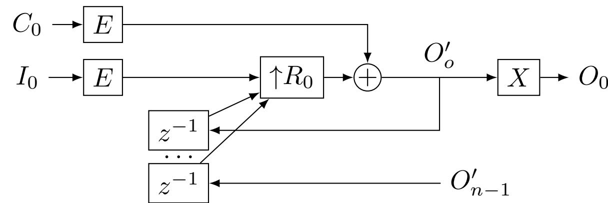

Theorem 8.3. The program defined by the previous circuit computes the relations O_i as a function of the input relations I_j and C_i.

Proof. TODO. □

Example: mutually recursive relations Consider the Datalog program below computing the transitive closure of a graph having two kinds of edges, blue (B) and red (R):

```
P(x,y) :- B(x,y).
Q(x,y) :- R(x,y).
P(x,y) :- B(x,z), Q(z,y).
Q(z,y) :- R(x,z), P(z,y).
O(x,y) :- P(x,y).
O(x,y) :- Q(x,y).
```

The program defined by the following circuit computes the relation `O` as a function of the input relations `R`, `B`:

VAL: I will add in Section 3.3 a consequence of Corollary 3.17 to justify the well-definedness of this.

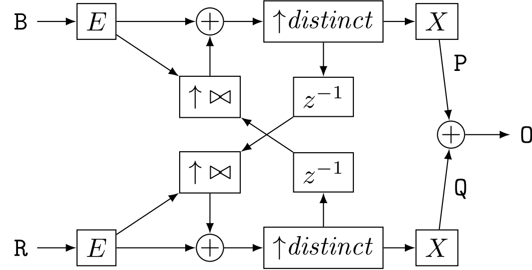

9 Incremental recursive queries

In Section 3–4 we showed how to incrementalize a relational query by compiling it into a circuit, lifting the circuit to compute on streams, and applying the $\cdot^{\Delta}$ operator to the lifted circuit. In Section 14 we showed how to compile a recursive query into a circuit that employs incremental computation internally to compute the fixed point. Here we combine these results to construct a circuit that evaluates a *recursive query incrementally*. The circuit receives a stream of updates to input relations, and for every update recomputes the fixed point. To do this incrementally, it preserves the stream of changes to recursive relations produced by the iterative fixed point computation, and adjusts this stream to account for the modified inputs. Thus, every element of the input stream yields a stream of adjustments to the fixed point computation, using *nested streams*.

This proposition gives the ability to lift entire circuits, including circuits computing on streams and having feedback edges, which are well-defined, due to Proposition 7.6. With this machinery we can now apply Algorithm 6.4 to arbitrary circuits, even circuits built for recursively-defined relations. Consider the "semi-naive" circuit from Section 8: and denote $distinct \circ R$ with T:

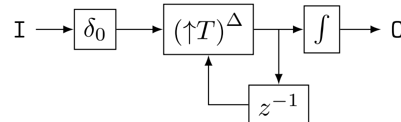

Lift the entire circuit using Proposition 7.7 and incrementalize it:

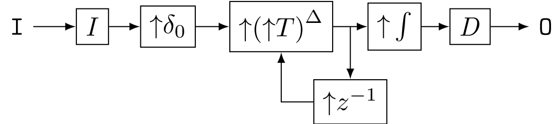

Now apply the chain rule to this circuit, and use the linearity of δ_0 and $\int$:

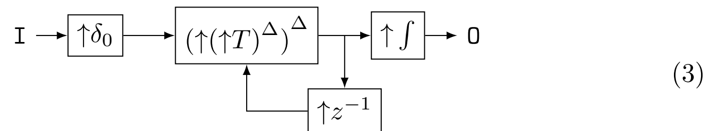

(3)

This is the incremental version of an arbitrary recursive query.

9.1 Incrementalizing a recursive query

Here we take the DBSP circuit for the transitive closure of a graph generated in Section 8.2 and convert it to an incremental circuit using algorithm 6.4. The resulting circuit maintains the transitive closure as edges are inserted or removed.

First we lift the circuit entirely, using Proposition 7.7:

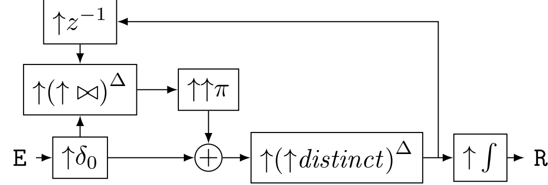

We convert this circuit into an incremental circuit, which receives in each transaction the changes to relation `E` and produces the corresponding changes to relation `R`:

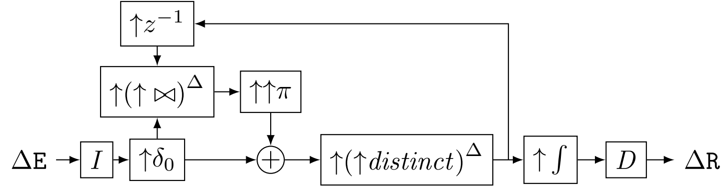

We can now apply again the chain and cycle rules to this circuit:

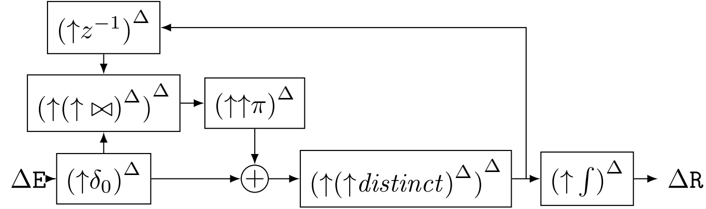

We now take advantage of the linearity of $\uparrow\delta_0$, $\uparrow\int$, $\uparrow z^{-1}$, and $\uparrow\uparrow\pi$ to simplify the circuit by removing some $\cdot^\Delta$ invocations:

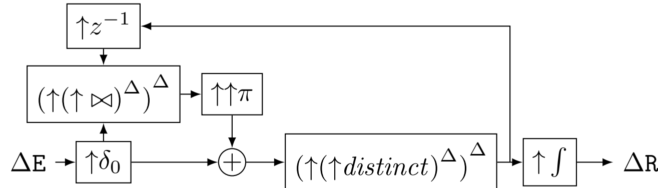

There are two applications of $\cdot^\Delta$ remaining in this circuit: $(\uparrow(\uparrow \bowtie)^\Delta)^\Delta$ and $(\uparrow(\uparrow distinct)^\Delta)^\Delta$. We expand their implementations separately, and we stitch

them into the global circuit at the end. This ability to reason about sub-circuits highlights the modularity of DBSP.

The join is expanded twice, using the bilinearity of $\uparrow\bowtie$ and $\uparrow\uparrow\bowtie$. Let's start with the inner circuit, implementing $(\uparrow\bowtie)^\Delta$, given by Theorem 5.5:

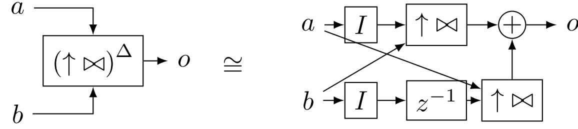

Now we lift and incrementalize to get the circuit for $(\uparrow(\uparrow\bowtie)^\Delta)^\Delta$:

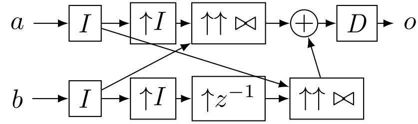

Applying the chain rule and the linearity of $\uparrow I$ and $\uparrow z^{-1}$:

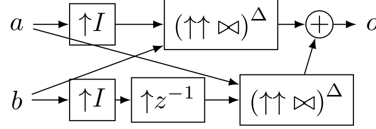

We now have two applications of $(\uparrow\uparrow\bowtie)^\Delta$. Each of these is the incremental form of a bilinear operator, so it in the end we will have $2 \times 2 = 4$ applications of $\uparrow\uparrow\bowtie$. Here is the final form of the expanded join circuit:

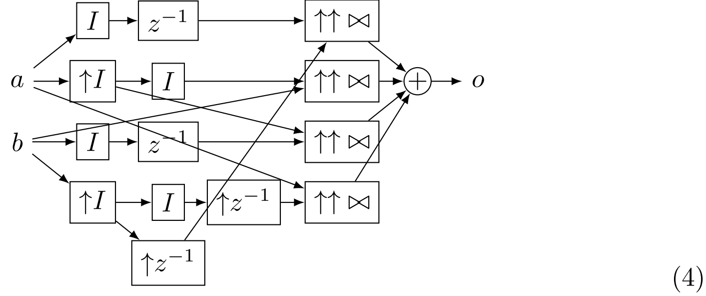

(4)

Returning to $(\uparrow(\uparrow distinct)^\Delta)^\Delta$, we can compute its circuit by expanding once using Proposition 6.3:

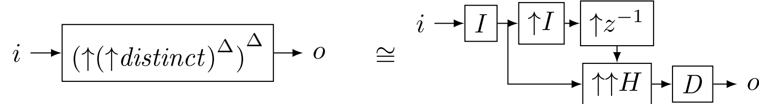

Finally, stitching all these pieces together we get the final circuit shown in Figure 1.

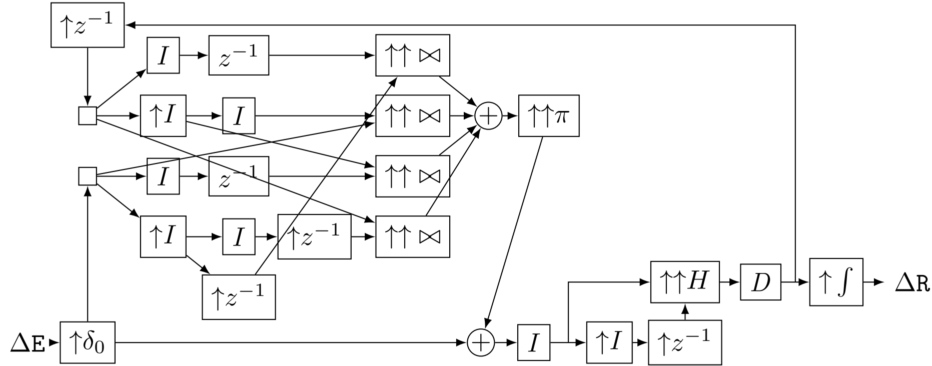

Figure 1: Final form of circuit from Section 8.2 which is incrementally maintaining the transitive closure of a graph.

10 Complexity of recursive incremental circuits

Time complexity The time complexity of an incremental recursive query can be estimated as a product of the number of fixed point iterations and the complexity of each iteration. The incrementalized circuit (3) never performs more iterations than the non-incremental circuit (1): once the non-incremental circuit reaches the fixed point, its output is constant, and the derivative of corresponding value in the incrementalized circuit becomes 0.

Moreover, the work performed by each operator in the incremental circuit is asymptotically better than the non-incremental one. As a concrete example, consider a join in a recursive circuit. A non-incremental implementation is shown in the Appendix in example 2. The incremental implementation of the same circuit is in circuit 5, and contains 4 join operators. The work performed by the non-incremental join is $O(|DB|^2)$ for each iteration. The size of the inputs of each of the joins in the incremental circuit is shown below. We notice that the four join operators perform work $O(|\Delta DB|^2)$, $O(|DB||\Delta DB|)$, $(O|DB||\Delta DB|)$, and $O(|DB|0)$ respectively (the last operator performs work only in the first iteration), so each of them is asymptotically better than the non-incremental version.

In this diagram we annotate edges with the size of the collections flowing along the edge. The a stream is produced by a δ_0 operator; it contains initially a change to the database, but afterwards all elements are 0 – we show this with $|\Delta DB|, 0, 0$.

