


# 1 Verified Time-Aware Stream Processing

2 Rafael Castro Gonçalves Silva ✉ 

3 University of Copenhagen

4 Dmitriy Traytel ✉ 

5 University of Copenhagen

## 6 — Abstract —

7 Stream processing frameworks provide programming abstractions that allow their users to express  
8 the desired time-dependent data analysis. The frameworks organize the computation as a directed  
9 graph of interconnected operators that perform event-wise transformations. We tackle the correctness  
10 question for programs expressed in this way. To this end, we model (possibly stateful) operators and  
11 define their composition, model data streams with time-stamps and watermarks, define reusable,  
12 modular operators, and prove their correctness in the Isabelle/HOL proof assistant, taking advantage  
13 of its advanced coinductive methods infrastructure. We demonstrate the usefulness of our model by  
14 verifying stream processing algorithms computing incremental histograms and relational joins.

15 **2012 ACM Subject Classification** Security and privacy → Logic and verification; Computing  
16 methodologies → Distributed algorithms; Software and its engineering → Data flow languages

17 **Keywords and phrases** stream processing, verification, coinduction, Isabelle/HOL

18 **Digital Object Identifier** 10.4230/LIPIcs...

19 **Funding** This work is supported by a Novo Nordisk Fonden start package grant (NNF20OC0063462).

## 20 **1** Introduction

21 Data stream processing frameworks, such as Apache Flink [10], Apache Samza [18], Apache  
22 Spark [38], Google Cloud Dataflow [4], and Timely Dataflow [29, 31, 32] are widely used tools  
23 for the low-latency parallel processing of large quantities of data arriving at a high velocity and  
24 possibly out of order at a software system. For many developers, these frameworks take the  
25 role of both the programming language and the operating system as they provide high-level  
26 abstractions that transparently hide much of the complexity of parallel programming.

27 Most stream processing frameworks follow the dataflow paradigm in which *operators*  
28 transform streams of time-stamped events and the entire program, called *dataflow*, consists  
29 of a graph of operators that together orchestrate the desired computation. Dataflows are  
30 programmed at this *logical* level of abstraction, yet they are deployed and executed in parallel  
31 on multiple *workers*, which can be separate processes or even separate machines.

32 Despite their popularity, stream processing frameworks are notorious for correctness  
33 issues [17, 21–23, 36, 37]. The predominant countermeasure to errors in the stream processing  
34 (and the wider distributed systems) community is testing [5, 17, 37]. For certain kinds of  
35 errors, random testing is effective [26]. Others are hard to detect [21] and require specialized  
36 approaches [23]. But no matter how elaborated, testing remains fundamentally incomplete.

37 How to formally reason about and prove the correctness of programs implemented in these  
38 frameworks? To answer this question, in this paper we focus on the logical level where the  
39 question becomes: “How to reason about operators and their composition?” More specifically,  
40 we use the Isabelle/HOL proof assistant (Section 2) to model operators as coinductive  
41 input/output state machines, the event streams they manipulate as lazy lists (i.e., potentially  
42 infinite sequences), and to define composition of operators by corecursion (Section 3).

43 Stream processing frameworks manage out-of-order data streams through *logical time-*  
44 *stamps* and *completeness metadata*. Logical time-stamps do not necessarily represent chrono-



45 logical time; rather, they can encompass application-specific semantics, such as grouping of  
 46 data items based on different versions or origins. The completeness metadata informs operat-  
 47 ors when a group is complete, i.e., no further data for this group will be arriving in the future,  
 48 which may prompt the operator to send final results for this group and clean up its state.

49 We instantiate our core model with event streams that incorporate partially ordered  
 50 time-stamps and use watermarks [6] as the completeness metadata (Section 4). We define fun-  
 51 damental correctness notions for such event streams, which we call *monotonicity* (watermark  
 52 give correct predictions about future events) and *productivity* (all data items are eventually  
 53 followed by corresponding watermarks). We then identify and prove the correctness of basic,  
 54 reusable operators for time-dependent batching of data and incremental computations. To  
 55 prove correctness of an operator, we develop a blueprint consisting of four properties: (1)  
 56 operator-specific soundness, which describes the desired properties of an operator’s output, (2)  
 57 operator-specific completeness, which characterizes what must be contained in an operator’s  
 58 output, (3) monotonicity preservation, and (4) productivity preservation. The properties  
 59 (3) and (4) are crucial for reasoning about composed operators.

60 To validate our model we conduct two case studies (Section 5). First, we define an  
 61 incremental histogram computation as the composition of our above reusable operators. We  
 62 show how the building blocks’ correctness properties compose to yield the overall correctness  
 63 property. In addition, we define an optimized incremental histogram computation directly as a  
 64 single operator and establish the equivalence between the modular and the optimized variant.  
 65 Second, we use partially ordered time-stamps, which allow us to model operators with multiple  
 66 inputs as operators with a single input using the disjoint sum order on time-stamps. We use  
 67 this insight to define and verify a streaming relational join operator using our building blocks.

68 Our formalization is not tied to any particular stream processing framework, although  
 69 many of the design choices were inspired by Timely Dataflow [31]. Our work is best understood  
 70 as a model for specifying and proving correctness of abstract algorithms in the dataflow  
 71 paradigm. For example, we represent data streams as lazy lists instead of actual data channels,  
 72 e.g., the network layer is not part of our model. Nonetheless, our operators are executable, in  
 73 that we can run small examples before engaging with the formal proofs and test our conjectures.  
 74 Our formalization is publicly available [1]. In summary, we make the following contributions:

- 75 ■ We develop a framework for formally specifying time-aware stream processing programs,  
 76 which supports the modular composition of building block operators, provides executable  
 77 specifications, and is validated in two case studies.
- 78 ■ We distill a four-ingredient proof methodology for operator correctness: soundness,  
 79 completeness, monotonicity preservation, and productivity preservation.
- 80 ■ Our work constitutes a case study in advanced coinductive methods in a proof assistant  
 81 including mixing friendly corecursion [8] and monadic recursion [19], coinduction up to  
 82 friends [8] and other generalizations of (co)inductive proofs.

83 **Related Work** Our long term objective is to formalize and verify a simplified variant of the  
 84 Timely Dataflow [32] stream processing framework in Isabelle. The present work arose from  
 85 the need to specify what Timely Dataflow programs are expected to do, without delving into  
 86 framework-specific details. This background explains some of our design choices and justifies  
 87 why we could not reuse an existing model from the literature. As a distinguishing feature,  
 88 Timely Dataflow uses partially ordered time-stamps to support cyclic dataflows. We inherit  
 89 this design choice, and while we have not yet formalized cyclic operator graphs, our framework  
 90 is prepared to do so. Moreover, we benefit from partially ordered time-stamps when modeling  
 91 operators with multiple inputs. This design choice also requires us to generalize the standard

notion of watermarks [6], which traditional stream processing frameworks with totally ordered time-stamps use [4, 10, 18]. After receiving a traditional watermark  $wm$ , the operator is aware that no event with time-stamp  $t < wm$  will be received in the future. This notion of monotonicity [6] inspired our generalization to partial orders shown in Subsection 4.1.

Next, we discuss related pen-and-paper formalizations of stream processing. Kahn networks [16] are often considered as the origin of dataflow programming but lack a notion of time. Hancock et al. [13] develop *stream processors* as a programmable representation of continuous stream transformations. Continuity in this context means that a finite prefix of the output is determined by a finite prefix of the input. Stream processors are a popular example of mixing recursion and corecursion and have been formalized in Agda [3, 11] and Isabelle/HOL [9]. They lack a notion of time and the fact that they include the continuity proof in their very structure complicates the definition of time-aware operators. In contrast, our approach decouples the definition of operators from proofs of their desirable properties. Mamouras [27] organizes abstractions of streams as monoids and operators as deterministic stream transducers. We follow a similar approach, albeit being less general (representing streams as lazy lists) and equating transducers with their semantics using a shallow embedding via a coinductive datatype. Mamouras' framework has a notion of time, but restricts time-stamps to be natural numbers.

In the context of proof assistants, the Isabelle formalization of the FOCUS system [34, 35] models time-aware, but in-order possibly infinite streams and operators transforming them. Once again, it uses natural numbers as time-stamps. Hinrichsen et al. [15] model the seminal MapReduce framework [12] in Coq. MapReduce processes finite data batches in parallel and without a notion of time. Lochbihler and Hölzl [24] define in Isabelle recursive functions on lazy lists as fixpoints in chain-complete partial orders [19]. We reuse their library of lazy lists.

## 2 Codatatypes, Coinduction, and Corecursion

We introduce Isabelle's infrastructure for defining and reasoning about coinductive datatypes, short codatatypes [7]. The following codatatype `enat` of extended natural numbers has the infinity element `ESucc (ESucc (...))`. If `datatype` was used instead of `codatatype` below, the infinity element would not exist and the type would be isomorphic to `nat`.

```
codatatype enat = EZero | ESucc enat
```

To represent finite and infinite data streams in a single type, we work with the codatatype of lazy lists, i.e., finite or infinite sequences of elements:

```
codatatype (lset: 'a) llist = Inull: LNil ([#] | LCons (lhd: 'a) (ltl: 'a llist) (infixr ## 65)
for map: lmap where ltl LNil = LNil
```

This command introduces the type `'a llist` of potentially infinite sequences along with:

- constructors `LNil :: 'a llist`, written `[#]`, and `LCons :: 'a ⇒ 'a llist ⇒ 'a llist`, written `##`,
- selectors `lhd :: 'a llist ⇒ 'a` (where `lhd [#] = undefined`, i.e., a fixed, unspecified value of type `'a`) and `ltl :: 'a llist ⇒ 'a llist` (where `ltl [#] = LNil` as specified),
- discriminator `Inull :: 'a llist ⇒ bool` returning true iff its argument is `[#]`,
- the functorial action `lmap :: ('a ⇒ 'b) ⇒ 'a llist ⇒ 'b llist`, and
- the function `lset :: 'a llist ⇒ 'a set` extracting the set of elements contained in the lazy list.

