

Stream Types

JOSEPH W. CUTLER, University of Pennsylvania, USA

CHRISTOPHER WATSON, University of Pennsylvania, USA

PHILLIP HILLIARD, University of Pennsylvania, USA

HARRISON GOLDSTEIN, University of Pennsylvania, USA

CALEB STANFORD, University of California, Davis, USA

BENJAMIN C. PIERCE, University of Pennsylvania, USA

We propose a rich foundational theory of typed data streams and stream transformers, motivated by two high-level goals: (1) the type of a stream should be able to express complex *sequential patterns* of events over time, and (2) it should describe the *parallel structure* of the stream to enable deterministic stream processing on parallel and distributed systems. To this end, we introduce *stream types*, with operators capturing sequential composition, parallel composition, and iteration, plus a core calculus of *transformers* over typed streams which naturally supports a number of common streaming idioms, including punctuation, windowing, and parallel partitioning, as first-class constructions. The calculus exploits a Curry-Howard-like correspondence with an ordered variant of the logic of Bunched Implication to program with streams compositionally and uses Brzozowski-style derivatives to enable an incremental, event-based operational semantics. To validate our design, we provide a reference interpreter and machine-checked proofs of the main results.

ACM Reference Format:

Joseph W. Cutler, Christopher Watson, Phillip Hilliard, Harrison Goldstein, Caleb Stanford, and Benjamin C. Pierce. 2024. Stream Types. In *Proceedings of The 51st ACM SIGPLAN Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 68 pages.

1 INTRODUCTION

What is the type of a stream? A straightforward answer, dating back to the early days of functional programming [16], is that a stream is an unbounded sequence of items of a single fixed type, produced by one part of a system (or the external world) and consumed by another. This perspective has been immensely successful, with the programming models exposed by the most popular distributed stream processing eDSLs (e.g., Flink [17, 25], Beam [30], Storm [29], and Heron [26]) typically offering just one type, `Stream t`.

However, this homogeneous treatment of streams leaves something to be desired in the world of modern stream processing. For one thing, streaming data sometimes arrives at a processing node from multiple sources *in parallel*. Using arrival times to impose an “incidental” order on such parallel data can make it difficult to ensure that processing is deterministic because downstream results may then depend on external factors like network latency [49, 55]. Another issue is that temporal patterns like *bracketedness* (every “begin” event has a following “end”) or streams with exactly k events are invisible in the stream’s type. Programmers cannot trust the type system to ensure these properties when producing a stream, nor can they rely on them when consuming a stream; see Section 2 for examples.

We offer a new logical foundation for typed stream processing that can describe streams with *both* complex sequential patterns and parallel structure. On this foundation, we build a calculus called λ^{ST} that is (a) expressive and type-safe for streams with complex temporal patterns and (b) deterministic, even in the presence of parallel inputs.

Programs in λ^{ST} are intuitively batch processors that operate over entire streams at once. But, since streams are in general unbounded, stream transformers can't actually wait for "the entire input stream" to arrive before producing any output. The operational semantics of the λ^{ST} calculus is therefore *incremental*, producing partial outputs from partial inputs on the fly. Concretely, a λ^{ST} program is interpreted as a function mapping any prefix of its input(s) to a prefix of its output plus a "derivative" term that is ready to transform the rest of the inputs to the rest of the output.

Our stream types include two kinds of products, one representing a pair of streams in temporal sequence, the other a pair of streams in parallel. This structure is inspired both by Concurrent Kleene Algebras [35, 45], which syntactically describe partially ordered series/parallel data, and work by Alur et al. [3] and Mamouras et al. [49], where streams are modeled as partially ordered sets. We find a suitable proof theory for this two-product formalism in a variant of O'Hearn and Pym's Logic of Bunched Implications (BI) [52]. BI has famously been used as a foundation for separation logic [54], where the "separating conjunction" allows for local reasoning about separate regions of the heap in imperative programs. In λ^{ST} , we replace spatial separation with *temporal separation*: one product describes pairs of streams separated sequentially in time; the other, independent pairs of streams whose elements may arrive in interleaved fashion.