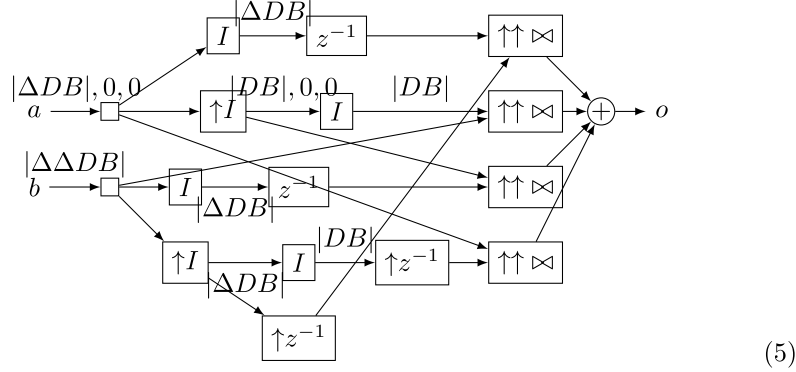

(5)

Space complexity Integration (I) and differentiation (D) of a stream $\Delta s \in \mathcal{S}_{\mathcal{S}_A}$ use memory proportional to $\sum_{t_2} \sum_{t_1} |s[t_1][t_2]|$, i.e., the total size of changes aggregated over columns of the matrix. The unoptimized circuit integrates and differentiates respectively inputs and outputs of the recursive program fragment. As we move I and D inside the circuit using the chain rule, we additionally store changes to intermediate streams. Effectively we cache results of fixed point iterations from earlier timestamps to update them efficiently as new input changes arrive. Notice that space usage is proportional to the number of iterations of the inner loop that computes the fixed-point. Fortunately, many recursive algorithms converge in a relatively small number of steps (for example, transitive closure requires a number of steps that is the log of the diameter of the graph).

11 Additional query languages

In this section we describe several query models that go behind stratified Datalog and show how they can be implemented in DBSP.

11.1 Aggregation

Aggregation in SQL applies a function a to a whole set producing a "scalar" result with some type R: $a : 2^A \to R$. We convert such aggregation functions to operate on $\mathbb{Z}$-sets, so in DBSP an aggregation function has a signature $a : \mathbb{Z}[A] \to R$.

The SQL `COUNT` aggregation function is implemented on $\mathbb{Z}$-sets by $a_{\texttt{COUNT}} : \mathbb{Z}[A] \to \mathbb{Z}$, which computes a *sum* of all the element weights: $a_{\texttt{COUNT}}(s) = \sum_{x \in s} s[x]$. The SQL `SUM` aggregation function is implemented on $\mathbb{Z}$-sets by $a_{\texttt{SUM}} : \mathbb{Z}[\mathbb{R}] \to \mathbb{R}$ which performs a *weighted sum* of all (real) values: $a_{\texttt{SUM}}(s) = \sum_{x \in s} x \times s[x]$.

With this definition the aggregation functions $a_{\texttt{COUNT}}$ and $a_{\texttt{SUM}}$ are in fact linear transformations between the group $\mathbb{Z}[A]$ and the result group ($\mathbb{Z}$, and $\mathbb{R}$ respectively).

If the output of the DBSP circuit can be such a "scalar" value, then aggregation with a linear function is simply function application, and thus it is automatically incremental. However, in general, for composing multiple queries we require the result of an aggregation to be a singleton $\mathbb{Z}$-set (containing a single value), and not a scalar value. In this case the aggregation function is implemented in DBSP as the composition of the actual aggregation and the makeset : $A \to \mathbb{Z}[A]$ function, which converts a scalar value of type A to a singleton $\mathbb{Z}$-set, defined as follows: $\text{makeset}(x) \stackrel{\text{def}}{=} 1 \cdot x$.

In conclusion, the following SQL query: `SELECT SUM(c) FROM I` is implemented as the following circuit:

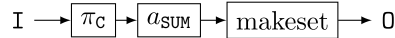

The lifted incremental version of this circuit is interesting: since π and $a_{\texttt{SUM}}$ are linear, they are equivalent to their own incremental versions. Although $(\uparrow\text{makeset})^{\Delta} = D \circ \uparrow\text{makeset} \circ I$ cannot be simplified, it is nevertheless efficient, doing only O(1) work per invocation, since its input and output are singleton values.

An aggregation function such as `AVG` can be written as the composition of a more complex linear function that computes a pair of values using `SUM` and `COUNT`, followed by a makeset and a selection operation that divides the two columns.

```
SELECT AVG(c) FROM I
```

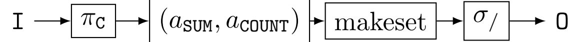

Finally, some aggregate functions, such as `MIN`, are *not* incremental in general, since for handling deletions they may need to know the full set, and not just its changes. The lifted incremental version of such aggregate functions is implemented essentially by "brute force", using the formula $(\uparrow a_{\texttt{MIN}})^{\Delta} = D \circ \uparrow a_{\texttt{MIN}} \circ I$. Such functions perform work proportional to $R(s)$ at each invocation.

Note that the SQL `ORDER BY` directive can be modeled as a non-linear aggregate function that emits a list. However, such an implementation it is not efficiently incrementalizable in DBSP. We leave the efficient handling of ORDER BY to future work.

Even when aggregation results do not form a group, they usually form a structure with a zero element. We expect that a well-defined aggregation function maps empty $\mathbb{Z}$-sets to zeros in the target domain.

11.2 Nested relations

11.2.1 Indexed partitions

Let $A[K]$ be the set of functions with finite support from K to A. Consider a group A, an arbitrary set of **key values** K, and a *partitioning function* $k : A \to A[K]$ with the property that $\forall a \in A . a = \sum k(a)$. We call elements of $A[K]$ *indexed* values of A — indexed by a key value.

Notice that $A[K]$ also has a group structure, and k itself is a linear function (homomorphism). As an example, if $A = \mathbb{Z}[B_0 \times B_1]$, we can use for k the first projection $k : A \to \mathbb{Z}[A][B_0]$, where $k(a)[b] = \sum_{t \in a, t|0=b} a[t] \cdot t$. In other words, k projects the elements in $\mathbb{Z}[B_0 \times B_1]$ on their first component. This enables *incremental computations on nested relations*. This is how operators such as group-by are implemented: the result of group-by is an indexed $\mathbb{Z}$-set, where each element is indexed by the key of the group it belongs to. Since indexing is linear, its incremental version is very efficient. Notice that the structure $\mathbb{Z}[A][K]$ represents a form of *nested relation*.

11.3 Grouping; indexed relations

Pick an arbitrary set K of "key values." Consider the mathematical structure of finite maps from K to $\mathbb{Z}$-sets over some other domain A: $K \to \mathbb{Z}[A] = \mathbb{Z}[A][K]$. We call values i of this structure **indexed $\mathbb{Z}$-sets**: for each key $k \in K$, $i[k]$ is a $\mathbb{Z}$-set. Because the codomain $\mathbb{Z}[A]$ is an abelian group, this structure is itself an abelian group.

We use this structure to model the SQL `GROUP BY` operator in DBSP. Consider a **partitioning function** $p : A \to K$ that assigns a key to any value in A. We define the grouping function $G_p : \mathbb{Z}[A] \to (K \to \mathbb{Z}[A])$ as $G_p(a)[k] \stackrel{\text{def}}{=} \sum_{x \in a. p(x)=k} a[x] \cdot x$. When applied to a $\mathbb{Z}$-set a this function returns a indexed $\mathbb{Z}$-set, where each element is called a **grouping**[8]: for each key k a grouping is a $\mathbb{Z}$-set containing all elements of a that map to k (as in SQL, groupings are multisets, represented by $\mathbb{Z}$-sets). Consider our example $\mathbb{Z}$-set R from Section 4, and a key function $p(s)$ that returns the first letter of the string s. Then we have that $G_p(R) = \{\texttt{j} \mapsto \{\texttt{joe} \mapsto 1\}, \texttt{a} \mapsto \{\texttt{anne} \mapsto -1\}\}$, i.e., grouping with this key function produces an indexed $\mathbb{Z}$-set with two groupings, each of which contains a $\mathbb{Z}$-set with one element.

The grouping function G_p is linear for any p. It follows that the group-by implementation in DBSP is automatically incremental: given some changes to the input relation we can apply the partitioning function to each change separately to compute how each grouping changes.

11.4 `GROUP BY-AGGREGATE`

Grouping in SQL is almost always followed by aggregation. Let us consider an aggregation function $a : (K \times \mathbb{Z}[A]) \to B$ that produces values in some group B, and an indexed relation of type $\mathbb{Z}[A][K]$, as defined above in Section 14.2.8. The nested relation aggregation operator $Agg_a : \mathbb{Z}[A][K] \to B$ applies a to the contents of each grouping independently and adds the results: $Agg_a(g) \stackrel{\text{def}}{=} \sum_{k \in K} a(k, g[k])$. To apply this to our example, let us compute the equivalent of GROUP-BY count; we use the following aggregation function $count : K \times \mathbb{Z}[A]$, $count(k, s) = \text{makeset}((k, a_{\texttt{COUNT}}(s)))$, using the $\mathbb{Z}$-set counting function $a_{\texttt{COUNT}}$

[8] We use "group" for the algebraic structure and "grouping" for the result of `GROUP BY`.

from Section 14.2.9; the notation (a, b) is a pair of values a and b. Then we have $Agg_{count}(G_p(R)) = \{(\texttt{j}, 1) \mapsto 1, (\texttt{a}, -1) \mapsto 1\}$.

Notice that, unlike SQL, DBSP can express naturally computations on indexed $\mathbb{Z}$-sets, they are just an instance of a group structure. One can even implement queries that operate on each grouping in an indexed $\mathbb{Z}$-set. However, our definition of incremental computation is only concerned with incrementality in the *outermost* structures. We leave it to future work to explore an appropriate definition of incremental computation that operates on the *inner* relations.

A very useful operation on nested relations is **flatmap**, which is essentially the inverse of partitioning, converting an indexed $\mathbb{Z}$-set into a $\mathbb{Z}$-set: $\text{flatmap} : \mathbb{Z}[A][K] \to \mathbb{Z}[A \times K]$. flatmap is in fact a particular instance of aggregation, using the aggregation function $a : K \times \mathbb{Z}[A] \to \mathbb{Z}[A \times K]$ defined by $a(k, s) = \sum_{x \in s[k]} s[k][x] \cdot (k, x)$. For our previous example, $\text{flatmap}(G_p(R)) = \{(\texttt{j}, \texttt{joe}) \mapsto 1, (\texttt{a}, \texttt{anne}) \mapsto -1\}$.

If we use an aggregation function $a : K \times Z[A]$ that is linear in its second argument, then the aggregation operator Agg_a is linear, and thus fully incremental. As a consequence, flatmap is linear. However, many practical aggregation functions for nested relations are in fact not linear; an example is the *count* function above, which is not linear since it uses the makeset non-linear function. Nevertheless, while the incremental evaluation of such functions is not fully incremental, it is at least partly incremental: when applying a change to groupings, the aggregation function only needs to be re-evaluated *for groupings that have changed*.

11.5 Streaming joins

Consider a binary query $T(s, t) = I(t) \ \uparrow\bowtie \ s$. This is the *relation-to-stream join* operator supported by streaming databases like ksqlDB [32]. Stream t carries changes to a relation, while s carries arbitrary data, e.g., logs or telemetry data points. T "discards" values from s after matching them against the accumulated contents of the relation.

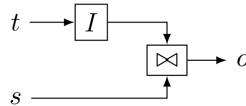

11.6 Explicit delay

So far the z^{-1} operator was confined to its implicit use in integration or differentiation. However, it can be exposed as a primitive operation that can be applied to streams or collections. This enables programs that can perform time-based window computations over streams, and convolution-like operators.

11.7 Multisets/bags

Since $\mathbb{Z}$-sets generalize multisets and bags, it is easy to implement query operators that compute on such structures. For example, while SQL `UNION` is $\mathbb{Z}$-set addition followed by *distinct*, `UNION ALL` is just $\mathbb{Z}$-set addition.

11.8 Window aggregates

Streaming databases often organize the contents of streams into windows, which store a subset of data points with a predefined range of timestamps. The circuit below (a convolution filter in DSP) computes a *fixed-size sliding-window aggregate* over the last four timestamps defined by the T_i functions.

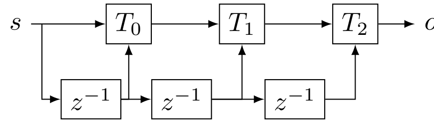

In practice, windowing is usually based on physical timestamps attached to stream values rather than logical time. For instance, the CQL [9] query "`SELECT * FROM events [RANGE 1 hour]`" returns all events received within the last hour. The corresponding circuit (on the left) takes input stream $s \in \mathcal{S}_{\mathbb{Z}[A]}$ and an additional input $\theta \in \mathcal{S}_{\mathbb{R}}$ that carries the value of the current time.