The `codatatype` command provides a coinduction principle to reason about equality. Proving that two lazy lists are equal reduces to exhibiting a *bisimulation*  $R$  that relates them. A bisimulation is a relation that is stable under discriminators and destructors. Formally, we obtain the following coinduction principle for `llist`:

## XX:4 Verified Time-Aware Stream Processing

$$R \text{ } lxs \text{ } lys \implies (\bigwedge lxs \text{ } lys . R \text{ } lxs \text{ } lys \implies \text{Inull } lxs \longleftrightarrow \text{Inull } lys \wedge \neg \text{Inull } lxs \implies \neg \text{Inull } lys \implies \text{Lhd } lxs = \text{Lhd } lys \wedge R \text{ } (\text{Itl } lxs) \text{ } (\text{Itl } lys)) \implies lxs = lys \quad (1)$$

Here,  $\bigwedge$  is universal quantification. An induction principle for `lset` membership is also derived:

$$x \in \text{lset } lxs \implies (\bigwedge x \text{ } lxs . P \text{ } x \text{ } (x \text{ } \#\# \text{ } lxs)) \implies (\bigwedge x \text{ } lxs \text{ } y . y \in \text{lset } lxs \implies P \text{ } y \text{ } lxs \implies P \text{ } y \text{ } (x \text{ } \#\# \text{ } lxs)) \implies P \text{ } x \text{ } lxs \quad (2)$$

All functions in Isabelle/HOL must be total. This is guaranteed for terminating recursive functions on datatypes. In contrast, corecursive functions produce elements of codatatypes are total when they are productive, i.e., they always eventually output a codatatype constructor. Corecursive functions are defined with the `corec` command, where the corecursive calls must be guarded by a codatatype constructor, which ensures productivity. For example:

```
corec lshift :: 'a list  $\Rightarrow$  'a llist  $\Rightarrow$  'a llist (infixr @@@ 65) where
  xs @@@ lys = case xs of []  $\Rightarrow$  lys | x # xs'  $\Rightarrow$  x ## (xs' @@@ lys)
```

where `[]` and infix `#` are the ordinary constructors of the (finite) list datatype.

Many total corecursive functions involve other operations than the guarding constructor in the context surrounding the corecursive call. An example is `lconcat` characterized by:

```
lconcat lxs = case lxs of [#]  $\Rightarrow$  [#] | xs ## lxs'  $\Rightarrow$  xs @@@ lconcat lxs'
```

The `corec` command permits functions to be used freely in definitions if they have been previously registered as (and proved to be) *friendly* [8]. A function is friendly when it preserves productivity in a rather rigid way: it may consume at most one constructor to produce one constructor. Isabelle can automatically prove that `lshift` is friendly by supplying the (*friend*) option to the above `lshift` definition (and adjusting it mildly). With `lshift` registered as a friend, we can define `lconcat`, although not using the above equation, which lacks a guarding constructor and hides the complication that all lists in `xss` may be empty (in which case `lconcat xss = [#]`). We refer for details to our formalization, which derives the above equation as a lemma.

The coinduction principle (1) is inconvenient for corecursive functions that use friends in their corecursive call contexts. The `corec` command automatically derives the *coinduction up to congruence* principle, which replaces  $R \text{ } (\text{Itl } xs) \text{ } (\text{Itl } ys)$  by `cong R (Itl xs) (Itl ys)` in (1), for the congruence closure `cong` with respect to currently known friends, and thus allows the bisimulation proof to descend under these friends.

Coinductive predicates are definitions as Knaster–Tarski greatest fixpoints on the predicate lattice [33]. They resemble inductive predicates, but allow the introduction rules to be applied infinitely often. Consider the prefix relation on lazy lists:

```
coinductive lprefix :: 'a llist  $\Rightarrow$  'a llist  $\Rightarrow$  bool where
  lprefix [#] lys
  | lprefix lxs lys  $\implies$  lprefix (x ## lxs) (x ## lys)
```

Any finite prefix will eventually reach the base case; the coinductive bit is relevant here to ensure that the definition is reflexive, also for infinite lazy lists. Coinductive predicates are accompanied by a coinduction principle. To prove that a coinductive predicate holds for some arguments, we must exhibit a *witness relation*  $R$  that relates them. For `lprefix`,  $R$  is a witness relation if for any  $R$ -related lazy lists either the left argument is empty, or both arguments are non-empty, their heads equal, and their tails related by either  $R$  or `lprefix`. Formally:

$$R \text{ } lxs \text{ } lys \implies (\bigwedge lxs \text{ } lys . R \text{ } lxs \text{ } lys \implies \text{Inull } lxs \vee (\neg \text{Inull } lxs \wedge \neg \text{Inull } lys \implies \text{Lhd } lxs = \text{Lhd } lys \wedge (R \text{ } (\text{Itl } lxs) \text{ } (\text{Itl } lys) \vee \text{lprefix } (\text{Itl } lxs) \text{ } (\text{Itl } lys)))) \implies \text{lprefix } lxs \text{ } lys \quad (3)$$

### 3 Lazy Lists Processors

We introduce operators as a codatatype, give them a semantics as lazy list transformers, define sequential composition, and prove that its correctness with respect to the given semantics.

#### 3.1 Operators on Lazy Lists

We define the codatatype  $(i, o)$  *op* of operators with a single constructor parametrized by a user defined function, which we call logic, and a lazy list:

```
codatatype (i, o) op = Logic (apply: (i ⇒ (i, o) op × o list)) (exit: o llist)
```

The type variable  $i$  represents the operator's inputs, and  $o$  its outputs. The selector **apply** applies the operator's logic to an input, yielding a pair consisting of a new (evolved) operator and a list of outputs. The new operator can again be applied to another input, etc. The selector **exit** handles the situation when there are no more inputs. In this case, the operator may produce final outputs. An operator can be seen as an infinitely deep tree branching over  $i$  (which also may be infinite): **apply** follows a branch, depending on the provided input, to the next node.

It may seem that our operators are stateless, but corecursive functions may have additional parameters representing the state. Our first example operator has state: it buffers data until a buffer limit  $n$  is reached, and then it performs a bulk operation with the function  $f$ . At the end of the input stream any remaining data in the buffer is output as well:

```
corec bulk_op where bulk_op f n buf = Logic (λe. if length (e # buf) < n
  then (bulk_op f n (e # buf), []) else (bulk_op f n [], f (e # buf))) (llist_of (f buf))
```

We define next how an operator can be applied to an entire, possibly infinite lazy list. We start with the auxiliary function **produce<sub>1</sub>**, which iterates **apply** until the operator produces a non-empty output  $x \# xs$ , in which case **Some (Inl (op', x, xs, lxs'))** is returned. The operator may fail to produce a non-empty output, either because the input lazy list ends (**Some (Inr op')**), or nothing in the infinite lazy list causes any output to be produced (**None**).

```
partial_function (option) produce1 where
  produce1 op lxs = (case lxs of [#] ⇒ Some (Inr op)
    | h ## lxs' ⇒ (case apply op h of (op', []) ⇒ produce1 op' lxs'
      | (op', x # xs) ⇒ Some (Inl (op', x, xs, lxs'))))
```

The **partial\_function** command [19] defines functions by recursion in the option monad (with non-termination represented by **None**). The command automatically derives an induction principle for the terminating executions of **produce<sub>1</sub>**, which we rewrite as:

$$\begin{aligned} \text{produce}_1 \text{ op } lxs = \text{Some } y \implies & \\ (\bigwedge \text{op } h \text{ lxs}' \text{ op}' \text{ zs. } \text{apply } \text{op } h = (\text{op}', []) \implies Q \text{ op}' \text{ lxs}' \text{ zs} \implies Q \text{ op } (h \text{ ## } \text{lxs}') \text{ zs}) \implies & \\ (\bigwedge \text{op } h \text{ x } \text{xs} \text{ lxs}' \text{ op}'. \text{produce}_1 \text{ op } (h \text{ ## } \text{lxs}) = \text{Some } (\text{Inl } (\text{op}', \text{x}, \text{xs}, \text{lxs}')) \implies & \quad (4) \\ \text{apply } \text{op } h = (\text{op}', \text{x} \# \text{xs}) \implies Q \text{ op } (h \text{ ## } \text{lxs}') (\text{Inl } (\text{op}', \text{x}, \text{xs}, \text{lxs}')) \implies & \\ (\bigwedge \text{op}. Q \text{ op } [\#] (\text{Inr } \text{op})) \implies Q \text{ op } lxs \text{ y} & \end{aligned}$$

We define **produce** to corecursively repeat this processes and concatenate the outputs. The following definition requires **@@** to be (registered as) friendly.

```
corec produce where produce op lxs = (case produce1 op lxs of None ⇒ [#]
  | Some (Inr op') ⇒ exit op'
  | Some (Inl (op', x, xs, lxs')) ⇒ x ## (xs @@ produce op' lxs'))
```

### 229 3.2 Sequential Composition

230 The (sequential) composition of two operators means giving the first operator's output as input  
 231 to the second operator. As our operators output lists, this requires folding the second operator  
 232 over the output of the first, for which we use the following finite list variant of `produce`:  
 233

234 **definition** `fproduce`  $op\ xs = \text{fold } (\lambda e\ (op,\ out)).$   
 235 `let`  $(op',\ out') = \text{apply } op\ e\ \text{in } (op',\ out\ @\ out')$   $x\ (op,\ [])$   
 236

237 The infix `@` is the list append function. The composition operator is:

238 **corec** `comp_op` **where** `comp_op`  $op_1\ op_2 = \text{Logic } (\lambda ev.$   
 239 `let`  $(op_1',\ out) = \text{apply } op_1\ ev;$   $(op_2',\ out') = \text{fproduce } op_2\ out$   
 240 `in`  $(\text{comp\_op } op_1'\ op_2',\ out')$   $(\text{produce } op_2\ (\text{exit } op_1))$   
 241 242