Concretely, our contributions are:

- (1) We propose *stream types*, a type discipline for distributed stream processing that generalizes the traditional homogeneous view to a richer nested-parallel-and-sequential structure.
- (2) We define a calculus λ^{ST} of stream processing transformers, inspired by a Curry-Howard-like correspondence with an ordered variant of BI. Terms in λ^{ST} are high level programs in a functional style that conceptually transform whole streams.
- (3) We equip λ^{ST} with an operational semantics interpreting terms as incremental transformers that accept and produce finite prefixes of streams. Our main technical result is a powerful *homomorphism theorem* (Theorem 3.3) guaranteeing that the result of a transformer does not depend on how the input stream is divided into successive prefixes.
- (4) We formalize the type system and semantics in Coq and prove that the incremental semantics is deterministic: all interleavings of parallel substreams yield the same final result.
- (5) We demonstrate by example how λ^{ST} enables type-safe programming for streams with complex patterns and prevents nondeterminism and how programming patterns from stream processing practice are elegantly supported by this richer model, including MapReduce-like pipelines, temporal integrity constraints, windowing, punctuation, parallelism control, routing, and side outputs.
- (6) We show that the high-level operational semantics of λ^{ST} can be implemented atop conventional sequential streams by compiling stream-typed terms to string transducers, opening a path to a future distributed implementation based on an existing industrial-strength runtime.

The rest of the paper is structured as follows. Section 2 explores some concrete cases where λ^{ST} 's structured types can prevent common stream processing bugs and enable cleaner programming patterns. Section 3 presents Kernel λ^{ST} , a minimal subset with just the features needed to state and understand our main results. Section 4 extends this presentation to the full λ^{ST} . Section 5 develops several examples and argues that λ^{ST} solves the problems presented in Section 2. Section 6 shows how to compile λ^{ST} terms to string transducers. Sections 7 and 8 discuss related and future work. Overviews of the Coq formalization of our main results and our prototype interpreter in Haskell can be found in Appendix A, technical details omitted from the main paper in Appendix B, and expanded versions of the examples in Appendix C.

2 MOTIVATING EXAMPLES

Types for temporal invariants. Consider a stream of brightness data coming from a motion sensor, where each event in the stream is a number between 0 and 100. Suppose we want a stream transformer that acts as a threshold filter, sending out a “Start” event when the brightness level goes above level 50, forwarding along brightness values until the level dips below the threshold, and sending a final “Stop” event. For example:

$$11, 30, 52, 56, 53, 30, 10, 60, 10, \dots \implies \text{Start}, 52, 56, 53, \text{Stop}, \text{Start}, 60, \text{Stop}, \dots$$

The output of the transformer should satisfy the following *temporal invariant*: each start event must be followed by one or more data events, and then one end event. Unfortunately, usual stream processing systems would give this transformation a type that looks like $\text{Stream Int} \rightarrow \text{Stream} (\text{Start} + \text{Int} + \text{Stop})$, which expresses only the types of events in the output, not the temporal invariant on them.

These simple types become even more of a problem when *consuming* streams. Suppose a second transformer wanted to consume the output stream of type $\text{Stream} (\text{Start} + \text{Int} + \text{Stop})$ and compute the average brightness between each start/end pair. We know a priori that the stream is well bracketed, but the type does not say so. Thus, the transformer must *re-parse* the stream to compute the averages, requiring additional logic for various special cases (e.g., empty start/end pairs) that cannot actually occur in the real stream!

In λ^{ST} , we can express this invariant with the type $(\text{Start} \cdot \text{Int} \cdot \text{Int}^* \cdot \text{End})^*$, specifying that the stream consists of a start message, at least one Int , and an end message, repeatedly. A well-typed transformer with this output type is guaranteed to enforce the desired invariant; conversely, a downstream transformer can assume that its input will adhere to the invariant.

Enforcing deterministic parallelism. A second limitation of homogeneous streams is that they impose a *total ordering* on their component events.¹ In other words, for each pair of events in the stream, the transformer can tell which came first. Imposing a total ordering like this is problematic in a world where stream transformers work over data that is logically partially ordered—e.g., because it comes from separate sources. For example, consider a pair of sensors, each producing one reading per second and sending them different network connections to a single transformer that averages them pairwise, producing a composite reading each second.

A natural way to do this is to merge the two streams, group adjacent pairs of elements (i.e., impose a size-two tumbling window), and average the pairs, but this is subtly wrong: a network delay could cause a pair of consecutive elements in the merged stream to be from the same sensor, after which the averages will all be bogus. The problem here is that this transformer is not deterministic: its result can depend on external factors like network latency. Bugs of this type can easily occur in practice [49, 55] and can be very difficult to track down, since they may only manifest under rare conditions [43].

Once again, this is a failure of type structure. In λ^{ST} , we would give the merged stream the type $(\text{Sensor1} \parallel \text{Sensor2})^*$, capturing the fact that it is a stream of *parallel pairs* of readings from the two sensors. We can write a strongly typed merge operator that produces this type, given parallel streams of type Sensor1^* and Sensor2^* . This merge operator is deterministic because all well-typed λ^{ST} programs are (Section 3.4); operationally, it waits for events to arrive on both of its input streams before sending them along as a pair.