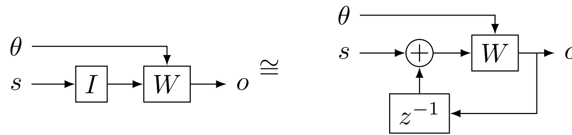

where the *window operator* W prunes input $\mathbb{Z}$-sets, only keeping values with timestamps less than an hour behind $\theta[t]$. Assuming $ts : A \to \mathbb{R}$ returns the physical timestamp of a value, W is defined as $W(v, \theta)[t] \stackrel{\text{def}}{=} \{x \in v[t].ts(x) \geq \theta[t] - 1hr\}$. Assuming θ increases monotonically, W can be moved inside integration, resulting in the circuit on the right, which uses bounded memory to compute a window of an unbounded stream. This circuit is a building block of a large family of window queries, including window joins and aggregation. We conjecture that DBSP can express any CQL query.

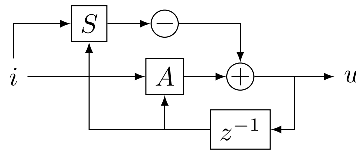

11.9 Relational while queries

(See also non-monotonic semantics for Datalog$^{\neg}$ and Datalog$^{\neg\neg}$[5].) To illustrate the power of DBSP we implement the following "while" program, where Q is an arbitrary relational algebra query:

```
x := i;
while (x changes)
    x := Q(x);
```

The DBSP implementation of this program is:

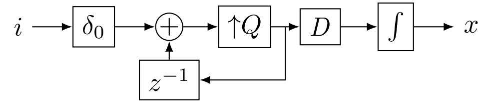

This circuit can be converted to a streaming circuit that computes a stream of values i by lifting it; it can be incrementalized using Algorithm 6.4 to compute on changes of i:

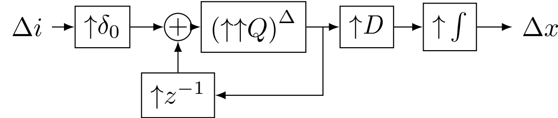

Part II
Implementation

12 Well-formed circuits

In this section we formalize the shape legal computations allowed in our framework. We model computations as circuits. A circuit is a directed graph where each vertex is a primitive computation node (a function) and each edge has a type; at circuit evaluation time each edge will represent one value of that type.

We provide a recursive set of circuit construction rules. We call a circuit **well-formed** (**WFC**) if it can be constructed by a sequence of applications of these rules. For each WFC construction rule of a circuit C from simpler parts we also provide typing derivation rules and a denotational semantics for $[\![C]\!]$ that reduces the meaning of C to its components. The semantics of a circuit expresses each circuit output as a function of the circuit inputs.

12.1 Primitive nodes

We assume a set of base types that represent abelian groups: $A_1, A_2, \ldots,$ $B_1, B_2, \ldots$. (The base types do *not* include stream types; stream types are derived.)

We are given a fixed set of **primitive computation nodes P**; each primitive node has an input arity k. Iin our applications we only use unary and binary primitive nodes, so we will restrict ourselves to such nodes, but these constructions can be generalized to nodes with any arity. A binary node $n \in \mathbf{P}$ has a type of the form $n : A_0 \times A_1 \to A$, where all As are base types. We assume

the existence of a total function "meaning" $[\![\cdot]\!]$ that gives the semantics of each node: $[\![n]\!] : [\![A_0]\!] \times [\![A_1]\!] \to [\![A]\!]$.

In Section 3.1 we have seen many generic primitive nodes: the identity function $id : A \to A$, scalar function nodes (where all inputs are base types A_i), sum $\bigoplus : A \times A \to A$, negation $- : A \to A$, pairs $\langle \cdot, \cdot \rangle : A \times B \to \langle A, B \rangle$, $\text{fst} : A \times B \to A$ and $\text{snd} : A \times B \to B$. Their typing and semantics is standard. We will introduce more domain-specific primitive nodes when discussing specific applications, in Section 14.

12.2 Circuits as graphs

A circuit is a 5-tuple $C = (I, O, V, E, M)$.

- I is an ordered list of input ports. An input port indicates a value produced by the environment. We use letters like i, j for input ports. Each input port has a type $i : A$ for some abelian group A (where A can be a stream type).
- O is an ordered list of output ports. An output port represents a value produced by the circuit. as a function of the values of the input ports. We use letters like o, l for output ports. Each output port has a type $o : A$ for some abelian group A (where A can be a stream type).
- V is a set of internal vertices.
- E is a set of edges; $E \subseteq (V \times V) \cup (I \times V) \cup (V \times O)$. We call an edge of the form (i, v) for $i \in I, v \in V$ an **input edge**, and an edge of the form (v, o) for $v \in V, o \in O$ an **output edge**.
- Each internal vertex v is associated with a primitive computation node or with another circuit; for primitive nodes the **implementation function** $M : V \to \mathbf{P}$ gives the primitive computation associated with a vertex.

Given a circuit C we use the suffix notation to indicate its various components, e.g., $C.O$ is the list of output edges, and $C.V$ is the set of vertexes.

12.3 Circuit semantics

Given a circuit C with k input ports $C.I = (i_j, j \in [k])$ with types $i_j : A_j, j \in [k]$, and m output ports $C.O = (o_j, j \in [m])$ with types $o_j : B_j, j \in [m]$, the semantics of the circuit is given by the semantics of all its output ports: $[\![C]\!] = \prod_{j \in [m]} [\![C.o_j]\!]$. The semantics of output port $C.o_j$ is a total function $[\![C.o_j]\!] : [\![A_0]\!] \times \ldots [\![A_{k-1}]\!] \to [\![B_j]\!]$.

12.4 Circuit construction rules

We give a set of inductive rules for constructing WFCs. In the inductive definitions we always combine or WFCs to produce a new WFC.

These rules maintain the following invariants about all constructed circuits, which can proved by structural induction:

- All input and outputs are either all scalars or they are all streams of the same "depth".
- All circuits are time-invariant.
- For circuits operating over streams, all circuit outputs are causal in all of the circuit inputs.

12.4.1 Single node

Given a primitive (unary or binary) node n with type $n : A_0 \times A_1 \to B$, (where each A_i, B is a base type), we can construct a circuit $C(n)$ with a single vertex. The circuit has exactly 2 input ports and any number m of output ports having all the same type B; the output ports will carry all the same value.

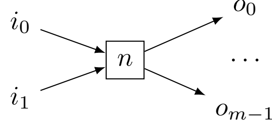

Formally:

- $C(n).V = \{v\}$.
- $C(n).I = (i_j, j \in [k])$
- $C(n).O = (o_j, j \in [m])$
- $C(n).E = \{(i_j, v), j \in [k]\} \cup \{(v, o_j), j \in [m]\}$
- $C(n).M(v) = n$.

$$\frac{i_0 : A_0, \ldots, i_{k-1} : A_{k-1}, n : A_0 \times A_1 \times \ldots A_{k-1} \to B}{\forall j \in [m].n.o_j : B}$$

All output edges of the circuit produce the same value. $[\![C(n).o_j]\!] : [\![A_0]\!] \times \ldots [\![A_{k_1}]\!] \to [\![A]\!]$ given by $[\![C(n).o_j]\!] = [\![n]\!]$.

12.4.2 Delay node

A similar circuit construction rule is applicable to the z^{-1} node, with the difference that z^{-1} operates on streams. Given a type A we can construct the circuit Cz with a single vertex. The circuit has 1 input port of type $i : \mathcal{S}_A$ and any number m of output ports with the same type $\mathcal{S}_A$.

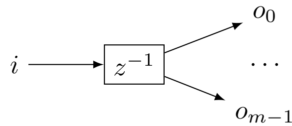

Formally:

- $Cz.V = \{v\}$.
- $Cz.I = (i)$.
- $Cz.O = (o_j, j \in [m])$.
- $Cz.E = \{(i, v)\} \cup \{(v, o_j), j \in [m]\}$.
- $Cz.M(v) = z^{-1} : \mathcal{S}_A \to \mathcal{S}_A$

$$\frac{i : \mathcal{S}_A}{Cz.o_j : \mathcal{S}_A}$$

All output edges of the circuit produce the same value. $[\![Cz.o_j]\!] : [\![\mathcal{S}_A]\!] \to [\![\mathcal{S}_A]\!]$ given by $[\![Cz.o_j]\!] = [\![z^{-1}]\!]$.

12.4.3 Sequential composition

Given two WFCs C and D with inputs (and necessarily outputs as well) in the same clock domain, their sequential composition is specified by a set of pairs of ports; each pair has an output port of C and an input port of D: $P = \{(o_j, i_j), j \in [n]\}$, where $o_j \in C.O$ and $i_j \in D.I$ with respectively matching types. The sequential composition $C +_P D$ is a new circuit where each output port of C is "connected" with the corresponding input port of D from P.

To simplify the definition, we can assume WLOG that each of C has exactly 2 inputs and outputs and D has 2 inputs and 1 output, and also that the second output of C is connected to first input of D (i.e., P contains at most one pair or ports). Circuits and connections with more inputs and outputs can be built by "bundling" multiple edges using the pair operator $\langle \cdot, \cdot \rangle$. Sequential composition is given by the following diagram:

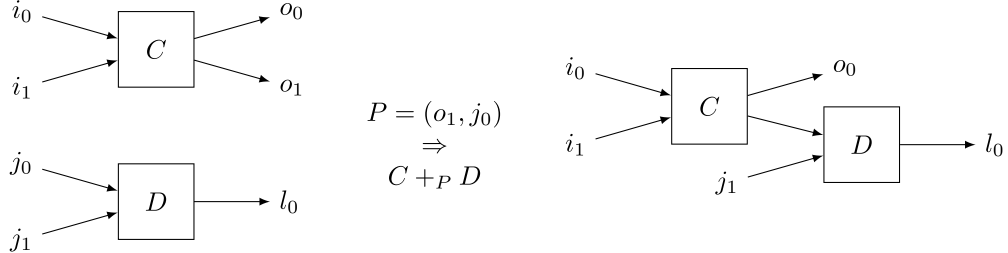

Formally the definition is given by:

- $(C +_P D).I = C.I \cup D.I \setminus \{i \mid \exists o.(o, i) \in P\}$.
- $(C +_P D).O = C.O \setminus \{o \mid \exists i.(o, i) \in P\} \cup D.O$.
- $(C +_P D).V = C.V \cup D.V$.
- $(C +_P D).E = C.E \cup D.E \setminus \{(v, o) \mid \exists i.(o, i) \in P, v \in C.V\}$ $\setminus \{(i, v) \mid \exists i.(o, i) \in P, v \in D.V\} \cup \{(u, v) \mid \exists (o, i) \in P, (u, o) \in C.E, (i, v) \in D.E\}$.

- $(C +_P D).M = \{v \mapsto C.M(v) \mid v \in C.V\} \cup \{v \mapsto D.M(v) \mid v \in D.V\}$.

$$\frac{\begin{array}{c} C : A_0 \times A_1 \to B_0 \times B_1 \\ D : B_1 \times A_2 \to B_2 \\ P = \{(C.O_0, D.I_0)\} \end{array}}{D +_P C : A_0 \times A_1 \times A_2 \to B_0 \times B_2}$$

$$[\![(C +_P D).O_0]\!] = [\![C.O_0]\!]$$
$$[\![(C +_P D).O_1]\!] = \lambda i_0, i_1, j_1.[\![D]\!]([\![C]\!](i_0, i_1), j_1)$$

This transformation combines acyclic graphs into an acyclic graph.

MIHAI: We may also need a parallel composition

12.4.4 Adding a back-edge

We can assume WLOG that we are given a circuit C with two inputs and two outputs. Assume that C's output port $o \in C.O$ has a stream type $o : \mathcal{S}_A$ and the input port $i \in C.I$ has the same type $i : \mathcal{S}_A$. We can create a new circuit $C \overset{\frown}{io}$ by adding two nodes: one z node implemented by z^{-1} and one p node implemented by $+_{\mathcal{S}_A}$, and a new input port $i' : \mathcal{S}_A$, connected as in the following diagram:

VAL: According to the model, the two outputs can produce different streams. Algebraically, this means the circuit C is implementing two operators, one for each output, say $T(i_0, i)$ for output o and $T_0(i_0, i)$ for output o_0.

MIHAI: Yes, that is correct. In fact, this is the main difference between circuits and operators: a circuit can have many different outputs, while an operator always has exactly 1.