243 The composed exit value iteratively applies  $op_2$  to  $op_1$ 's exit. The correctness of `comp_op`  
 244 states that its production corresponds to the functional composition of its arguments' produc-  
 245 tions: `produce`  $(\text{comp\_op } op_1\ op_2)\ lxs = \text{produce } op_2\ (\text{produce } op_1\ lxs)$ . Unfortunately, this  
 246 equality does hold unconditionally: if  $op_1$  enters an *unproductive* state (`produce`<sub>1</sub> returns  
 247 `None`), then the equality reduces to  $\#[\#] = \text{exit } op_2$ . Therefore, we add an assumption stating  
 248 that either  $op_1$  and all its evolutions are eventually productive or  $op_2$ 's and all its evolutions'  
 249 exit is  $\#[\#]$ . To this end, we introduce another operator that *skips* the first  $n$  outputs of an  
 250 operator, but does not alter that operator's evolution:

251 **corec** `skip_op` **where** `skip_op`  $op\ n = \text{Logic } (\lambda ev.$  `let`  $(op',\ out) = \text{apply } op\ ev\ \text{in}$   
 252 `if` `length`  $out < n$  `then`  $(\text{skip\_op } op'\ (n - \text{length } out),\ [])$   
 253 `else`  $(op',\ \text{drop } n\ out)$   $(\text{ldropn } n\ (\text{exit } op))$   
 254 255

256 Using `ldropn`, which drops the first  $n$  elements of a lazy list, the correctness of `skip_op` is:

$$257 \text{produce } (\text{skip\_op } op\ n)\ lxs = \text{ldropn } n\ (\text{produce } op\ lxs) \quad (5)$$

258 We prove (5) by coinduction up to congruence, instantiating the bisimulation candidate  $R$  with

$$259 \lambda l\ r. \exists op\ n\ lxs. l = \text{produce } (\text{skip\_op } op\ n)\ lxs \wedge r = \text{ldropn } n\ (\text{produce } op\ lxs)$$

260 (Isabelle's coinduction proof method automatically constructs this canonical instantiation.)  
 261 Proving that  $R$  is a bisimulation up to congruence yields three subgoals about arbitrary  $R$ -  
 262 related lazy lists  $l$  and  $r$ : (i) `lnull`  $l \longleftrightarrow \text{lnull } r$ , (ii) `lhd`  $l = \text{lhd } r$ , and (iii) `cong`  $R\ (\text{!}l\ l)\ (\text{!}l\ r)$ ,  
 263 where (ii) and (iii) may additionally assume that  $l$  and  $r$  do not satisfy `lnull`. All subgoals  
 264 are proved with lemmas about `produce`<sub>1</sub> that are proved by its induction principle (4).

265 Correctness of `comp_op` is then expressed as

$$266 (\forall n. \text{produce}_1\ (\text{skip\_op } op_1\ n)\ lxs \neq \text{None}) \vee (\forall xs. \text{exit}\ (\text{fst}\ (\text{fproduce } op_2\ xs)) = \#[\#]) \implies$$

$$\text{produce } (\text{comp\_op } op_1\ op_2)\ lxs = \text{produce } op_2\ (\text{produce } op_1\ lxs) \quad (6)$$

267 The proof of (6) also proceeds by coinduction up to congruence. Again, different lemmas  
 268 about `produce`<sub>1</sub> must be proved, one relevant example being:

$$269 \text{produce}_1\ (\text{comp\_op } op_1\ op_2)\ lxs = \text{None} \implies$$

$$\text{produce } op_1\ lxs = ys\ @@\ lys \implies \exists op_2'. \text{fproduce } op_2\ ys = (op_2',\ []) \quad (7)$$

270 We proceed by an induction on  $ys$ . Its induction step case is:


$$271 (\bigwedge op_1\ op_2\ lxs\ lys. \text{produce}_1\ (\text{comp\_op } op_1\ op_2)\ lxs = \text{None} \implies$$

$$\text{produce } op_1\ lxs = ys\ @@\ lys \implies \exists op_2'. \text{fproduce } op_2\ ys = (op_2',\ [])) \implies$$

$$\text{produce}_1\ (\text{comp\_op } op_1\ op_2)\ lxs = \text{None} \implies \text{produce } op_1\ lxs = (y\ \#\ ys)\ @@\ lys \implies$$

$$\exists op_2'. \text{fproduce } op_2\ (y\ \#\ ys)\ [] = (op_2',\ [])$$





■ **Figure 1** An example stream and its collections (ordered by their time-stamps)

272 There is a mismatch between  $ys$  in the induction hypothesis and  $y \# ys$  in the assumption  
 273 about `produce`  $op_1$   $xs$ , so that we cannot use  $op_1$  to instantiate the induction hypothesis.  
 274 Our solution is to instantiate  $op_1$  in the induction hypothesis with `skip_op`  $op_1$   $1$ . Using  
 275 `skip_op` as a means of generalization is a common pattern in our development. It allows us  
 276 to formulate statements about arbitrary positions in the output of `produce`. Namely, when  
 277 `produce1` (`skip_op`  $op$   $n$ )  $xs$  returns `Some (Inl (op, x, xs, xs))`, the value  $x$  is  $op$ 's  $n$ th output.

## 278 4 Time-Aware Operators

279 We model time-aware data streams. Events are either time-stamped data items or watermarks:

280 **datatype** ( $'t::order, 'd$ )  $event = DT$  (tmp:  $'t$ ) (data:  $'d$ ) | **WM** (tmp:  $'t$ )  
 281  
 282

283 Here, time-stamps  $'t$  are assumed to form a partial order via Isabelle's `order` type class.  
 284 A (time-aware data) stream is a lazy list of type ( $'t, 'd$ )  $event$   $llist$ . A collection  $C_t$  is the  
 285 multiset containing all the data items with the time-stamp  $t$  from a time-aware data stream.  
 286 A watermark **WM**  $wm$  *completes* a time-stamp **DT**  $t$   $d$  if  $wm \geq t$ . A complete collection is a  
 287 collection from the lazy list that is completed by some watermark from that lazy list. Figure 1a  
 288 shows a prefix of an out-of-order stream, named  $stream_1$ , and Figure 1b shows the associated  
 289 set of collections for all time-stamps occurring in the prefix. We arrange the set of collections in  
 290 a Hasse diagram, partially ordered by the collections' time-stamps. That is in this example, we  
 291 assume that  $t_0 < t_1 < t_4$ ,  $t_0 < t_2 < t_5$ , and  $t_0 < t_3 < t_5$  are the inequalities that hold for  $'t$ .

### 292 4.1 Monotonicity and Productivity

293 Watermarks denote that some collection is completed, allowing operators to safely send out-  
 294 puts and clean up their state. If the stream violates the property that after a watermark  $wm$   
 295 no data with time-stamp  $t \leq wm$  will arrive, then operators may no longer work as expected.  
 296 In other words, the time-stamps on data items must respect a monotonicity property, which  
 297 intuitively means “never moving backwards” in relation to the already seen watermarks [6].

298 It is not enough to prohibit time-stamps below the last received watermark due to partially-  
 299 ordered time-stamps: e.g., given the order from Figure 1b, if **WM**  $t_4$  was received just after  
 300 **WM**  $t_5$  tracking only one of these incomparable time-stamps would necessarily lose information.  
 301 Instead, we track a set  $W$  of received watermarks and ensure that for all  $wm'$  in  $W$ , every  
 302 future time-stamp is not less or equal than  $wm'$ . Timely Dataflow [32] uses a similar (dual)  
 303 notion of frontiers, which are antichains of time-stamps (i.e., sets of incomparable time-stamps)  
 304 that may be encountered in the future. We use plain sets rather than antichains for simplicity.

305 As streams may be infinite, the monotonicity property is defined coinductively:

306 **coinductive mono where**  $mono$  [ $\#$ ]  $W$   
 307 |  $mono$   $xs$  ( $\{wm\} \cup W$ )  $\implies mono$  (**WM**  $wm$   $\#\#$   $xs$ )  $W$   
 308 | ( $\forall wm \in W. \neg wm \geq t$ )  $\implies mono$   $xs$   $W \implies mono$  (**DT**  $t$   $d$   $\#\#$   $xs$ )  $W$   
 309  
 310

311 The prefix of  $stream_1$  shown in Figure 1a satisfies  $mono$   $stream_1$   $\{\}$ .

```

inductive mono_cong for  $R$  where  $R\ lxs\ W \implies \text{mono\_cong}\ R\ lxs\ W$ 
| mono_cong  $R\ lxs\ (WMs\ xs\ \cup\ W) \implies \text{mono}\ (\text{llist\_of}\ xs)\ W \implies \text{mono\_cong}\ R\ (xs\ @\@ \ lxs)\ W$ 
inductive prod_cong for  $R$  where  $R\ lxs \implies \text{prod\_cong}\ R\ lxs$ 
| prod_cong  $R\ lxs \implies (\forall n < \text{length}\ xs. \forall t\ d. \text{nth}\ xs\ n = \text{DT}\ t\ d \longrightarrow$ 
   $(\exists wm \geq t. \text{WM}\ wm \in \text{lset}\ (\text{drop}\ (\text{Suc}\ n)\ xs\ @\@ \ lxs))) \implies \text{prod\_cong}\ R\ (xs\ @\@ \ lxs)$ 
 $R\ lxs\ W \implies (\bigwedge lxs\ W. R\ lxs\ W \implies \text{lnull}\ lxs \vee (\exists wm\ lxs'. lxs = \text{WM}\ wm\ \#\#\ lxs' \wedge$ 
   $(\text{mono\_cong}\ R\ lxs'\ (\{wm\} \cup W) \vee \text{mono}\ lxs'\ (\{wm\} \cup W))) \vee$ 
   $(\exists t\ d\ lxs'. lxs = \text{DT}\ t\ d\ \#\#\ lxs' \wedge (\forall wm \in W. \neg t \leq wm) \wedge$ 
   $(\text{mono\_cong}\ R\ lxs'\ W \vee \text{mono}\ lxs'\ W))) \implies \text{mono}\ lxs\ W$ 
 $R\ lxs \implies (\bigwedge lxs. R\ lxs \implies \text{lfinite}\ lxs \vee$ 
   $(\exists wm\ lxs'. lxs = \text{WM}\ wm\ \#\#\ lxs' \wedge \neg \text{lfinite}\ lxs' \wedge (\text{prod\_cong}\ R\ lxs' \vee \text{prod}\ lxs')) \vee$ 
   $(\exists t\ d\ lxs'. lxs = \text{DT}\ t\ d\ \#\#\ lxs' \wedge \neg \text{lfinite}\ lxs \wedge (\exists wm \geq t. \text{WM}\ wm \in \text{lset}\ lxs') \wedge$ 
   $(\text{prod\_cong}\ R\ lxs' \vee \text{prod}\ lxs')))) \implies \text{prod}\ lxs$ 