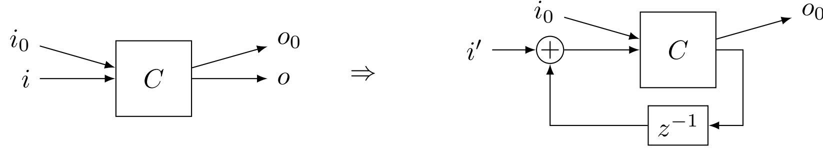

Formally the definition is given by:

- $(C \overset{\frown}{io}).I = C.I \setminus \{i\} \cup \{i'\}$.
- $(C \overset{\frown}{io}).O = C.O \setminus \{o\}$.
- $(C \overset{\frown}{io}).V = C.V \cup \{z, p\}$.
- $(C \overset{\frown}{io}).E = C.E \cup \{(i', p)\} \setminus \{(u, o) \in C.E \mid u \in C.V\} \cup \{(p, v) \mid \exists i.(i, v) \in C.E\} \cup \{(u, p) \mid \exists (u, o) \in C.E\}$.
- $(C \overset{\frown}{io}).M = C.M \cup \{p \mapsto +_{\mathcal{S}_A}, z \mapsto z^{-1}\}$.

We call the edge connecting z^{-1} to $\bigoplus$ a *back-edge*. One can prove by induction on the structure of C that the graph of C ignoring all back-edges is acyclic.

$$\frac{C : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_A \times \mathcal{S}_C}{C \overset{\frown}{io} : \mathcal{S}_A \times \mathcal{S}_B \to \mathcal{S}_C}$$

Since the source of a back-edge is always z^{-1}, strict operator, and any circuit is causal, the composition has a well-defined semantics, according to Corollary 3.17.

$$[\![C \overset{\frown}{io}]\!] = \lambda i_0, i'.\text{fix}\, i.[\![C.O_0]\!](i_0, i' + [\![z^{-1}]\!](i)).$$

VAL: For notation T, T_0 please see my comment above. Which one of T or T_0 is $[\![C.O._0]\!]$ in this definition? I think you intend it to be T_0 in order to produce the right output. But then, the well-definedness of the semantics of this construction does not follow from Corollary 3.17 because that result uses T rather than T_0. I am working on the more general result that we need to justify this case.

12.4.5 Lifting a circuit

Given a circuit C with scalar inputs and outputs, we can lift the entire circuit to operate on streams. As before, we can assume WLOG that C has a single input and output. If the circuit is a function: $C : A \to B$, the lifted circuit $\uparrow C$ operates time-wise on streams: $\uparrow C : \mathcal{S}_A \to \mathcal{S}_B$.

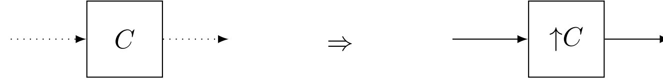

Formally the definition is given by:

- $(\uparrow C).I = C.I$.
- $(\uparrow C).O = C.O$.
- $(\uparrow C).V = C.V$.
- $(\uparrow C).E = C.E$.
- $(\uparrow C).M = \{\uparrow(C.M(v)) \mid v \in C.V\}$.

$$\frac{C : A \to B}{\uparrow C : \mathcal{S}_A \to \mathcal{S}_B}$$

If C is a WFC, then $[\![\uparrow C]\!] = \lambda s.[\![C]\!] \circ s$, for $s : \mathbb{N} \to B = \mathcal{S}_B$, as described in Section 3.1.

12.4.6 Bracketing

This construction uses nodes $\delta_0 : A \to \mathcal{S}_A$ and $\int : \mathcal{S}_A \to A$ which "create" and "eliminate" streams, as defined in Section 7.1. These nodes are always used in pairs.

Given a WFC C computing on streams with a single input $i : \mathcal{S}_A$ and a single output of the exact same type $o : \mathcal{S}_A$, we can "bracket" this circuit with a pair of nodes δ_0 and $\int$ as follows:

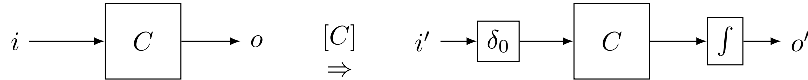

The types of the resulting input and given by $i' : A$, and $o' : A$.

It is very important for C to have a single input and output; this prevents connections to C from going "around" the bracketing nodes. If multiple input

or outputs are needed to C, they can be "bundled" into a single one using a pairing operator $\langle \cdot, \cdot \rangle$.

Formally the definition is given by:

- $[C].I = \{i'\}$.
- $[C].O = \{o'\}$.
- $[C].V = C.V \cup \{d, s\}$.
- $[C].E = C.E \cup \{(i', d), (s, o')\} \setminus \{(i, v) \mid v \in V\} \setminus \{(v, o) \mid v \in V\} \cup \{(d, v) \mid (i, v) \in C.E\} \cup \{(v, s) \mid (v, o) \in C.E\}$.
- $[C].M = C.M \cup \{d \mapsto \delta_0, s \mapsto \int\}$.

$$\frac{C : \mathcal{S}_A \to \mathcal{S}_B}{[C] : A \to B}$$

The semantics of the resulting circuit is just the composition of the three functions: $[\![[C]]\!] = [\![\int]\!] \circ [\![C]\!] \circ [\![\delta_0]\!]$. (However, note that the semantics of $\int$ is only defined for streams that are zero almost everywhere.)

Theorem 12.1. For any WFCs all inputs and outputs are streams of the same "depth".

Proof. The proof proceeds by induction on the structure of the circuit. All construction rules maintain this invariant, assuming that it is true for all primitive nodes. □

13 Implementing WFC as Dataflow Machines

In this section we give a compilation scheme that translates a WFC into a set of cooperating state machines that implement the WFC behavior. Each primitive node is translated into a state machine, each circuit is translated into a control element, and each edge is translated into a communication channel between two state machines, storing at most one value at any one time.

There are essentially 6 kinds of nodes in our circuits:

MIHAI: I realize that this would be much simpler if we force the toplevel circuit to have exactly 1 input and output edge. I will work on that.

- Lifted scalar nodes.
- Delay nodes z^{-1} operating on streams.
- Delay nodes z^{-1} operating on nested streams.
- "Loop entry" nodes, corresponding to δ_0.
- "Loop exit" nodes, corresponding to $\int$.
- Controller nodes, corresponding to circuits.

There are 4 types of events in our implementation:

Reset events: cause a circuit to be initialized. The main effect is to cause z^{-1} nodes to initialize their internal state to 0. The reset events have no effect on any other node.

Latch events: these events cause z^{-1} nodes to emit their internal state as an output. The latch events have no effect on any other node.

Data events: these events signal to a node or circuit that data is present on one of the input channels.

Repeat events: signal that a loop has to perform one more iteration. Only sent between $\int$ and corresponding δ_0 nodes.

Here are the state machines of each of these nodes:

MIHAI: This is pseudocode, but it would probably look better as real code.

Environment state machine The environment feeds data to the input edges of a top-level circuit and retrieves results from the output edges. The environment is expected to operate in epochs, executing the following infinite loop:

- Send a reset event to circuit
- Repeat forever
  - Assign data to all input edges.
  - Wait for all output edges to receive a "data" event.
  - Collect results from all output edges.

Circuit state machine

1. On receipt of a reset event:
   - Send a reset event to all nodes in the circuit.
   - Send a latch event to all nodes in the circuit.
2. On receipt of data on an input edge send a data event to the destinations connected to the input edge. The environment of the circuit should send

Primitive node state machine

1. On receipt of a "data" event check if all inputs have received data. If they have, compute the output, and send it as a "data" event on all output wires.

z^{-1} **node state machine** Stores a value in the internal state.

1. On receipt of "reset" event set internal state to 0.
2. On receipt of "latch" event, send a data event on the output channel with the value that is already present there.
3. On receipt of "data" event, copy internal state to output channel and input to internal state.

z^{-1} node operating on a nested stream state machine This node internally stores a potentially unbounded *list* of values as internal state. It also maintains a counter "time" to index within this list.

1. On receipt of a "reset" even set time to 0.
2. On receipt of a "latch" event, send the value in list[time] as a "data" event to the output channel.
3. On receipt of a "data" event
   - Increment "time"
   - Set the output channel to the list[time] value (zero if the list is not long enough, and grow the list)
   - Store data value received in list[time].

Loop entry node state machine

1. On receipt of a "data" event
   - Send a "reset" event to circuit that is connected as output
   - Send a "data" event to the circuit connected as output with the data received
2. On receipt of a "repeat" event send a "data" event with value 0 to the circuit connected as output.

Loop exit node state machine Maintain an internal accumulator.

1. On receipt of a "reset" event set the accumulator to 0.
2. On receipt of a "data" event:
   - if the data value is 0, send a "data" event to the output channel with the current value of the accumulator.
   - otherwise add the input value to the accumulator and send a "repeat" event to the corresponding $delta_0$ node.

14 Implementing Differential Datalog in DBSP

This section gives an implementation of DDlog in terms of DBSP circuits. This is a precise specification of the semantics of DDlog.

```
1  DatalogProgram := TypeDeclaration*
2                    RelationDeclaration*
3                    Rule*
4  TypeDeclaration := ...
5  Type := ...
6  Expression := ...
7  Id := ... // identifiers
8  RelationDeclaration := ("input"|"output")? "relation"
9                         Id "(" ColumnTypes? ")"
10 ColumnTypes := ColumnType ( "," ColumnType)*
11 ColumnType := Id ":" Type
12 Rule := Head ":-" Body
13 Head := RelationTerm
14 RelationTerm := Id "(" Columns? ")"
15 Columns := Variable ( "," Variable )*
16 Body := RelationTerm ( "," Term )?
17 Term := RelationTerm
18      | Predicate
19      | NegatedTerm
20      | VariableDefinitionTerm
21      | FlatmapTerm
22      | GroupByTerm
23 NegatedTerm := "not" RelationTerm
24 VariableDefinitionTerm := "var" Variable "=" Expression
25 FlatmapTerm := "var" Variable "=" "Flatmap" "(" Id ")"
26 GroupByTerm := "var" Variable "="
27      Expression "." "group_by" "(" VariableList? ")"
28 Variable := Id
29 VariableList := Id (, Id)*
30 Predicate := Expression
```

Figure 2: Datalog rules grammar.

14.1 Differential Datalog syntax and semantics

We describe the syntax and semantics of a dialect of Datalog called Differential Datalog, or DDlog. We start by ignoring the differential aspects, we will return to these in Section 14.3. In defining the syntax and semantics of Datalog we mostly follow standard definitions, e.g., [5]. In this section we give a syntax-directed translation of Datalog programs into circuits. We model a *core* of the language (ignoring constructs that can be viewed as syntactic sugar), to simplify the description. We argue informally that the resulting circuits implement the standard semantics of Datalog.

Our Datalog is strongly-typed, supports stratified negation and recursion, and is enhanced with additional operators, such as grouping (which we will describe below). Figure 2 shows the EBNF-like grammar of the core DDlog language (omitting types and expressions).

Types Assume that we are given a set of basic types, including `integer`, Booleans, `string`, but also structures (product types), unions (sum types), tuples, vectors, and sets. Each such type must support an operation to compare values for equality. This is the only requirement to store values of a particular type in a set.

We allow arbitrary computations over the base types (e.g., arithmetic) through the use of built-in functions (e.g., addition, subtraction, equality comparison, etc.). A standard language of expressions can be used to combine built-in functions into more complex functions. We require all such functions to be total and deterministic. We treat such computations as uninterpreted functions (black boxes) and no longer concern ourselves with them in this document.

All DDlog programs must be strongly typed, but we don't specify the typing rules in this document (e.g., predicates must produce Boolean results). The typing rules are standard. Only the semantics of well-typed programs is defined.

Relations Datalog programs compute over relations. The inputs and outputs of a Datalog program are relations. The standard Datalog semantics of a relation is a *set* of values from some domain.

DDlog programs continuously interact with their environment. Thus they distinguish relations by their roles. Some relations are `input` relations; their contents is supplied by the environment. Some relations are declared as `output` relations. The contents of these relations is visible to external observers.

A DDlog program must have a declaration for each relation, specifying the type of its elements, as in this example:

```
input relation People(name: string, age: integer)
relation Ages(age: integer)
output relation Names(name: string)
```

This declares three relations. The first relation is an `input` relation, named `People`, and it has 2 columns, `name` of type `string`, and `age` of type `integer`. Its elements are 2-tuples of type (`string`, `integer`).

The value of relation `People` is a *set* of such tuples. As a running example, let us assume that the input relation `People` has the value (supplied by the program's environment) $\{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$, containing three tuples.

The relation `Names` has a single column, and it is an `output` relation. This means that the environment can observe it's contents. The relation `Ages` is neither `input` nor `output`. The Datalog program must contain *rules* that show how `Ages` and `Names` are computed from the `input` relations, i.e., `People`.

Rules Besides type and relation declarations, the most important part of a Datalog program is a set of rules. A **Datalog rule** defines the contents of a relation as a function of other relations. A rule has a **head** and a **body**, as shown in the grammar from Figure 2. In the following rule:

```
Names(n) :- People(n, a).
```

the head is to the left of the turnstyle symbol `:-`, and it is always a relation with variables standing for the tuple fields. This rule defines how `Names` is computed from the contents of `People`. In this example the rule's head is `Names(n)`. You may guess that the value of `Names` is the set {`bob`, `john`, `amy`}. We will explain how this result is computed. The turnstyle symbol can be read as "if". `n` is a variable of type `string`, standing for the column `name` (the type of `n` is `string` because it stands positionally for the declared column `name` with type `string` of relation `Names`). The values that `n` may take are defined by the body of the rule — each variable in the head must appear in the body of the rule.

The body of a rule consists of one or two **terms** separated by commas; a comma is read as an "and", and such a rule is form of a *conjunctive query*. The valuation computed by the body is defined recursively on the list of terms in the body. Datalog allows an arbitrary number of terms in a rule body, but our grammar allows only two. We argue in the paragraph below that this does not reduce the power of the language, but it simplifies the description.

The grammar of terms is shown in line 16 in Figure 2. We will discuss the semantics of the various terms and their implementation as circuits in the rest of this section.

Valuations The body of the above rule is `People(n, a)`. The body defines two variables, `n` and `a`. The semantics of a rule body is a **valuation**: a set of values that the variables defined by the body may *jointly* take. Our example body defines a valuation for the tuple of variables (`n`, `a`). Since the body is `People(n,a)`, the valuation defines the set of values for $(\texttt{n}, \texttt{a})$ to be the contents of the relation `People`, that is $(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$.

We can assume without loss of generality that every rule body has at most two terms. A rule with n terms can be decomposed into n - 1 rules with 2 terms each by introducing a temporary relation that stores the entire valuation. For example, consider the following rule:

```
input relation Lives(name:string, country:string)
output relation USAges(age: integer)
USAges(a) :- People(n, a), Lives(n, c), c == "USA".
```

The rule prefix `People(n, a), Lives(n, c)` defines a valuation for variables `n, a, c`. By introducing a temporary relation `Temp(n, a, c)` we can rewrite this rule as two rules, producing the same result for the visible relatin `USAges`:

```
relation Temp(name:string, age:integer, country:string)
Temp(n, a, c) :- People(n, a), Lives(n, c).
USAges(a) :- Temp(n, a, c), country == "USA".
```

14.2 Compiling Datalog programs to circuits

A Datalog program is a set of rules computing over valuations and relations. Both relations and valuations are represented as $\mathbb{Z}$-sets in a translation to cir-

cuits, by representing each set by a $\mathbb{Z}$-set with weights 1. (Subsequent optimizations can relax this requirement for internal relations as long as the semantics of output relations is preserved, since only `output` relations are observable from the environment. We expand on this in Section 6.1.)

14.2.1 Rule compilation

Each rule is compiled to a circuit with the following properties:

- Each relation and valuation is represented by a $\mathbb{Z}$-set.
- The head of the rule is the output edge of the circuit;
- The relations that appear in the body are inputs of the circuit;
- The circuit structure is defined by the terms that appear in the relation body (as discussed in the rest of this section);
- The turnstyle is compiled into a projection operator (described in Section 14.2.4);
- The *valuation* at a conjunction in the rule body is translated into an edge in the circuit, carrying $\mathbb{Z}$-set values.

This translation is illustrated in Figure 3.

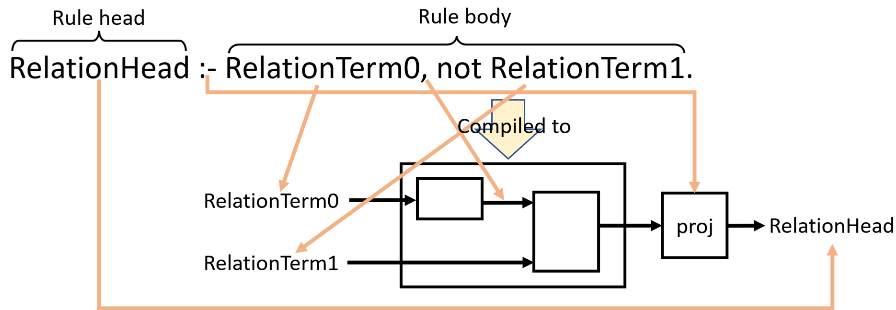

Figure 3: Compilation of a DDlog rule into a circuit.

The compilation of a program containing a set of rules produces a "toplevel" circuit, composed of the circuits for the rules interconnected with each other (as described in Section 14.2.2). For the toplevel circuit the `input` relations correspond to the input edges. (`input` relations cannot appear in the head of any rule). Similarly, the `output` relations will correspond to output edges of the toplevel circuit, (each `output` relation must appear in some rule head).

14.2.2 Relation terms in rule bodies

A `RelationTerm` is a term in a rule body containing a relation with variables substituted for the columns: `People(n, a)` is such an example. This term *defines a valuation* for all variables that appear in the columns; the valuation associates the variables with the contents of the relation itself. In our example the term `People(n, a)` defines the following valuation: $(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$. Each Datalog relation is represented by an edge in a circuit carrying a $\mathbb{Z}$-set.

As Figure 3 shows, a relation in the head of a rule is compiled into the output edge of the circuit corresponding to the rule, and that a `RelationTerm` in the body of a rule is compiled into an input edge of the circuit corresponding to the rule.

A relation that appears in the body of a rule and in the head of another rule is compiled into an edge connecting the circuits representing the two rules. This is in fact just a form of function composition.

(This rule does not apply directly for recursive rules; the translation for recursive or mutually recursive rules (which define a relation in terms of itself), is described in Section 8.1.1).

For example, for the following Datalog program structure, where `R` is used within a body and within a separate head:

```
R(y) :- I(x), ....
O(y) :- ..., R(y).
```

and, given circuits C_R implementing the first rule and C_O implementing the second rule:

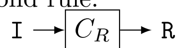

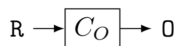

the translation of the program with both rules is:

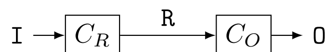

This construction is repeated for all rules, translating a program with n rules into n circuits connected to each other. If the rules are not recursive the resulting circuit is acyclic.

14.2.3 Repeated rule heads (set union)

The same relation may appear in the head of multiple rules. In this case the contents of the head relation is the *set union* of the values assigned by all heads. Consider the following example, where `I1` and `I2` are rule bodies of arbitrary complexity providing a valuation for variable `v`:

```
O(v) :- I1(v).
O(v) :- I2(v).
```

The following circuit implements the Datalog program with both rules:

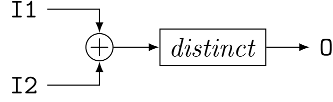

Given two $\mathbb{Z}$-sets $a \in \mathbb{Z}[I]$ and $b \in \mathbb{Z}[I]$ which are sets (i.e., $\text{isset}(a)$ and $\text{isset}(b)$), their *set union* can be computed as: $\cup : \mathbb{Z}[I] \times \mathbb{Z}[I] \to \mathbb{Z}[I]$. $a \cup b \stackrel{\text{def}}{=} distinct(a +_{\mathbb{Z}[I]} b)$. The *distinct* application is necessary to provide the set semantics of Datalog. We have $\text{isset}(a) \wedge \text{isset}(b) \Rightarrow \text{isset}(a \cup b)$ and $\text{ispositive}(a) \wedge \text{ispositive}(b) \Rightarrow \text{ispositive}(a \cup b)$.

Consider a concrete example for the above program where the value of `I1(v)` is $\texttt{v} \in \{\texttt{bob} \mapsto 1, \texttt{mike} \mapsto 1\}$ and the value of `I2(v)` is $\texttt{v} \in \{\texttt{bob} \mapsto 1, \texttt{john} \mapsto 1\}$. In terms of $\mathbb{Z}$-sets we are performing the following addition:

v	**W**
bob	1
mike	1

$+$

v	**W**
bob	1
john	1

$=$

v	**W**
bob	2
mike	1
john	1

It is apparent why the *distinct* operator is needed.

14.2.4 Projection

Given a valuation produced by the body of a rule, the head of the rule defines the contents of a relation as *the projection* of the valuation on the variables used in the head. For our example rule `Names(n) :- People(n, a)`, the body defines a valuation for $(\texttt{n}, \texttt{a})$, but the head uses only `n`. The projection of the valuation $(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$ on the variable `n` is the valuation $\texttt{n} \in \{\texttt{bob}, \texttt{john}, \texttt{amy}\}$. This defines the contents of the relation in the head: $\texttt{Names} = \{\texttt{bob}, \texttt{john}, \texttt{amy}\}$.

Thus, in Datalog projection is used when some of the bound variables in the body of a rule are not used in the head. We can assume without loss of generality that a single variable is removed in a projection (by bundling multiple variables in a single tuple-valued variable). Let us consider the following example, where `I` stands for a rule body producing a valuation for `(v, v1)`.

```
O(v) :- I(v, v1).
```

Here the type of the implementation of `I` is $\mathbb{Z}[A_0 \times A_1]$ (a $\mathbb{Z}$-set of tuples with two elements), while the type of the implementation of `O` is $\mathbb{Z}[A_0]$. In terms of $\mathbb{Z}$-sets, the projection of a $\mathbb{Z}$-set i on A_0 is defined as: $\pi_0(i)[t] = \sum_{x \in i, x|_0 = t} i[x]$, where $x|_0$ is first component of the tuple x. The multiplicity of a tuple in the result is the sum of the multiplicities of all tuples that project to it.

As a concrete example of projection, consider the $\mathbb{Z}$-set corresponding to the `People` relation and it's projection on the `Age` column. The projection is $\pi_{Age}(\texttt{People}) = \{10 \mapsto 1 + 1, 20 \mapsto 1\}$. Notice that in the projection the weight of 10 is the sum of all weight of the tuples that have age 10, i.e., 2.

The circuit for such a rule is:

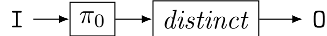

Note that $\uparrow\pi$ is time-invariant, since π has the zero-preservation property. We have $\text{isset}(i) \Rightarrow \text{isset}(\pi_A(i))$ and $\text{ispositive}(\pi_0)$.

14.2.5 Flatmap in DDlog

Recall that in DDlog the type of a column in a relation can be any of the types supported by the language, including complex types, such as vectors, sets, or even maps. For example, the following declaration indicates that the values in relation `I` are sets of integers.

```
relation I(set: Set<integer>)
```

Flatmap is an operator that can expand the data in such a collection value stored in a relation into the contents of a relation. Classic Datalog does not support flatmaps. In DDlog `Flatmap` is an explicit keyword. It appears in rules in the form of a `FlatmapTerm` in the grammar in Figure 2. The DDlog type system ensures that the `Flatmap` operator can only be applied to an expression whose type is a collection.

Here is an example of a program using `Flatmap`:

```
relation I(set: Set<integer>)
relation O(integer)
O(v) :- I(set), var v = Flatmap(set).
// O = union of all sets in I
```

Each element in relation `I` is a set of integers. The DDlog `Flatmap` operator implements a restricted form of the functionality of the general mathematical operator from above (unlike the mathematical flatmap, which is parameterized by a function f, the DDlog `Flatmap` uses a hardwired function, essentially the identity function.). The semantics is as follows: the `Flatmap` rule body term extends the existing valuation with a new variable, `v` in this example. `Flatmap`'s argument is an expression that depends on the current valuation (`set` in this example) whose value is a collection. Let us assume that the contents of the `I` relation is: $\{\{1,2\},\{2,3\}\}$.

The valuation produced by the rule `I(set), var v = Flatmap(set)` is the following: $(\texttt{set}, \texttt{v}) \in \{(\{1,2\},1),(\{1,2\},2),(\{2,3\},2),(\{2,3\},3)\}$.