```

■ **Figure 2** Coinduction up to lshift-congruence principles for mono and prod

Streams must be productive, i.e., always eventually make it possible to produce outputs. We require that for all seen data items  $\text{DT}\ t\ d$ , there eventually is a watermark  $\text{WM}\ wm$  with  $wm \geq t$ . This is crucial for operators' completeness: for every consumed time-stamp there must be an associated output. For finite streams ( $\text{lfinite}$ ), the requirement can be ignored as the stream's end represents the computation's end. We define productivity coinductively:

```

coinductive prod where  $\text{lfinite}\ lxs \implies \text{prod}\ lxs$ 
|  $\neg \text{lfinite}\ lxs \implies (\exists u \geq t. \text{WM}\ u \in \text{lset}\ lxs) \implies \text{prod}\ lxs \implies \text{prod}\ (\text{DT}\ t\ d\ \#\#\ lxs)$ 
|  $\text{prod}\ lxs \implies \text{prod}\ (\text{WM}\ t\ \#\#\ lxs)$ 

```

In Figure 1a,  $\text{stream}_1$  may respect this property if it is finite, or if there is a watermark completing  $t_4$  at some later position and the rest of the lazy list respects  $\text{prod}$ .

The automatically derived coinduction principles for both coinductive predicates are inconvenient for proving some of the properties in Subsection 4.3 due to the occurrence of  $\text{lshift}$  in  $\text{produce}$ . The `corec` command [8] derives the coinduction up to congruence principle only for lazy list equality. Therefore, we manually prove coinduction up to congruence for both predicates using their respective regular coinduction principles. To this end, we define inductively in Figure 2 custom congruence closure (under  $\text{lshift}$ ) predicates for both  $\text{mono}$  and  $\text{prod}$ . The congruence closure predicates are parameterized by a relation  $R$  and allow us to descend under  $\text{lshift}$  provided that the prefixed list  $xs$  preserves the respective properties. In Figure 2,  $\text{WMs}$  returns the set of all  $wm$  values from all watermarks  $\text{WM}\ wm$  in a list,  $\text{nth}$  returns the  $n$ th element of a list and  $\text{drop}$  removes the first  $n$  elements of a list. The respective coinduction up to congruence principles (8) and (9) are also shown in Figure 2.

## 4.2 Building Blocks

We introduce two operators, which we call building blocks, for batching and incremental computations. The batching operator creates *batches*, which are lists of pairs of time-stamps and data, consisting of completed time-stamps that were not part of a previous batch. We introduce an auxiliary function for creating batches given a list of watermarks:

```

fun batches where batches [] tds = ([], tds)
| batches (wm # wms) tds = let (bs, tds') = batches wms [(t, d) ← tds. ¬ t ≤ wm]
in (DT wm [(t, d) ← tds. t ≤ wm] # bs, tds')

```



produce (batch\_op [(t<sub>0</sub>, a)]) stream<sub>1</sub> =  

DT t <sub>1</sub> [(t <sub>0</sub> , a), (t <sub>0</sub> , a), (t <sub>1</sub> , c)]	WM t <sub>1</sub>	DT t <sub>5</sub> [(t <sub>2</sub> , b), (t <sub>3</sub> , a), (t <sub>5</sub> , a), (t <sub>5</sub> , c)]	WM t <sub>5</sub>	WM t <sub>2</sub>	⋯
--	-------------------	--	-------------------	-------------------	---

■ **Figure 3** A prefix of produce (batch\_op [(t<sub>0</sub>, a)]) stream<sub>1</sub>

produce (incr\_op []) (produce (batch\_op [(t<sub>0</sub>, a)]) stream<sub>1</sub>) =  

DT t <sub>0</sub> [(t <sub>0</sub> , a), (t <sub>0</sub> , a), (t <sub>1</sub> , c)]	DT t <sub>1</sub> [(t <sub>0</sub> , a), (t <sub>0</sub> , a), (t <sub>1</sub> , c)]	WM t <sub>1</sub>	⋯
DT t <sub>5</sub> [(t <sub>0</sub> , a), (t <sub>0</sub> , a), (t <sub>1</sub> , c), (t <sub>2</sub> , b), (t <sub>3</sub> , a), (t <sub>5</sub> , a), (t <sub>5</sub> , c)]		WM t <sub>5</sub>	WM t <sub>2</sub>

■ **Figure 4** A prefix of produce (incr\_op []) (produce (batch\_op [(t<sub>0</sub>, a)]) stream<sub>1</sub>)

345 For every watermark in its first argument, `batches` computes a batch consisting of pairs from  
 346 its second argument that have a time-stamp below that watermark. The batch operator is:

```
347 corec batch_op where
348   batch_op tds = Logic (λev. case ev of DT t d ⇒ (batch_op (tds @ [(t, d)]), [])
349   | WM wm ⇒ let (out, tds') = batches [wm] tds
350   in (batch_op tds', [x ← out. data x ≠ []] @ [WM wm]))
351   (let wms = maximal_antichain_list (map fst tds) in llist_of (fst (batches wms tds)))
352
353
```

354 Here, `maximal_antichain_list` computes a maximal antichain, i.e., a list of distinct, maximal,  
 355 and incomparable elements from its argument. The operator's buffer `tds` accumulates received  
 356 DT items and only outputs them (and removes them from the buffer) when a watermark with  
 357 a greater or equal time-stamp arrives. The end of the input stream is interpreted by the  
 358 operator's exit as a final watermark which completes all time-stamps in the buffer. Figure 3  
 359 shows the output of `batch_op [(t0, a)]` after consuming the stream prefix from Figure 1a.  
 360 The watermark `WM t1` causes the output of a batch containing the collections  $C_{t_0}$  and  $C_{t_1}$ ,  
 361 whereas `WM t5` outputs all other newly completed collections. If the stream ends at that  
 362 point, then there will be a final output with `DT t4 [(t4, d)]` caused by the operator's exit.

363 The incremental computation operator appends batches, creating *accumulated batches*:

```
364 corec incr_op where
365   incr_op abs = Logic (λev. case ev of WM wm ⇒ (incr_op abs, [WM wm])
366   | DT wm b ⇒ (incr_op (abs @ b), map (λt. DT t (abs @ b)) (remdups (map fst b)))) [#]
367
368
```

369 Here, `remdups` removes duplicates in a list. The `incr_op` operator receives batches and  
 370 produces batches that are incrementally accumulated. For each time-stamp of the received  
 371 batch, it outputs a new batch that is concatenated with all previously received batches.  
 372 The batches are produced without inspecting the time-stamps. Hence, they may include  
 373 incomparable time-stamps. This design choice simplifies the soundness property. Figure 4  
 374 shows the output stream of `incr_op []` after consuming the stream prefix from Figure 3.

### 375 4.3 Correctness

376 We identify four *correctness* properties of time-aware operators: soundness, completeness,  
 377 preservation of monotonicity and preservation of productivity. Soundness means that each  
 378 operator output meets the operator-specific specification. Completeness means that each  
 379 input is somehow represented in the operator's output under operator-specific conditions.

380 For `batch_op` and `incr_op`, soundness and completeness statements rely on the auxiliary  
 381 definitions in Figure 5. There, `list_of` converts a finite lazy list into a list (and returns  
 382 `undefined` for infinite lazy lists), `lfilter` filters lazy lists, `mset` transforms a list into a multiset,  
 383 and `ltakeWhile` takes the elements from a lazy list while the given predicate holds.

```

definition map_filter :: ('a ⇒ 'b option) ⇒ 'a list ⇒ 'b list where
  map_filter f xs = map (λx . the (f x)) (filter (λx. f x ≠ None) xs)
definition DTs lxs t = map_filter (λev. case ev of DT t d ⇒ Some d | WM wm ⇒ None)
  (list_of (lfilter (λev. case ev of DT t' d ⇒ t' = t | WM wm ⇒ False) lxs))
definition lcoll lxs t = mset (DTs lxs t)
definition coll xs t = mset (map snd (filter (λ(t', d). t' = t) xs))
definition set_t xs = set (map fst xs)
definition ts lxs t = {t'. ∃d'. DT t' d ∈ lset lxs ∧ t' ≤ t}
definition ws lxs wm = {wm'. WM wm' ∈ lset (ltakeWhile ((≠) (WM wm)) lxs)}
definition incompletes lxs = (let xs = filter (λev. case ev of
  DT t d ⇒ ¬ (∃wm ≥ t. WM wm ∈ lset lxs) | WM _ ⇒ False) (list_of lxs)
  in maximal_antichain_list (map tmp xs))
definition ws2 lxs wm = set (takeWhile ((≠) wm) (incompletes lxs))
definition batch_ts lxs wm = if WM wm ∈ lset lxs
  then {t' ∈ ts lxs wm. ¬ (∃wm' ≥ t'. wm' ∈ ws lxs wm)}
  else {t' ∈ ts lxs wm. ¬ (∃wm' ≥ t'. wm' ∈ ws2 lxs wm ∨ WM wm' ∈ lset lxs)}

```

■ **Figure 5** Auxiliary definitions for the soundness and completeness of the building blocks

384 Collections are formalized for both *list* and *llist* types, respectively, by `coll` and `lcoll`.  
 385 As exemplified by Figure 3, the time-stamps of separate batches are disjoint. Under the  
 386 `mono` assumption the `list_of` used in DTs is only applied to finite lazy lists because `lfilter` is  
 387 guaranteed to only find finitely many elements. The functions `ws` and `ws2` represent the  
 388 time-stamps of outputs that possibly happened prior to receiving `wm`. Here, `ws2` is the special  
 389 case for `exit`, in which there are no actual watermarks completing these time-stamps, so we  
 390 take the maximal antichain of time-stamps that do not have completing watermarks instead.