The circuit-based implementation of `Flatmap`, operating on $\mathbb{Z}$-sets, can be defined as follows:

$$\texttt{I} \rightarrow \boxed{\text{flatmap}(e)} \rightarrow \boxed{\mathit{distinct}} \rightarrow \texttt{O}$$

where the function e extends each tuple in a $\mathbb{Z}$-set with the newly introduced variable and each set value with a Cartesian product between the collection and all it's elements. For our example: $e : \texttt{Set<integer>} \rightarrow \mathbb{Z}[\texttt{Set<integer>} \times \texttt{integer}]$ defined by: $e(\texttt{set}) = \sum_{x \in \texttt{set}} (\texttt{set}, x) \mapsto 1$.

The `distinct` operator is needed because some collections (e.g., vectors) may contain duplicate values.

Proposition 14.1. ispositive(`Flatmap`).

14.2.6 Map in DDlog

Given a function $f : A \to B$, the mathematical **map** operator "lifts" the function f to operate on $\mathbb{Z}$-sets: $\text{map}(f) : \mathbb{Z}[A] \to \mathbb{Z}[B]$. map can be defined in terms of flatmap: $\text{map}(f) \stackrel{\text{def}}{=} \text{flatmap}(x \mapsto 1 \cdot f(x))$.

Classic Datalog does not support map computations, but many practical implementations do. DDlog programs perform map computations when using `VariableDefinitionTerm` in a rule, by using an expression to computing a value for a new variable, that is added to a valuation, as in the following example:

```
O(v) :- I(x), var v = x + 1.
```

The `VariableDefinitionTerm`, `var v = x + 1`, extends the current valuation, which contains just `x`, to include the newly defined variable `v`.

The circuit implementation of the previous rule is:

I → map(e) → O

The function e extends the current valuation tuple with a new column (corresponding to `v` in the example) and evaluates the expression in the term (`x + 1`) for each row of the valuation to compute the corresponding value for the new column. In our example, $e : \mathbb{Z}[\texttt{integer}] \to \mathbb{Z}[\texttt{integer} \times \texttt{integer}]$, $e(x) = (x, x{+}1)$.

Note that $\text{ispositive}(\text{map}(f))$ for any function f. From the linearity of flatmap it follows that map is linear as well. Moreover, the operator $\uparrow\!\text{map}(f)$ is time-invariant for any f.

14.2.7 Filtering

Filtering occurs in Datalog whenever a `TermPredicate` appears in the body of a rule, in the guise of a Boolean expression, as in the following example:

```
relation Minors(n: string, a: integer)
Minors(n, a) :- People(n, a), a < 18.
```

(A predicate may not appear in the first position in the body of a rule.) The predicate must only use variables in the current valuation. The produced valuation contains the same variables as the source valuation. The value of the valuation the set of tuples in the source valuation that satisfy the predicate.

Recall that the valuation of the term `People(n, a)` is $(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$. The valuation of the entire rule is the set of tuples $(\texttt{n}, \texttt{a})$ for which the predicate $\texttt{a} < 18$ holds. That valuation is $(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{amy}, 10)\}$. Thus the contents of relation `Minors` is $\{(\texttt{bob}, 10), (\texttt{amy}, 10)\}$.

To compute on $\mathbb{Z}$-sets, let us assume that we are filtering with a predicate $P : A \to \mathbb{B}$. We define the following function $\sigma_P : A \to \mathbb{Z}[A]$ as:

$$\sigma_P(x) = \begin{cases} 1 \cdot x & \text{if } P(x) \\ 0 & \text{otherwise} \end{cases}$$

The filter of a $\mathbb{Z}$-set is defined as $\text{filter}_P : \mathbb{Z}[A] \to \mathbb{Z}[A]$ by $\text{filter}_P \stackrel{\text{def}}{=} \text{flatmap}(\sigma_P)$. We have $\text{isset}(i) \Rightarrow \text{isset}(\text{filter}_P(i))$ and $\text{ispositive}(\text{filter}_P)$. Thus

a *distinct* is not needed. As a consequence of the linearity of flatmap, we have that filtering is also linear. The lifted version of filtering is also time-invariant.

The circuit for filtering with a predicate P can be implemented as:

14.2.8 Grouping

Classic Datalog does not support grouping. In DDlog grouping is the fundamental operator used for aggregation. Grouping is applied to a set and produces a partition of that set into a set of collections. The type of a partition is a built-in type in DDlog, called `Group`. The following example shows an example of grouping in DDlog:

```
output relation O(v: Group<integer, string>)
// Groups with key integer and values Vector<string>
ByAge(g) :- People(n, a), var g = (n).group_by(a).
// Each g is the group of all names that have the same age
```

The general syntax of a `GroupByTerm` in Figure 2 is given by:
`var g = (project-expression).group_by(key-expression)`. In DDlog the `key-expression` is restricted to be a tuple of variables in the current valuation.

The semantics of a `GroupByTerm` is given by the following algorithm:

1. We start with some input valuation.

2. The `project-expression` is evaluated for the input valuation V, adding a new (anonymous) variable to the valuation, storing the result of the `project-expression` for each row.

3. The resulting valuation is projected on a tuple of variables containing the new anonymous variable and all variables that appear in `key-expression`. The result of the projection is a new valuation P, a set (with no duplicates). This valuation *only includes the variables that appear in the projection and the key expression*; all other variables are removed. This behavior is unusual — this is the only DDlog operator that removes variables from a valuation.

4. Finally, the data in the valuation is grouped by key, and the result is a valuation that contains for the new variable a group.

As an example, let us evaluate the above DDlog rule according to these steps:

1. The term `Persons(n, a)` provides the following valuation:
$(\texttt{n}, \texttt{a}) \in \{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\}$.

2. In our example `project-expression` is just `n`, whose value will be assigned to the anonymous variable. This creates a new valuation:
$(\texttt{n}, \texttt{a}, \texttt{anonymous}) \in \{(\texttt{bob}, 10, \texttt{bob}), (\texttt{john}, 20, \texttt{john}), (\texttt{amy}, 10, \texttt{amy})\}$.

3. The valuation is projected on `a` (the group key) and `anonymous` (the projection key), obtaining: $(\texttt{a}, \texttt{anonymous}) \in \{(10, \texttt{bob}), (20, \texttt{john}), (10, \texttt{amy})\}$. In this case there are no duplicates, but any duplicates would be removed.

4. The values in the valuation are grouped by their `a` value, providing a new group for each value. The variable `g` is added to the resulting valuation. The result is $(\texttt{a}, \texttt{g}) \in \{(10, [\texttt{bob}, \texttt{amy}]), (20, [\texttt{john}])\}$. Notice how the value of `g` in the valuation is a collection, shown with square brackets.

Let us define this computation in terms of $\mathbb{Z}$-sets. Consider an arbitrary type of keys K, and a function that computes a key for a value $k : I \rightarrow K$. Then we define $\text{groupby}(k) : \mathbb{Z}[I] \rightarrow \mathbb{Z}[I][K]$, as $\text{groupby}(k)(i) = \sum_{x \in i} \{k(x) \mapsto 1 \cdot x\}$. Note that groupby always produces a set of $\mathbb{Z}$-sets. The weight of each group is always 1. Note that $\text{ispositive}(\text{groupby}(k))$. Also, $\uparrow\text{groupby}(k)$ is time-invariant for any function k, since $\text{groupby}(k)$ has the zero-preservation property.

The implementation of the `group_by` operator requires chaining the implementation of the projection and of groupby function just described: the resulting circuit is (using p as the translation of `project-expression` and k as the translation of `key-expression`):

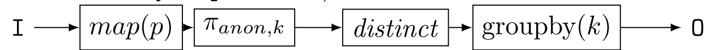

14.2.9 Aggregation

Classic Datalog does not support aggregations but many practical implementations have extended Datalog with a construct equivalent with a composition of groupby-aggregate.

Strictly speaking, DDlog does not support for aggregation – the only aggregate supported is a group. However, since DDlog allows users to apply arbitrary functions to a `Group` object, traditional aggregation can be performed by grouping and then applying a scalar-returning function using a `map`, as described Section 14.2.6. Consider the following example, an extension of the example in Section 14.2.8:

```
output relation NamesByAge(s: string)
NamesByAge(s) :- ByAge(g),
      var s = g.key + ": " + g.toString().
```

This example uses built-in functions `g.key` that obtains the key of a group, and `g.toString()`, which converts the group contents to a string.

Formally, given a function a : `Group<`K`,`I`>` $\rightarrow O$, aggregation is just $\text{map}(a)$ applied to a set of groups. As a consequence lifted aggregation is time-invariant.

14.2.10 Cartesian products

Cartesian products in Datalog appear from the use of in a rule body of a `TermRelation` where the relation arguments are all new variables (not already defined in the input valuation). The Datalog semantics of Cartesian products

is to produce a new valuation that includes all variables, and having as values the Cartesian product of the values of the input valuation and the relation in the term.

The following program shows an example Cartesian product:

```
O(v1, v2) :- I1(v1), I2(v2).
```

A Cartesian product is implemented as a circuit using the product operation on $\mathbb{Z}$-sets. For $i_1 \in \mathbb{Z}[A]$ and $i_2 \in \mathbb{Z}[B]$ we define $i_1 \times i_2 \in \mathbb{Z}[A \times B]$ by $(i_1 \times i_2)(\langle x, y \rangle) \stackrel{\text{def}}{=} i_1[x] \times i_2[y].\forall x \in i_1, y \in i_2$. The circuit computing the Cartesian product is given by:

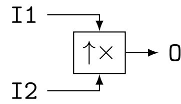

As an example, let us consider the product of the following two $\mathbb{Z}$-sets:

x	**W**
bob	1
mike	2

$\times$

y	**W**
bob	1
john	-1

$=$

(x, y)	**W**
(bob, bob)	1
(mike, bob)	2
(bob, john)	-1
(mike, john)	-2

Notice that $\text{isset}(x) \wedge \text{isset}(y) \Rightarrow \text{isset}(x \times y)$. The Cartesian product as defined on $\mathbb{Z}$-sets is a bilinear operator. Also, $\uparrow\times$ is time-invariant.

14.2.11 Joins

A join appears in a Datalog program by using a `TermRelation` that has as arguments some variables that are already defined in the current valuation. The following example shows a join: since the second relation reuses variable `v`, which is already bound by the valuation, this is a join, and not a Cartesian product:

```
O(x, y) :- I1(x, v), I2(v, y).
```

The semantics of a join can be modeled as a cartesian product followed by a sequence of filters. This is achieved by using fresh variable names for the arguments of each `TermRelation`, and adding predicates that require these fresh variables to be equal to the bound variables they replace. For example, the following program is equivalent to the one above.

```
O(x, y) :- I1(x, v), I2(v1, y), v = v1.
```

Since a join is a composition of a bilinear (the Cartesian product) and a linear (filtering) operator, it is also a bilinear operator, and thus its lifted version is time-invariant.

In practice joins are very important computationally, and they are implemented by a custom operator on $\mathbb{Z}$-sets denoted by $\bowtie$.

The circuit computing the join product is given by:

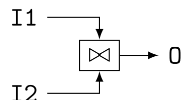

14.2.12 Set intersection

Set intersection in Datalog is just a particular case of join where a `TermRelation` uses *all* the variables defined in the current valuation as arguments. In the following example relation `O` is the intersection of relations `I1` and `I2`:

```
O(v) = I1(v), I2(v).
```

The join implementation using circuits immediately applies to set intersections. It follows that set intersection is a bilinear operator, and thus time-invariant when lifted.

14.2.13 Negation

We only support Datalog programs with stratified negation. See [5] for a precise definition. Negation in a Datalog program is introduced syntactically by a `NegatedTerm` from the grammar in Figure 2. A negated term cannot appear first in a rule body. All variables that appear in the negated term must have been already defined by previous terms in the body.

With these syntactic constraints the are two different meanings to negation:

- If the negated term uses *all* variables already in the valuation, it is modeled as a set difference.
- If the negated term uses a subset of all the variables in the existing valuation, it is modeled as an antijoin.

We describe each of these two cases.

14.2.14 Set difference

If the `NegatedTerm` contains as arguments all variables in the current valuation, the meaning of negation is just set difference: the resulting valuation will *exclude* all tuples from the negated relation.

For example, consider the rule:

```
relation Major(name: string, age: integer)
Major(n, a) :- People(n, a), not Minor(n, a).
```

The valuation computed by the rule's body is $\{(\texttt{bob}, 10), (\texttt{john}, 20), (\texttt{amy}, 10)\} \setminus \{(\texttt{bob}, 10), (\texttt{amy}, 10)\} = \{(\texttt{john}, 20)\}$.

In terms of $\mathbb{Z}$-sets, let us consider the following program:

```
O(v) :- I1(v), not I2(v).
```

We define the set difference on $\mathbb{Z}$-sets as follows: $\setminus : \mathbb{Z}[I] \times \mathbb{Z}[I] \rightarrow \mathbb{Z}[I]$, where $i_1 \setminus i_2 = distinct(i_1 - i_2)$. Note that we have $\forall i_1, i_2, \text{ispositive}(i_1 \setminus i_2)$ due to the application of the *distinct* operator.

The circuit computing the valuation of the body of this rule is:

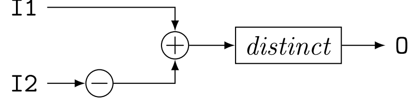

This whole circuit is time-invariant, since it composed only of time-invariant operators.

14.2.15 Antijoin

Antijoin is the semantics of a Datalog `NegatedTerm` that uses a relation which does not use some of the variables in the current valuation. Consider the following program:

```
O(v) :- I1(v, z), not I2(v).
```

The semantics of such a rule can be defined in terms of joins and set difference. This rule is equivalent with the following pair of rules:

```
C(v, z) :- I1(v, z), I2(v).
O(v) :- I1(v, z), not C(v, z).
```

This transformation reduces an antijoin to a join (using all variables in the current valuation), followed by a set difference. The translation of these rules is covered by Sections 14.2.11 and Section 14.2.14. In terms of circuits we can just build the circuit for the pair of rules:

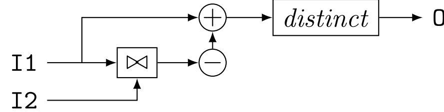

14.3 Streaming Differential Datalog

In this section we have shown how Given a Datalog (or SQL) query Q, can be converted into a circuit C_Q that computes the same input-output function as Q.

We can perform two simple transformations to this circuit: we can lift it to convert it into a streaming program, and then we can incrementalize it, to convert it into a differential program.

14.3.1 Streaming Datalog

Given the circuit $C_Q : A \rightarrow B$, We can lift it to compute on *streams* of relations. $\uparrow C_Q : \mathcal{S}_A \rightarrow \mathcal{S}_B$ interacts with its environment in "epochs," corresponding to

the time dimension of the streams. In each epoch the circuit receives a new set of values for the inputs relations and it provides the corresponding values for the output relations.

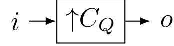

14.3.2 Streaming Differential Datalog

Furthermore, we can apply the $\cdot^{\Delta}$ operator to the streaming circuit $\uparrow C_Q$, converting it into an incremental streaming circuit: $(\uparrow C_Q)^{\Delta} : \mathcal{S}_A \rightarrow \mathcal{S}_B$.

This is a differential streaming version of the circuit C_Q. This circuit interacts with its environment in "ephochs," corresponding to the time dimension of the streams. In each epoch the circuit receives a new set of *changes* to the inputs relations and it provides the corresponding *change* for the output relations.

This is in essence the service provided by the DDlog compiler: given a query Q it provides a streaming implementation of $\uparrow C_Q{}^{\Delta}$. However, the DDlog runtime provides some additional services, described in the next section.

15 Implementations

15.1 DBSP Rust library

We have built an implementation of DBSP as part of an open-source project with an MIT license: https://github.com/vmware/database-stream-processor. The implementation consists of a Rust library and a runtime. The library provides APIs for basic algebraic data types: such as groups, finite maps, $\mathbb{Z}$-set, indexed $\mathbb{Z}$-set. A separate circuit construction API allows users to create DBSP circuits by placing operator nodes (corresponding to boxes in our diagrams) and connecting them with streams, which correspond to the arrows in our diagrams. The library provides pre-built generic operators for integration, differentiation, delay, nested integration and differentiation, and a rich library of $\mathbb{Z}$-set basic incremental operators: corresponding to plus, negation, grouping, joining, aggregation, *distinct*, flatmap, window aggregates, etc.

For iterative computations the library provides the δ_0 operator and an operator that approximates $\int$ by terminating iteration of a loop at a user-specified condition (usually the condition is the requirement for a zero to appear in a specified stream). The low level library allows users to construct incremental circuits manually by stitching together incremental versions of various primitive operators.

The library has also support for multicore execution of $\mathbb{Z}$-set operators (using a natural sharding strategy), and a variety of adaptors for external data sources (e.g., Kafka, CSV files, etc). The library can also spill internal operator state to persistent storage.

15.2 SQL compiler

We have also built a SQL to DBSP compiler, which can translate standard SQL queries into DBSP circuits. The compiler implements Algorithm 6.4, which can be used to generate the streaming version of any expressible SQL query. The compiler is also an open-source project `https://github.com/vmware/sql-to-dbsp-compiler` with an MIT license. The compiler front-end parser and optimizer is based on the Apache Calcite [11] infrastructure. The project is mature enough to pass essentially all 7 million SQL Logic Tests [1]. The compiler handles all aspects of SQL, including NULLs, ternary logic, grouping, aggregation, multiset queries, etc.

15.3 Formal verification

As a third implementation, we have formalized and verified all the definitions, lemmas, propositions, theorems, and examples in this paper using the Lean theorem prover; we make these proofs available at . This amounted to roughly 5K lines of Lean code.

MIHAI: Tej?

15.4 Additional Implementation Observations

15.4.1 Checkpoint/restore

DBSP programs are stateful streaming systems. Fault-tolerance and migration for such programs requires state migration. We claim that it is sufficient to checkpoint and restore the "contents" of all z^{-1} operator in order to migrate the state of a Ddlog computation.

15.4.2 Maintaining a database

DBSP is not a database, it is just a streaming view maintenance system. In particular, DBSP will not maintain more state than absolutely necessary to compute the changes to the views. There is no way to find out whether a specific value exists at a specific time moment within a DBSP relation. However, a simple extension to DBSP runtime can be made to provide a *view query* API: essentially all relations that may be queried have to be maintained internally in an integrated form as well. The system can then provide an API to enumerate or query a view about element membership between input updates.

15.4.3 Materialized views

An incremental view maintenance system is not a database – it only computes changes to views when given changes to tables. However, it can be integrated with a database, by providing capabilities for *querying* both tables and views. An input table is just the integral of all the changes to the table. This makes possible building a system that is both stateful (like a database) and streaming (like an incremental view maintenance system).

15.4.4 Maintaining input invariants

For relational query systems there is however an important caveat: the proofs about the correctness of the C_Q implementing the same semantics as Q all require some preconditions on the circuit inputs. In particular, the semantics of Q is only defined for sets. In order for C_Q to faithfully emulate the behavior of Q we must enforce the invariant that the input relations are in fact sets.

However, the differential streaming version of the circuits accepts an arbitrary stream of changes to the input relations. Not all such streams define input relations that are sets! For example, consider an input stream where the first element removes a tuple from an input relation. The resulting $\mathbb{Z}$-set does not represent a set, and thus the proof of correctness does not hold. This problem has been well understood in the context of the relational algebra: it is the same as the notion of positivity from [23].

We propose three different solutions to this problem, in increasing degrees of complexity.

Assume that the environment is well-behaved The simplest solution is to do nothing and assume that at any point in time the integral of the input stream of changes i is a set: $\forall t \in \mathbb{N}.\text{isset}(I(i)[t])$. This may be a reasonable assumption if the changes come from a controlled medium, e.g., a traditional database, where they represent legal database changes.

Normalize input relations to sets In order to enforce that the input relations are always sets it is sufficient to apply a *distinct* operator after integration.

$$i \longrightarrow \boxed{I} \longrightarrow \boxed{distinct} \rightarrow \boxed{\uparrow C_Q} \rightarrow \boxed{D} \longrightarrow o$$

The semantics of the resulting circuit is identical to the semantics of ${\uparrow C_Q}^{\Delta}$ for well-behaved input streams. For non-well behaved input streams one can give a reasonable definition: a change is applied to the input relations, and then non-relations are normalized into relations. Removing a non-existent element is a no-op, and adding twice an element is the same as adding it once.

Use a "change manager" In this solution we interpose a separate software component between the environment and the circuit. Let us call this a "change manager" (CM). The CM is responsible for accepting commands from the environment that perform updates on the input relations, validating them, and building incrementally an input change, by computing the effect of the commands. The CM needs to maintain enough internal state to validate all commands; this will most likely entail maintaining the full contents of the input tables. Note that the input tables can be computed as the I of all input deltas ever applied. Once all commands producing a change have been accepted, the environment can apply the produced input change atomically, and obtain from the circuit the corresponding output changes.

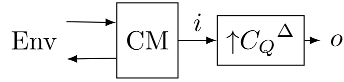

Part III
Appendixes

A Z-transform and stream convolutions

Definition A.1. The **Z-transform** of a stream, as traditionally defined in signal processing, is a function that associates with any stream over a group a formal power series in the indeterminate z: $\mathcal{Z} : \mathcal{S}_A \to A[\![z]\!]$ defined as $\mathcal{Z}(s) \stackrel{\text{def}}{=} \sum_{t \geq 0} s[t]z^{-t}$.

For example, the Z-transform of the id stream is the power series $0 + z^{-1} + z^{-2} + z^{-3} + \ldots$.

Definition A.2. Let $(R, +, \cdot, 0, 1)$ be a commutative ring. The **Cauchy product** (also called discrete convolution) of two streams $* : \mathcal{S}_R \times \mathcal{S}_R \to \mathcal{S}_R$ is defined as:

$$(a * b)[t] \;=\; \sum_{i=0}^{t} a[i] \cdot b[t-i]$$

For example, the convolution of the id stream with itself is the stream $id * id$ containing the sequence of values $0, (0 \cdot 1 + 1 \cdot 0), (0 \cdot 2 + 1 \cdot 1 + 2 \cdot 0), \ldots = 0, 0, 1, 4, 8, \ldots$.

Proposition A.3. The structure $(\mathcal{S}_R, +, *, 0, 1)$ is also a commutative ring. This ring is isomorphic to the ring of formal power series in one indeterminate $R[\![z]\!]$ with coefficients from R.

Sometimes it is more convenient to use the formal power series notation. Notice that we have $z^{-1}(s) = z^{-1} * s$, justifying the traditional notation for the delay operator z^{-1}. It follows that the differentiation of a stream s is $D(s) = (1 - z^{-1}) * s$.

Moreover, the equation that defines the integration of a stream s, $\xi = z^{-1}(\xi) + s$ is equivalent to $\xi = z^{-1} * \xi + s$ and then to $(1 - z^{-1}) * \xi = s$. Since $1 - z^{-1}$ has multiplicative inverse

$$(1 - z^{-1})^{-1} \;=\; 1 + z^{-1} + z^{-2} + z^{-3} + \cdots$$

we can express the integration operator by $I(s) = (1 - z^{-1})^{-1} * s$. Theorem 3.30 now follows by algebraic manipulations in the ring of formal power series. Similarly for the time-invariance and linearity properties of D and I. Even causality can be treated algebraically, once we note that, like addition, convolution is causal.

Observation As shown above, there are two proof styles for equations over streams: one is (usually) by induction over the time dimension, and the other one is equational, by operating with polynomials over z. The theory of digital signal processing posits that these two proof styles give the same results.

References

[1] sqllogictest. https://www.sqlite.org/sqllogictest/doc/trunk/about.wiki. Retrieved December 2022.

[2] The aurora project. http://cs.brown.edu/research/aurora/, December 2004.

[3] Martín Abadi, Frank McSherry, and Gordon Plotkin. Foundations of differential dataflow. In *Foundations of Software Science and Computation Structures (FoSSaCS)*, 2015. URL: http://homepages.inf.ed.ac.uk/gdp/publications/differentialweb.pdf.

[4] Supun Abeysinghe, Qiyang He, and Tiark Rompf. Efficient incrementialization of correlated nested aggregate queries using relative partial aggregate indexes (rpai). In *ACM SIGMOD International conference on Management of data (SIGMOD)*, page 136–149, 2022. URL: https://doi.org/10.1145/3514221.3517889.

[5] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. URL: http://webdam.inria.fr/Alice/.