### 391 4.3.1 Correctness of `batch_op`.

392 For `batch_op`, the intuitive meaning of soundness is: for an output DT `wm b`, the batch `b`  
 393 is formed by all completed collections with time-stamps  $\leq wm$  that were not yet output. We  
 394 characterize the time-stamps belonging to `b` by `batch_ts` in Figure 5. Thereby, we distinguish  
 395 whether the time-stamps are below a watermark from the stream. The full soundness lemma is:

$$\begin{aligned}
 \text{mono } lxs \ W \implies \text{DT } wm \ b \in \text{lset } (\text{produce } (\text{batch\_op } tds) \ lxs) \implies \\
 (\forall t \in \text{set\_t } b. \text{coll } b \ t = \text{lcoll } lxs \ t + \text{coll } tds \ t) \wedge \\
 \text{set\_t } b = \text{batch\_ts } ((\text{map } (\lambda (t,d). \text{DT } t \ d) \ tds) \ \text{@@} \ lxs) \ wm
 \end{aligned} \tag{10}$$

397 The soundness proof proceeds by induction on `lset` (2). However, another mismatch with the  
 398 induction hypothesis forces us to generalize the operator using `skip_op`. To finish the proof,  
 399 we use several lemmas about `produce1` that follow by the respective induction principle (4).

400 Completeness of `batch_op` states that every input DT `t d` must be present as `(t, d)` in  
 401 some output batch. From `batch_op`'s soundness (10) we already know that the time-stamps  
 402 in the output batches form completed collections. Hence we only need to show that the  
 403 time-stamp is present in some batch. Because of the `exit` call, we know that every time-stamp  
 404 eventually will be in some batch, even if there is no watermark completing it. We formally  
 405 state the completeness of `batch_op` as follows:

$$\begin{aligned}
 \text{prod } lxs \implies \text{DT } t \ d \in \text{lset } lxs \vee t \in \text{set\_t } tds \implies \\
 \text{lfinite } lxs \vee (\forall t' \in \text{set\_t } tds. \exists wm \geq t'. \text{WM } wm \in \text{lset } lxs) \implies \\
 \exists wm \ b. \text{DT } wm \ b \in \text{lset } (\text{produce } (\text{batch\_op } tds) \ lxs) \wedge t \in \text{set\_t } b
 \end{aligned} \tag{11}$$

406

407 The third assumption extends the concept of productivity to also include the buffer *tds*. Completeness follows by induction on `lset` (2). In the base case, the time-stamp is already in the buffer, or enters the buffer after one call of `produce1`. We prove an (omitted) auxiliary lemma that  
 408  
 409 if a time-stamp is in the buffer, it will eventually be part of the output, regardless whether there  
 410 is a completing watermark in the stream or the stream is finite. The induction step follows by  
 411 showing that *t* eventually enters the buffer, which reduces the problem to the auxiliary lemma.  
 412

413 The preservation of `mono` by `batch_op` requires an additional assumption about the buffer, which may be output in `batch_op`'s exit, and follows by coinduction up to congruence (8).  
 414 The preservation `prod` follows similarly by coinduction up to congruence (9).  
 415

$$\text{mono } lxs \ W \implies (\forall t \in \text{set\_t } tds. \forall wm \in W. \neg t \leq wm) \implies \text{mono } (\text{produce } (\text{batch\_op } tds) \ lxs) \ W \quad (12)$$

$$\text{prod } lxs \implies \text{prod } (\text{produce } (\text{batch\_op } tds) \ lxs) \quad (13)$$

### 4.3.2 Correctness of `incr_op`.

417 The soundness for `incr_op` describes its outputs: an accumulated batch *ab* which is the concatenation of the buffer with accumulated batches from some prefix of the incoming data stream. The notion of accumulated batches is given by:

```
421 fun ltake_DT where
422   ltake_DT (Suc n) (WM _ ## lxs) = ltake_DT n lxs
423   | ltake_DT (Suc n) (DT t b ## lxs) = (t, b) # ltake_DT n lxs
424   | ltake_DT _ _ = []
425 definition acc_batches n lxs ≡ concat (map snd (ltake_DT n lxs))
426
427
```

428 The soundness of `incr_op` relates the produced batch `DT t ab` with the accumulated consumed prefix, and asserts that *t* originated from the accumulated batches:  
 429

$$\text{DT } t \ ab \in \text{lset } (\text{produce } (\text{incr\_op } abs) \ lxs) \implies \exists n. ab = abs @ \text{acc\_batches } n \ lxs \wedge t \in \text{set\_t } (\text{acc\_batches } n \ lxs) \quad (14)$$

430 As for `batch_op`, soundness requires a generalization with `skip_op` and proceeds by induction on `produce1` (4) after unfolding `produce`.

431 The completeness of `incr_op` says that every *t* in every received batch will result in an accumulated batch, and it follows by induction on `lset` (2):  
 432  
 433  
 434

$$\text{DT } wm \ b \in \text{lset } lxs \implies t \in \text{set\_t } b \implies \exists ab. \text{DT } t \ ab \in \text{produce } (\text{incr\_op } abs) \ lxs \quad (15)$$

435 The preservation of `mono` and `prod` by `incr_op` also follow, respectively, from their coinduction up to congruence, lemmas (8) and (9), but each require an additional assumption:  
 436  
 437

$$\text{mono } lxs \ W \implies (\forall n \ wm \ b. n < \text{llength } lxs \longrightarrow \text{Inth } lxs \ n = \text{DT } wm \ b \longrightarrow (\forall t' \in \text{set\_t } b. t' \leq wm \wedge (\forall wm' \in W \cup \text{vimage } WM \ (\text{lset } (\text{ltake } n \ lxs)). \neg wm' \geq t')))) \implies \text{mono } (\text{produce } (\text{incr\_op } abs) \ lxs) \ W \quad (16)$$

$$\text{prod } lxs \implies (\forall n \ wm \ b. n < \text{llength } lxs \longrightarrow \text{Inth } lxs \ n = \text{DT } wm \ b \longrightarrow (\exists m > n. \text{Inth } lxs \ m = WM \ wm) \wedge (\forall t' \in \text{set\_t } b. t' \leq wm) \wedge \neg b \neq []) \implies \text{prod } (\text{produce } (\text{incr\_op } abs) \ lxs) \quad (17)$$

438 Here, `llength` is the length function returning an `enat`, `Inth` returns the *n*th element of a lazy list, `ltake` takes the first *n* `enat` elements from a lazy list, and `vimage f B = {x. f x ∈ B}`  
 439 is the inverse image of a function. The additional assumptions of the preservation lemmas extend the concepts of `mono` and `prod` to the time-stamps inside of the incoming batches.  
 440  
 441  
 442

#### 4.4 Compositional Reasoning

The operators `incr_op` and `batch_op` are composable:

**definition** `incr_batch_op tds abs = comp_op (batch_op tds) (incr_op abs)`

We prove the four correctness properties for the composed operator by *compositional reasoning*: the proved properties of `batch_op` prove the assumptions of `incr_op`. All proofs start by unfolding `incr_batch_op` and rewriting with `comp_op`'s correctness (6). We note that `exit` of `incr_op` always remains `[#]`, which validates the assumption of (6). For preservation of `mono`, we inherit the additional assumption of (12) which extends monotonicity to the buffer `tds`. Formally:

$$\text{mono } lxs \ W \implies (\forall t \in \text{set\_t } tds. \forall wm \in W. \neg t \leq wm) \implies \text{mono } (\text{produce } (\text{incr\_batch\_op } tds \ abs) \ lxs) \ W \quad (18)$$

Similarly, preservation of productivity is:

$$\text{prod } lxs \implies \text{prod } (\text{produce } (\text{incr\_batch\_op } tds \ abs) \ lxs) \quad (19)$$

Both lemmas are proved by backwards reasoning with the preservation of monotonicity (16) and productivity (17) for `incr_op`, followed by the preservation of monotonicity (12) and productivity (17) for `batch_op`, and using `batch_op`'s soundness (10) to discharge the additional assumptions of the preservation lemmas for `incr_op`.

Each accumulated batch produced by `incr_batch_op` has some prefix of concatenate DT productions of (`batch_op`). This fact follows as a corollary of the soundness of `batch_op` (10) which is then combined with the soundness of `incr_op` (14) to derive soundness of `incr_batch_op`:

$$\begin{aligned} \text{mono } lxs \ W \implies & \text{DT } t \ ab \in \text{lset } (\text{produce } (\text{incr\_batch\_op } tds \ abs) \ lxs) \implies \\ & \exists n. \ ab = \text{abs} \ @ \ \text{acc\_batches } n \ (\text{produce } (\text{batch\_op } tds) \ lxs) \ \wedge \\ & \{t' \in \text{set\_t} \ (\text{acc\_batches } n \ (\text{produce } (\text{batch\_op } tds) \ lxs)). \ t' \leq t\} = \\ & \text{ts } lxs \ t \cup \{t' \in \text{set\_t } tds. \ t' \leq t\} \ \wedge \\ & (\forall t' \leq t. \ \text{coll } (\text{acc\_batches } n \ (\text{produce } (\text{batch\_op } tds) \ lxs)) \ t' = \text{lcoll } lxs \ t' + \text{coll } tds \ t') \end{aligned} \quad (20)$$

Completeness of `incr_batch_op` follows from that of `batch_op` (11) and `incr_op` (15):

$$\text{prod } lxs \implies \text{DT } t \ d \in \text{lset } lxs \implies \exists b. \ \text{DT } t \ b \in \text{lset } (\text{produce } (\text{incr\_batch\_op } [] \ [])) \ lxs \quad (21)$$


## 5 Case Study

We introduce two other operators using the building blocks, and showcase how their correctness properties follow compositionally from those of the basic components.

### 5.1 Histogram

A histogram counts a collection's elements. We compute incremental histograms, which for a time-stamp  $t$  count not only elements belonging to the collection  $C_t$ , but also elements of collections  $C_{t'}$  for all smaller time-stamps  $t' < t$ . We write  $H_t$  for an incremental histogram that counts all elements with time-stamps  $\leq t$ . We represent (incremental) histograms as multisets.