[6] Yanif Ahmad and Christoph Koch. DBToaster: A SQL compiler for high-performance delta processing in main-memory databases. *Proc. VLDB Endow.*, 2(2):1566–1569, August 2009. URL: https://doi.org/10.14778/1687553.1687592.

[7] Mario Alvarez-Picallo, Alex Eyers-Taylor, Michael Peyton Jones, and C.-H. Luke Ong. Fixing incremental computation. In *European Symposium on Programming Languages and Systems (ESOP)*, pages 525–552, 2019. URL: https://link.springer.com/chapter/10.1007/978-3-030-17184-1_19.

[8] Krzysztof R. Apt and Jean-Marc Pugin. Maintenance of stratified databases viewed as a belief revision system. In Moshe Y. Vardi, editor, *ACM SIGMOD International conference on Management of data (SIGMOD)*, pages 136–145, San Diego, California, March 23-25 1987. doi:10.1145/28659.28674.

[9] Arvind Arasu, Shivnath Babu, and Jennifer Widom. An abstract semantics and concrete language for continuous queries over streams and relations. Technical Report 2002-57, Stanford InfoLab, 2002. URL: http://ilpubs.stanford.edu:8090/563/.

[10] W. Baker and A. Newton. The maximal VHDL subset with a cycle-level abstraction. In *Proceedings of the Conference on European Design Automation (DATE)*, pages 470–475, 1996.

[11] Edmon Begoli, Jesús Camacho-Rodríguez, Julian Hyde, Michael J. Mior, and Daniel Lemire. Apache Calcite: A foundational framework for optimized query processing over heterogeneous data sources. In *International Conference on Management of Data (IDMD)*, page 221–230, 2018. URL: https://doi.org/10.1145/3183713.3190662.

[12] Angela Bonifati, Stefania Dumbrava, and Emilio Jesús Gallego Arias. Certified graph view maintenance with regular Datalog. *Theory and Practice of Logic Programming*, 18(3-4):372–389, 2018. doi:10.1017/S1471068418000224.

[13] Stefano Ceri and Jennifer Widom. Deriving production rules for incremental view maintenance. In Guy M. Lohman, Amílcar Sernadas, and Rafael Camps, editors, *International Conference of Very Large Data Bases (VLDB)*, pages 577–589, Barcelona, Spain, 1991. URL: http://www.vldb.org/conf/1991/P577.PDF.

[14] Surajit Chaudhuri, Ravi Krishnamurthy, Spyros Potamianos, and Kyuseok Shim. Optimizing queries with materialized views. In *International Conference on Data Engineering (ICDE)*, page 190–200, 1995.

[15] Rada Chirkova and Jun Yang. *Materialized Views*. Now Publishers Inc., Hanover, MA, USA, 2012.

[16] Zaheer Chothia, John Liagouris, Frank McSherry, and Timothy Roscoe. Explaining outputs in modern data analytics. *Proc. VLDB Endow.*, 9(12):1137–1148, August 2016. URL: https://doi.org/10.14778/2994509.2994530.

[17] Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The Lean theorem prover. In *International Conference on Automated Deduction (CADE-25)*, Berlin, Germany, 2015.

[18] Hasanat M. Dewan, David Ohsie, Salvatore J. Stolfo, Ouri Wolfson, and Sushil Da Silva. Incremental database rule processing in PARADISER. *J. Intell. Inf. Syst.*, 1(2):177–209, 1992. doi:10.1007/BF00962282.

[19] J. Nathan Foster, Ravi Konuru, Jerome Simeon, and Lionel Villard. An algebraic approach to XQuery view maintenance. In *ACM SIGPLAN Workshop on Programming Languages Technologies for XML*, San Francisco, CA, January 9 2008.

[20] Peter Gammie. Synchronous digital circuits as functional programs. *ACM Comput. Surv.*, 46(2), November 2013. URL: https://doi.org/10.1145/2543581.2543588.

[21] Goetz Graefe. Encapsulation of parallelism in the Volcano query processing system. In *SIGMOD '90: Proceedings of the 1990 ACM SIGMOD international conference on Management of data*, pages 102–111, New York, NY, USA, 1990. ACM Press. doi:http://doi.acm.org/10.1145/93597.98720.

[22] Sergio Greco and Cristian Molinaro. Datalog and logic databases. *Synthesis Lectures on Data Management*, 7(2):1–169, 2015. URL: https://doi.org/10.2200/S00648ED1V01Y201505DTM041.

[23] Todd J Green, Zachary G Ives, and Val Tannen. Reconcilable differences. *Theory of Computing Systems*, 49(2):460–488, 2011. URL: https://web.cs.ucdavis.edu/~green/papers/tocs11_differences.pdf.

[24] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *Symposium on Principles of Database Systems (PODS)*, page 31–40, 2007. URL: https://doi.org/10.1145/1265530.1265535.

[25] Timothy Griffin and Leonid Libkin. Incremental maintenance of views with duplicates. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, page 328–339, 1995. URL: https://doi.org/10.1145/223784.223849.

[26] Ashish Gupta, Inderpal Singh Mumick, et al. Maintenance of materialized views: Problems, techniques, and applications. *IEEE Data Eng. Bull.*, 18(2):3–18, 1995.

[27] Ashish Gupta, Inderpal Singh Mumick, and V. S. Subrahmanian. Maintaining views incrementally. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, pages 157–166, Washington, DC, May 26-28 1993. ACM Press. doi:10.1145/170035.170066.

[28] John V. Harrison and Suzanne W. Dietrich. Maintenance of materialized views in a deductive database: An update propagation approach. In Kotagiri Ramamohanarao, James Harland, and Guozhu Dong, editors, *Workshop on Deductive Databases*, volume CITRI/TR-92-65 of *Technical Report*, pages 56–65, Washington, D.C., November 14 1992. Department of Computer Science, University of Melbourne.

[29] Muhammad Idris, Martin Ugarte, and Stijn Vansummeren. The dynamic yannakakis algorithm: Compact and efficient query processing under updates. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, page 1259–1274, 2017. URL: https://doi.org/10.1145/3035918.3064027.

[30] Muhammad Idris, Martín Ugarte, Stijn Vansummeren, Hannes Voigt, and Wolfgang Lehner. Conjunctive queries with inequalities under updates. *Proc. VLDB Endow.*, 11(7):733–745, mar 2018. URL: https://doi.org/10.14778/3192965.3192966.

[31] Muhammad Idris, Martín Ugarte, Stijn Vansummeren, Hannes Voigt, and Wolfgang Lehner. Efficient query processing for dynamically changing datasets. *SIGMOD Rec.*, 48(1):33–40, November 2019. URL: https://doi.org/10.1145/3371316.3371325.

[32] Hojjat Jafarpour, Rohan Desai, and Damian Guy. KSQL: Streaming SQL engine for Apache Kafka. In *International Conference on Extending Database Technology (EDBT)*, pages 524–533, Lisbon, Portugal, March 26-29 2019. URL: http://openproceedings.org/2019/conf/edbt/EDBT19_paper_329.pdf.

[33] Steven Dexter Johnson. *Synthesis of Digital Designs from Recursion Equations.* PhD thesis, Indiana University, May 1983. https://help.luddy.indiana.edu/techreports/TRNNN.cgi?trnum=TR141.

[34] Gilles Kahn. The semantics of a simple language for parallel programming. In *IFIP Congress on Information Processing*, 1974. URL: http://www1.cs.columbia.edu/~sedwards/papers/kahn1974semantics.pdf.

[35] Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. Maintaining triangle queries under updates. *ACM Trans. Database Syst.*, 45(3), aug 2020. URL: https://doi.org/10.1145/3396375.

[36] Christoph Koch. Incremental query evaluation in a ring of databases. In *Symposium on Principles of Database Systems (PODS)*, page 87–98, 2010. URL: https://doi.org/10.1145/1807085.1807100.

[37] Christoph Koch, Daniel Lupei, and Val Tannen. Incremental view maintenance for collection programming. In *Symposium on Principles of Database Systems (PODS)*, page 75–90, 2016. URL: https://doi.org/10.1145/2902251.2902286.

[38] Jakub Kotowski, François Bry, and Simon Brodt. Reasoning as axioms change - incremental view maintenance reconsidered. In *Web Reasoning and Rule Systems RR*, volume 6902 of *Lecture Notes in Computer Science*, pages 139–154, Galway, Ireland, August 29-30 2011. Springer. doi:10.1007/978-3-642-23580-1_11.

[39] Edward A. Lee. Multidimensional streams rooted in dataflow. In *IFIP Working Conference on Architectures and Compilation Techniques for Fine and Medium Grain Parallelism*, Orlando, FL, January 20-22 1993. URL: https://ptolemy.berkeley.edu/publications/papers/93/mdsdf/.

[40] Edward A. Lee and Thomas M. Parks. Dataflow process networks. *Proceedings of the IEEE*, pages 773–801, May 1995. URL: https://ptolemy.berkeley.edu/publications/papers/95/processNets/.

[41] Charles E. Leiserson and James B. Saxe. Retiming synchronous circuitry. *Algorithmica*, 6:5–35, 1991. URL: http://www.cs.columbia.

edu/~CS6861/handouts/leiserson-algorithmica-88.pdf, doi:https://doi.org/10.1007/BF01759032.

[42] James J. Lu, Guido Moerkotte, Joachim Schü, and V. S. Subrahmanian. Efficient maintenance of materialized mediated views. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, pages 340–351, San Jose, California, May 22-25 1995. doi:10.1145/223784.223850.

[43] Konstantinos Mamouras. Semantic foundations for deterministic dataflow and stream processing. In Peter Müller, editor, *European Symposium on Programming*, pages 394–427, Dublin, Ireland, April 25–30 2020.

[44] Frank McSherry, Andrea Lattuada, Malte Schwarzkopf, and Timothy Roscoe. Shared arrangements: Practical inter-query sharing for streaming dataflows. *Proc. VLDB Endow.*, 13(10):1793–1806, June 2020. URL: https://doi.org/10.14778/3401960.3401974.

[45] Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. Differential dataflow. In *Conference on Innovative Data Systems Research (CIDR)*, Asilomar, CA, January 6–9 2013. URL: http://cidrdb.org/cidr2013/Papers/CIDR13_Paper111.pdf.

[46] Boris Motik, Yavor Nenov, Robert Piro, and Ian Horrocks. Maintenance of Datalog materialisations revisited. *Artif. Intell.*, 269:76–136, 2019. URL: https://doi.org/10.1016/j.artint.2018.12.004.

[47] Boris Motik, Yavor Nenov, Robert Edgar Felix Piro, and Ian Horrocks. Incremental update of Datalog¡ materialisation: the backward/forward algorithm. In *Conference on Artificial Intelligence (AAAI)*, pages 1560–1568, Austin, Texas, January 25-30 2015. AAAI Press. URL: http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9660.

[48] Derek G. Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and Martín Abadi. Naiad: A timely dataflow system. In *ACM Symposium on Operating Systems Principles (SOSP)*, page 439–455, 2013. doi:10.1145/2517349.2522738.

[49] Milos Nikolic and Dan Olteanu. Incremental view maintenance with triple lock factorization benefits. In *International Conference on Management of Data (ICMD)*, page 365–380, 2018. URL: https://doi.org/10.1145/3183713.3183758.

[50] L. R. Rabiner and B. Gold, editors. *Theory and Application of Digital Signal Processing*. Prentice-Hall, 1975.

[51] Leonid Ryzhyk and Mihai Budiu. Differential datalog. In *Datalog 2.0*, Philadelphia, PA, June 4-5 2019. URL: http://budiu.info/work/ddlog.pdf.

[52] Martin Staudt and Matthias Jarke. Incremental maintenance of externally materialized views. In *International Conference of Very Large Data Bases (VLDB)*, pages 75–86, Mumbai (Bombay), India, September 3-6 1996. URL: http://www.vldb.org/conf/1996/P075.PDF.

[53] Qichen Wang and Ke Yi. Maintaining acyclic foreign-key joins under updates. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, page 1225–1239, 2020. URL: https://doi.org/10.1145/3318464.3380586.

[54] Ouri Wolfson, Hasanat M. Dewan, Salvatore J. Stolfo, and Yechiam Yemini. Incremental evaluation of rules and its relationship to parallelism. In *ACM SIGMOD International conference on Management of data (SIGMOD)*, pages 78–87, Denver, Colorado, May 29-31 1991. ACM Press. doi:10.1145/115790.115799.