Partially ordered time-stamps make incremental histograms somewhat counterintuitive. We recall the earlier example of time-stamp collections in Figure 6a. The time-stamp collections  $C_{t_0}$ ,  $C_{t_2}$ , and  $C_{t_3}$  result in the incremental histograms  $H_{t_0}$ ,  $H_{t_2}$  and  $H_{t_3}$  in Figure 6b. The question is “What precisely should  $H_{t_5}$  count and how to compute it?” since it has two immediate predecessor incremental histograms  $H_{t_2}$  and  $H_{t_3}$  that moreover share  $H_{t_0}$  as their predecessor. Here, we use as our specification the variant that counts every included collection only once:  $H_{t_5} = C_{t_0} + C_{t_2} + C_{t_3} + C_{t_5}$ . A similar example is



■ **Figure 6** Computing incremental histograms in partial order

481 present<sup>1</sup> in the Rust implementation of Differential Dataflow [2, 28, 30], a library for differ-  
 482 ential computation built on top of Timely Dataflow. There, pairs of numbers are used as  
 483 the partially ordered time-stamps, whereas we work with an arbitrary partial order. (In  
 484 our formalization, we additionally verify the more complex specification variant in which  
 485  $H_{t_5} = H_{t_2} + H_{t_3} + C_{t_5} = C_{t_0} + C_{t_0} + C_{t_2} + C_{t_3} + C_{t_5}$ .)

486 An implementation of our specification using `incr_batch_op` collects the data with time-  
 487 stamps  $\leq t$  from the accumulated batches at time-stamp  $t$ . It needs to compute the histogram  
 488 for each newly accumulated batch. Without further assumptions on the partial order, any  
 489 efficient implementation of this incremental histogram specification must keep track of *all*  
 490 previous incremental histograms, not only maximal ones. For example, in Figure 6a a new  
 491 collection with time-stamp  $t_6 > t_0$  could show up with  $t_6$  being incomparable to all other  
 492 time-stamps. To compute  $H_{t_6}$ , we must thus have access to  $H_{t_0}$  (or  $C_{t_0}$ ).

493 Our specification sums all collections with time-stamps less or equal than a given  $t$ . We use  
 494 the generic set summation  $\text{sum} :: 'a \text{ set} \Rightarrow ('a \Rightarrow 'b) \Rightarrow 'b :: \text{comm\_monoid\_add}$  for commut-  
 495 ative monoids (here multisets). We also introduce summation for finite accumulated batches.

496 **definition**  $S\_lcoll \text{ } lxs \text{ } t = \text{sum } (ts \text{ } lxs \text{ } t) (lcoll \text{ } lxs)$

497 **definition**  $S\_coll \text{ } t \text{ } xs = \text{sum } \{t' \in \text{set } xs. t' \leq t\} (coll \text{ } xs)$

500 To transform the accumulated batches from `incr_batch_op`, we use the following time-aware  
 501 map operator and define the incremental histogram operator by composition:

502 **corec** `map_op` **where** `map_op`  $f = \text{Logic } (\lambda ev. \text{case } ev \text{ of}$

503 `WM`  $wm \Rightarrow (\text{map\_op } f, [\text{WM } wm]) \mid \text{DT } t \text{ } d \Rightarrow (\text{map\_op } f, [\text{DT } t (f \text{ } t \text{ } d)])$ )

504 **definition** `ihist_op`  $tds \text{ } abs = \text{comp\_op } (\text{incr\_batch\_op } tds \text{ } abs) (\text{map\_op } S\_coll)$

507 From the preservation of `mono` (18) and `prod` (19) for `incr_batch_op` and `map_op` (omitted),  
 508 both preservation properties are also derived for `ihist_op`. The soundness and completeness  
 509 of the incremental histogram, simplified to use empty buffers, are:

510 **mono**  $lxs \text{ } W \Longrightarrow \text{DT } t \text{ } H \in \text{lset } (\text{produce } (\text{ihist\_op } [] \text{ } []) \text{ } lxs) \Longrightarrow H = S\_lcoll \text{ } lxs \text{ } t$  (22)

511 **prod**  $lxs \Longrightarrow \text{DT } t \text{ } d \in \text{lset } lxs \Longrightarrow \exists H. \text{DT } t \text{ } H \in \text{lset } (\text{produce } (\text{ihist\_op } [] \text{ } []) \text{ } lxs)$  (23)

512 They follow by `incr_batch_op`'s soundness (20) and completeness (21), respectively.

513 When time-stamps are natural numbers, an efficient implementation of incremental histo-  
 514 grams will compute each  $H_t$  by summing the *last stored incremental histogram*  $H_{t-1}$  with the  
 515 newest completed time-stamp collection  $C_t$  after seeing `WM`  $t$ . This is efficient because the last  
 516 histogram has been already computed. In contrast, our generic incremental histogram is ineffi-  
 cient: it has a buffer that is never cleaned, and it recomputes the entire histogram for each accu-

<sup>1</sup> <https://github.com/TimelyDataflow/differential-dataflow/blob/250e9c2ad2e9/examples/multitemporal.rs>

```

corec ihist_op' where ihist_op' H buf = Logic (λev. case ev of
  DT (t:::linorder) d ⇒ (ihist_op' H (buf @ [(t, d)], []))
  | WM wm ⇒ if ∃(t, d) ∈ set buf . t ≤ wm
  then let out = [(t, _) ← buf. t ≤ wm]; buf' = [(t, _) ← buf. t > wm];
  tss = remdups ((map fst out));
  Hs = map (λt. DT t (H + (mset (map snd [(t' _) ← out. t' ≤ t]))) tss in
    (ihist_op' (H + (mset (map snd out))) buf', Hs @ [WM wm])
  else (ihist_op' H buf, []))
  (lilst_of (map (λt. DT t (H + mset (map snd [(t' _) ← buf. t' ≤ t])))
    (remdups (map fst buf))))

```

■ **Figure 7** Efficient incremental histogram operator

produce (ihist_op' [] [])	DT 1 a	WM 1	DT 0 a	WM 0	=	DT 1 {a}	WM 1	DT 0 {a, a}	WM 0
produce (ihist_op [] [])	DT 1 a	WM 1	DT 0 a	WM 0	=	DT 1 {a}	WM 1	DT 0 {a}	WM 0

■ **Figure 8** Difference between ihist\_op' and ihist\_op

517 mulated batch from scratch. We optimize the incremental histogram computation, in Figure 7,  
 518 for a totally ordered (*linorder* typeclass) time-stamp type. The optimized operator has two  
 519 buffers: the first one keeps the last histogram (which makes sense for a total order); the second  
 520 one keeps the tuples not output yet. When new histograms are output, the operator replaces  
 521 the last stored histogram with the most recent one. The two implementations are equivalent  
 522 when applied to streams with totally ordered time-stamps respecting *mono*. Figure 8 demon-  
 523 strates that the implementations differ when the input is not monotone. The *ihist\_op'* operator  
 524 counts all previously seen events which it stores in its state *H*, whereas *ihist\_op* is aware of  
 525 all time-stamps because it uses accumulated batches including all previous time-stamps.

526 We show that the two operators produce the same results for monotone inputs:

mono <i>lxs</i> <i>W</i> ⇒ (∀t ∈ set_t <i>abs</i> . ∃wm ∈ <i>W</i> . t ≤ wm) ⇒	(24)
(∀t ∈ set_t <i>tds</i> . ∀wm ∈ <i>W</i> . t > wm) ⇒	
produce (ihist_op' (mset (map snd <i>abs</i> )) <i>tds</i> ) <i>lxs</i> = produce (ihist_op <i>tds</i> <i>abs</i> ) <i>lxs</i>	

527  
 528 Instead of directly proving this equality, we define a coinductive predicate between  
 529 operators that holds if they produce the same outputs when applied to the same monotone  
 530 input. The relation, similarly to *mono*, uses a set *W* of seen watermarks, and coinductively  
 531 checks that the time-stamp of the next event is not smaller or equal than any time-stamp in *W*:

<b>definition</b> eq_op_lifted <i>ev</i> <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> <i>R</i> = exit <i>op</i> <sub>1</sub> = exit <i>op</i> <sub>2</sub> ∧ (case <i>ev</i> of
DT <i>t</i> <i>d</i> ⇒ (∀wm ∈ <i>W</i> . ¬ t ≤ wm) ⇒ rel_prod ( <i>R</i> <i>W</i> ) (=) (apply <i>op</i> <sub>1</sub> <i>ev</i> ) (apply <i>op</i> <sub>2</sub> <i>ev</i> )
WM <i>wm</i> ⇒ rel_prod ( <i>R</i> ({ <i>wm</i> } ∪ <i>W</i> )) (=) (apply <i>op</i> <sub>1</sub> <i>ev</i> ) (apply <i>op</i> <sub>2</sub> <i>ev</i> ))
<b>coinductive</b> eq_op <b>where</b> (∀ <i>ev</i> . eq_op_lifted <i>ev</i> <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> eq_op) ⇒ eq_op <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub>

538 Here, *rel\_prod* relates the elements of two pairs point-wise using the given relations. The  
 539 coinduction principle of *eq\_op* is:

(∧ <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> <i>ev</i> . <i>R</i> <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> ⇒ eq_op_lifted <i>ev</i> <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> <i>R</i> ) ⇒	(25)
<i>R</i> <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> ⇒ eq_op <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub>	

541 The relation *eq\_op* is a reasonable equivalence of operators: two operators related by  
 542 *eq\_op* generate the same outputs via *produce* when applied to the same monotone input:

mono <i>lxs</i> <i>W</i> ⇒ eq_op <i>W</i> <i>op</i> <sub>1</sub> <i>op</i> <sub>2</sub> ⇒ produce <i>op</i> <sub>1</sub> <i>lxs</i> = produce <i>op</i> <sub>2</sub> <i>lxs</i>	(26)
---	------



544 This fact follows by the coinduction up to congruence (1). Finally, we prove that our two  
 545 histogram implementations are related using the coinduction principle for `eq_op` (25):

$$\forall t \in \text{set\_t } abs. \exists wm \in W. t \leq wm \implies \forall t \in \text{set\_t } tds. \forall wm \in W. t > wm \implies \text{eq\_op } W (\text{ihist\_op}' (\text{mset } (\text{map } \text{snd } abs)) tds) (\text{ihist\_op } tds abs) \quad (27)$$

546  
 547 Combining this fact with the soundness of `eq_op` (26) implies our desired property (24).

548 Our incremental histogram operators are executable using the lazy evaluation library  
 549 by Lochbihler and Stoop [25]. In particular, our definitions do not use any non-executable  
 550 constants such as quantifiers over infinite domains. Furthermore, `produce` has code equations,  
 551 which enables code generation [14]. Naturally, when working with infinite input lazy lists,  
 552 one can only observe finite prefixes of the output in finite time, e.g., by using `ltake_DT`.

## 553 5.2 Join

554 Our second case study illustrates how multiple disjoint time-aware data streams can be repre-  
 555 sented by a single one, which allows us to define operators with multiple inputs. The operator,  
 556 defined in Figure 9, joins data items using a given  $join :: 'a \Rightarrow 'a \Rightarrow 'b \text{ option}$  function, and  
 557 uses the sum type  $'t + 't$  as its time-stamp type: an input with time-stamp `lnl t1` is inter-  
 558 preted as coming the first data stream, whereas time-stamp `lnr t2` originates from the second  
 559 data stream. We use the point-wise partial order on sum types as the order for our definition.  
 560 That is, `lnl t1` is not related to any `lnr t2`, and `lnl t1 ≤ lnl t2` iff  $t_1 \leq t_2$  (and similarly for `lnr`).

561 Once again, we use the `incr_batch_op` operator, so that we only need to define how to com-  
 562 pute the results from the accumulated batches. In Figure 9, we apply `join_list` to each accumu-  
 563 lated batch, which combines `lnl` and `lnr` tuples when they share the same time-stamp, and apply  
 564 the provided  $join$  function to all possible combinations. Each joined result is output as an in-  
 565 dividual data item by the flatten operator `flatten_op`, which transforms batches into individual  
 566 data items assigning the time-stamp  $t$  from `DT t b` to each element in the batch  $b$ . Finally, we  
 567 remove the sum type from the output time-stamps using the union operator `union_op`, which  
 568 outputs watermarks as soon they are identified as *producible*, meaning that a greater or equal  
 569 watermark on the opposite side was already consumed. This identification is relevant for the  
 570 preservation of `mono`, as it is only safe to output a watermark  $wm$  after the data on both sides  
 571 is completed for it. The operator `union_op` has two buffers: the first keeps track of watermarks  
 572 not identified as producible yet; the second denotes the maximal antichain of seen watermarks.

573 Figure 10 shows the four correctness properties of `join_op` with empty initial buffers to sim-  
 574 plify the involved assumptions. The first two (soundness and completeness) express, respect-  
 575 ively, that each joined item is the result of joining elements from different sides of the input  
 576 data stream, and that every pair of items that can be joined will generate an item in the output.  
 577 Both proofs follow similarly to the soundness (22) and completeness (23) of `ihist_op`. Preserva-  
 578 tion of `prod` for `join_op` assumes all incoming watermarks to be producible. Under this assump-  
 579 tion, all watermarks will be output. Our proofs rely on four correctness properties of `union_op`  
 580 (omitted), that resemble in terms of additional assumptions closely the properties of `join_op`.

## 581 6 Conclusion

582 We used the Isabelle/HOL proof assistant to model (possibly non-terminating) time-aware  
 583 data stream processing. We identified two essential building block operators for batching and  
 584 incremental computations, which we reused in two case studies. Moreover, we established  
 585 the correctness for our operators, combining induction and coinduction. The correctness of  
 586 composed operators follows compositionally from the building block operators' correctness.

```

corec flatten_op where
  flatten_op = Logic ( $\lambda ev.$  case  $ev$  of WM  $wm \Rightarrow$  (flatten_op, [WM  $wm$ ])
    | DT  $t d \Rightarrow$  (flatten_op, map (DT  $t$ )  $d$ )) [#]
definition join_list where
  join_list join st xs = (let  $t =$  case_sum id id st;
    lefts = map_filter ( $\lambda(v, d).$  case  $v$  of Inr  $_ \Rightarrow$  None
      | Inl  $t' \Rightarrow$  if  $t = t'$  then Some  $d$  else None) xs;
    rights = map_filter ( $\lambda(v, d).$  case  $v$  of Inl  $_ \Rightarrow$  None
      | Inr  $t' \Rightarrow$  if  $t = t'$  then Some  $d$  else None) xs in
    concat (map ( $\lambda d1 .$  map_filter ( $\lambda d2 .$  join  $d1 d2$ ) rights) lefts))
definition producible  $wm MA =$  ( $\exists wm' \in MA.$   $wm \leq$  case_sum Inr Inl  $wm'$ )
corec union_op where union_op wms MA = Logic ( $\lambda ev.$  case  $ev$  of
  DT  $t d \Rightarrow$  (union_op wms MA, [DT (case_sum id id  $t$ )  $d$ ])
  | WM  $wm \Rightarrow$  let  $MA' =$  maximal_antichain_set (insert  $wm MA$ ) ;
    prds = { $wm' \in$  set ( $wm \# wms$ ). producible  $wm' MA'$ };
    (union_op [ $wm' \leftarrow wm \# wms.$   $wm' \notin$  prds]  $MA'$ ,
      map (case_sum WM WM) [ $wm' \leftarrow wm \# wms.$   $wm' \in$  prds]) [#]
definition join_op tds abs wms MA join = comp_op (comp_op (incr_batch_op tds abs)
  (comp_op (map_op (join_list join)) flatten_op)) (union_op wms MA)

```

■ **Figure 9** Flatten, union, and join operators

```

mono  $lxs W \Rightarrow$  DT  $t d \in$  lset (produce (join_op [] [] [] {} join)  $lxs$ )  $\Rightarrow$ 
 $\exists d_1 d_2.$  join  $d_1 d_2 =$  Some  $d \wedge d_1 \in$  set (DTs  $lxs$  (Inl  $t$ ))  $\wedge d_2 \in$  set (DTs  $lxs$  (Inr  $t$ )) (28)
mono  $lxs W \Rightarrow$  prod  $lxs \Rightarrow$  DT (Inl  $t$ )  $d_1 \in$  lset  $lxs \Rightarrow$  DT (Inr  $t$ )  $d_2 \in$  lset  $lxs \Rightarrow$ 
join  $d_1 d_2 =$  Some  $d \Rightarrow$  DT  $t d \in$  lset (produce (join_op [] [] [] {} join)  $lxs$ ) (29)
mono  $lxs \{\} \Rightarrow$  mono (produce (join_op [] [] [] {} join)  $\{\}$ ) (30)
prod  $lxs \Rightarrow$   $\forall wm \in$  vimage WM (lset  $lxs$ ). producible  $wm$  (vimage WM (lset  $lxs$ ))  $\Rightarrow$ 
prod (produce (join_op [] [] [] {} join)) (31)

```

■ **Figure 10** Correctness properties of join\_op

587 We further introduced a reusable generalization technique using the skip\_op operator that  
 588 allows (co)inductive reasoning about elements at arbitrary positions.

589 We benefited from Isabelle’s infrastructure for coinductive predicates, codatatypes, and  
 590 corecursive functions, especially the support for friendly corecursion and monadic recursion  
 591 and associated reasoning principles. A point for future work for Isabelle’s developers could be  
 592 to automatically derive coinduction up to congruence principles for coinductive predicates such  
 593 as our manually derived principles (8) and (9). Our formalization amounts to around 17 000  
 594 lines of definitions and proofs. Of these, the heavy lifting happens for basic libraries (6 000)  
 595 and reusable operators (9 000) with batch\_op being the main culprit (5 000). In contrast,  
 596 compositional reasoning in our case studies (Section 5) benefits from this groundwork (2 000).

597 As future work, we want to use partially-ordered time-stamps to introduce a feedback  
 598 loop operator as in the Timely Dataflow stream processing framework [29, 31, 32]. Moreover,  
 599 we currently do not support parallelism. A long-term goal to extend operators with a notion  
 600 of workers they run on, which will enable us to distribute the input stream across workers as  
 601 in Timely Dataflow and reason about the correctness of the resulting distributed streaming  
 602 computation. Our formalization is executable, but it is not efficient because relies on the code  
 603 extraction to purely functional languages. We plan to connect our work to the Isabelle-LLVM  
 604 refinement framework [20] to obtain efficient executable operators.

## References

- 605  
606 1 The Isabelle/HOL formalization of the results. [https://github.com/rafaelcgs10/verified\\_stream\\_processing](https://github.com/rafaelcgs10/verified_stream_processing). Accessed: 2024-03-18.  
607
- 608 2 Martín Abadi, Frank McSherry, and Gordon D. Plotkin. Foundations of differential dataflow. In  
609 Andrew M. Pitts, editor, *Foundations of Software Science and Computation Structures - 18th*  
610 *International Conference, FoSSaCS 2015, Held as Part of the European Joint Conferences on*  
611 *Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*,  
612 volume 9034 of *Lecture Notes in Computer Science*, pages 71–83. Springer, 2015. doi:  
613 10.1007/978-3-662-46678-0\_5.
- 614 3 Andreas Abel and Brigitte Pientka. Well-founded recursion with copatterns and sized types.  
615 *J. Funct. Program.*, 26:e2, 2016. doi:10.1017/S0956796816000022.
- 616 4 Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael Fernández-  
617 Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, and  
618 Sam Whittle. The dataflow model: A practical approach to balancing correctness, latency,  
619 and cost in massive-scale, unbounded, out-of-order data processing. *Proc. VLDB Endow.*,  
620 8(12):1792–1803, 2015. doi:10.14778/2824032.2824076.
- 621 5 Ahmed Alquraan, Hatem Takruri, Mohammed Alfatafta, and Samer Al-Kiswany. An analysis  
622 of network-partitioning failures in cloud systems. In Andrea C. Arpaci-Dusseau and Geoff  
623 Voelker, editors, *OSDI 2018*, pages 51–68. USENIX Association, 2018.
- 624 6 Edmon Begoli, Tyler Akidau, Slava Chernyak, Fabian Hueske, Kathryn Knight, Kenneth L.  
625 Knowles, Daniel Mills, and Dan Sotolongo. Watermarks in stream processing systems:  
626 Semantics and comparative analysis of apache flink and google cloud dataflow. *Proc. VLDB*  
627 *Endow.*, 14(12):3135–3147, 2021. doi:10.14778/3476311.3476389.
- 628 7 Julian Biendarra, Jasmin Christian Blanchette, Aymeric Bouzy, Martin Desharnais, Mathias  
629 Fleury, Johannes Hölzl, Ondrej Kuncar, Andreas Lochbihler, Fabian Meier, Lorenz Panny,  
630 Andrei Popescu, Christian Sternagel, René Thiemann, and Dmitriy Traytel. Foundational  
631 (co)datatypes and (co)recursion for higher-order logic. In Clare Dixon and Marcelo Finger,  
632 editors, *FroCoS 2017*, volume 10483 of *LNCS*, pages 3–21. Springer, 2017. doi:10.1007/  
633 978-3-319-66167-4\_1.
- 634 8 Jasmin Christian Blanchette, Aymeric Bouzy, Andreas Lochbihler, Andrei Popescu, and  
635 Dmitriy Traytel. Friends with benefits - implementing corecursion in foundational proof  
636 assistants. In Hongseok Yang, editor, *ESOP 2017*, volume 10201 of *LNCS*, pages 111–140.  
637 Springer, 2017. doi:10.1007/978-3-662-54434-1\_5.
- 638 9 Jasmin Christian Blanchette, Johannes Hölzl, Andreas Lochbihler, Lorenz Panny, Andrei  
639 Popescu, and Dmitriy Traytel. Truly modular (co)datatypes for Isabelle/HOL. In Gerwin  
640 Klein and Ruben Gamboa, editors, *ITP 2014*, volume 8558 of *LNCS*, pages 93–110. Springer,  
641 2014. doi:10.1007/978-3-319-08970-6\_7.
- 642 10 Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas  
643 Tzoumas. Apache flink™: Stream and batch processing in a single engine. *IEEE Data Eng.*  
644 *Bull.*, 38(4):28–38, 2015. URL: <http://sites.computer.org/debull/A15dec/p28.pdf>.
- 645 11 Nils Anders Danielsson and Thorsten Altenkirch. Subtyping, declaratively. In Claude Bolduc,  
646 Josée Desharnais, and Béchir Ktari, editors, *MPC 2010*, volume 6120 of *LNCS*, pages 100–118.  
647 Springer, 2010. doi:10.1007/978-3-642-13321-3\_8.
- 648 12 Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters.  
649 In Eric A. Brewer and Peter Chen, editors, *6th Symposium on Operating System Design and*  
650 *Implementation (OSDI 2004)*, San Francisco, California, USA, December 6-8, 2004, pages  
651 137–150. USENIX Association, 2004. URL: [http://www.usenix.org/events/osdi04/tech/  
652 dean.html](http://www.usenix.org/events/osdi04/tech/dean.html).
- 653 13 Neil Ghani, Peter G. Hancock, and Dirk Pattinson. Representations of stream processors  
654 using nested fixed points. *Log. Methods Comput. Sci.*, 5(3), 2009.

- 655 14 Florian Haftmann and Tobias Nipkow. Code generation via higher-order rewrite systems. In  
656 Matthias Blume, Naoki Kobayashi, and Germán Vidal, editors, *FLOPS 2010*, volume 6009 of  
657 *LNCS*, pages 103–117. Springer, 2010. doi:10.1007/978-3-642-12251-4\_9.
- 658 15 Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. Actris: session-type  
659 based reasoning in separation logic. *Proc. ACM Program. Lang.*, 4(POPL):6:1–6:30, 2020.  
660 doi:10.1145/3371074.
- 661 16 Gilles Kahn. The semantics of a simple language for parallel programming. In Jack L. Rosenfeld,  
662 editor, *Information Processing, Proceedings of the 6th IFIP Congress 1974, Stockholm, Sweden,*  
663 *August 5-10, 1974*, pages 471–475. North-Holland, 1974.
- 664 17 Kyle Kingsbury. Jepsen: Analyses. accessed: Oct 8, 2023. URL: [https://jepsen.io/  
665 analyses](https://jepsen.io/analyses).
- 666 18 Martin Kleppmann and Jay Kreps. Kafka, samza and the unix philosophy of distributed  
667 data. *IEEE Data Eng. Bull.*, 38(4):4–14, 2015. URL: [http://sites.computer.org/debull/  
668 A15dec/p4.pdf](http://sites.computer.org/debull/A15dec/p4.pdf).
- 669 19 Alexander Krauss. Recursive definitions of monadic functions. In Ekaterina Komendantskaya,  
670 Ana Bove, and Milad Niqui, editors, *Partiality and Recursion in Interactive Theorem Provers,*  
671 *PAR@ITP 2010*, volume 5 of *EPiC Series*, pages 1–13. EasyChair, 2010. doi:10.29007/1mdt.
- 672 20 Peter Lammich. Refinement of parallel algorithms down to LLVM. In June Andronick and  
673 Leonardo de Moura, editors, *13th International Conference on Interactive Theorem Proving,*  
674 *ITP 2022, August 7-10, 2022, Haifa, Israel*, volume 237 of *LIPICs*, pages 24:1–24:18. Schloss  
675 Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.ITP.2022.24.
- 676 21 Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. TaxDC:  
677 A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In Tom  
678 Conte and Yuanyuan Zhou, editors, *ASPLOS 2016*, pages 517–530. ACM, 2016.
- 679 22 Haopeng Liu, Guangpu Li, Jeffrey F. Lukman, Jiaxin Li, Shan Lu, Haryadi S. Gunawi, and  
680 Chen Tian. DCatch: Automatically detecting distributed concurrency bugs in cloud systems.  
681 In Yunji Chen, Olivier Temam, and John Carter, editors, *ASPLOS 2017*, pages 677–691. ACM,  
682 2017.
- 683 23 Haopeng Liu, Xu Wang, Guangpu Li, Shan Lu, Feng Ye, and Chen Tian. FCatch: Automat-  
684 ically detecting time-of-fault bugs in cloud systems. In Xipeng Shen, James Tuck, Ricardo  
685 Bianchini, and Vivek Sarkar, editors, *ASPLOS 2018*, pages 419–431. ACM, 2018.
- 686 24 Andreas Lochbihler and Johannes Hölzl. Recursive functions on lazy lists via domains and  
687 topologies. In Gerwin Klein and Ruben Gamboa, editors, *ITP 2014*, volume 8558 of *LNCS*,  
688 pages 341–357. Springer, 2014. doi:10.1007/978-3-319-08970-6\_22.
- 689 25 Andreas Lochbihler and Pascal Stoop. Lazy algebraic types in Isabelle/HOL, 2018.
- 690 26 Rupak Majumdar and Filip Niksic. Why is random testing effective for partition tolerance  
691 bugs? *PACMPL*, 2(POPL):46:1–46:24, 2018.
- 692 27 Konstantinos Mamouras. Semantic foundations for deterministic dataflow and stream pro-  
693 cessing. In Peter Müller, editor, *ESOP 2020*, volume 12075 of *LNCS*, pages 394–427. Springer,  
694 2020. doi:10.1007/978-3-030-44914-8\_15.
- 695 28 Frank McSherry. Github: Differential dataflow. URL: [https://github.com/timelydataflow/  
696 differential-dataflow/](https://github.com/timelydataflow/differential-dataflow/).
- 697 29 Frank McSherry. Github: Timely dataflow. URL: [https://github.com/TimelyDataflow/  
698 timely-dataflow/](https://github.com/TimelyDataflow/timely-dataflow/).
- 699 30 Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. Differential  
700 dataflow. In *Sixth Biennial Conference on Innovative Data Systems Research, CIDR 2013,*  
701 *Asilomar, CA, USA, January 6-9, 2013, Online Proceedings*. [www.cidrdb.org](http://www.cidrdb.org), 2013.
- 702 31 Derek Gordon Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and  
703 Martín Abadi. Naiad: a timely dataflow system. In Michael Kaminsky and Mike Dahlin,  
704 editors, *SOSP 2013*, pages 439–455. ACM, 2013. doi:10.1145/2517349.2522738.

- 705 32 Derek Gordon Murray, Frank McSherry, Michael Isard, Rebecca Isaacs, Paul Barham, and  
706 Martín Abadi. Incremental, iterative data processing with timely dataflow. *Commun. ACM*,  
707 59(10):75–83, 2016. doi:10.1145/2983551.
- 708 33 Lawrence C. Paulson. *A Fixedpoint Approach to (Co)Inductive and (Co)Datatype Definitions*,  
709 page 187–211. MIT Press, Cambridge, MA, USA, 2000.
- 710 34 Robert Sandner and Olaf Müller. Theorem prover support for the refinement of functions.  
711 In Ed Brinksma, editor, *TACAS 1997*, volume 1217 of *LNCS*, pages 351–365. Springer, 1997.  
712 doi:10.1007/BFb0035399.
- 713 35 Maria Spichkova. *Specification and seamless verification of embedded real-time systems:  
714 FOCUS on Isabelle*. PhD thesis, Technical University Munich, Germany, 2007. URL: <http://mediatum.ub.tum.de/doc/620981/document.pdf>.
- 716 36 Caleb Stanford, Konstantinos Kallas, and Rajeev Alur. Correctness in stream processing:  
717 Challenges and opportunities. In *CIDR 2022*. [www.cidrdb.org](http://www.cidrdb.org), 2022.
- 718 37 Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, Yongle Zhang,  
719 Pranay Jain, and Michael Stumm. Simple testing can prevent most critical failures: An  
720 analysis of production failures in distributed data-intensive systems. In Jason Flinn and Hank  
721 Levy, editors, *OSDI 2014*, pages 249–265. USENIX Association, 2014.
- 722 38 Matei Zaharia, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica.  
723 Discretized streams: fault-tolerant streaming computation at scale. In Michael Kaminsky and  
724 Mike Dahlin, editors, *SOSP 2013*, pages 423–438. ACM, 2013. doi:10.1145/2517349.2522737.