¹In practical stream processing systems, this is often a total ordering *per key* for a parallelized stream—cf. `KeyedStream` in Flink—but the same objections apply.

3 KERNEL λ^{ST}

In this section, we define the most important constructors of stream types together with the corresponding features of the term language; these form the “kernel” of the λ^{ST} calculus. The rest of the types and terms of full λ^{ST} will be layered on bit by bit in Section 4.

The *concatenation type* constructor describes streams that *vary over time*: if s and t are types of streams, then a stream of type $s \cdot t$ behaves first like a stream of type s , and then, if the s part finishes, like a stream of type t . Operationally, streams of type $s \cdot t$ can include a *punctuation marker* [60] indicating where they cross over from their s part to their t part. For example, streams of type s^* are distinguishable from streams of type $s^* \cdot s^*$ because a transformer accepting the latter can see when its input crosses from the first s^* to the second.

The *parallel type* constructor, written $s||t$, describes streams with two parallel substreams of types s and t . Each element in a parallel stream is tagged to indicate which substream it belongs to. This means that streams of type $s||t$ are isomorphic, but not identical, to streams of type $t||s$, and similarly $\text{Int}^*||\text{Int}^*$ is not the same as Int^* . Semantically, the s and t components are produced and consumed independently: a transformer that produces $s||t$ may send out an entire s first and then a t , or an entire t and then the s , or any interleaving of the two. Conversely, a transformer that accepts $s||t$ must handle all these possibilities uniformly by processing the s and t independently.

Parallel types can be combined with concatenation types in interesting ways. For example, a stream of type $(s||t) \cdot r$ consists of a stream of interleaved items from s and t , followed (once all the s 's and t 's have arrived) by a stream of type r . By contrast, a stream of type $(s \cdot t)||s' \cdot t'$ has two parallel components, one a stream described by s followed by a stream described by t and the other an s' followed by a t' . The fact that the parallel type is on the outside means that the change-over points from s to t and s' to t' are completely independent.

The base type 1 describes a stream containing just one data item, itself a unit value. The other base type is ε , the type of the empty stream containing no data; it is the unit for both the \cdot and $||$ constructors—i.e., $s \cdot \varepsilon$, $\varepsilon \cdot s$, $\varepsilon||s$ and $s||\varepsilon$ are all equivalent to s , in the sense that there are λ^{ST} transformers that convert between them. In summary, the Kernel λ^{ST} stream types are given by the grammar on the top left in Figure 1. (Of course, these kernel types can only describe streams of fixed, finite size. In Section 4.3 we will introduce unbounded streams via the Kleene star type s^* .)

What about terms? Our goal is to develop a language of core terms e , typed by stream types, where well-typed terms $x : s \vdash e : t$ are interpreted as incremental stream transformers accepting a stream described by s and producing one described by t . But one input stream is not enough: in order to be compositional, the type system needs to be able to handle stream transformers with multiple parallel and sequential inputs. To enable this, we draw upon results from proof theory for insight. Both the types $s \cdot t$ and $s||t$ are *product types*, in the sense that streams of these types should contain the data of a stream of type s and a stream of type t —although the temporal structure differs between the two. In situations where a logic (or type theory) includes two products, standard techniques from proof theory tell us that the corresponding typing judgment requires a context with two *context formers*.

The first of these, written with a comma (Γ, Δ) , is standard in intuitionistic logics. The second, written with a semicolon $(\Gamma; \Delta)$, is the interesting one, behaving intuitively like the parallel type constructor. A context Γ, Δ describes inputs to a transformer arriving from two different parts of the environment, one structured according to Γ and the other according to Δ . On the other hand, the semicolon corresponds to concatenation: a context $\Gamma; \Delta$ means that data will first arrive from the environment according to Γ , then according to Δ . Such *bunched* contexts were first introduced in the logic of Bunched Implication [52], the basis of modern separation logic [54]. In BI, the separating context former is commutative, while our semicolon context former is *not*.

$$\begin{array}{c}
s, t, r := 1 \mid \varepsilon \mid s \cdot t \mid s \parallel t \qquad \Gamma ::= \cdot \mid \Gamma, \Gamma \mid \Gamma; \Gamma \mid x : s \\
\\
\frac{\Gamma \vdash e_1 : s \quad \Gamma \vdash e_2 : t}{\Gamma \vdash (e_1, e_2) : s \parallel t} \text{PAR-R} \qquad \frac{\Gamma(x : s, y : t) \vdash e : r}{\Gamma(z : s \parallel t) \vdash \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e : r} \text{PAR-L} \\
\\
\frac{\Gamma \vdash e_1 : s \quad \Delta \vdash e_2 : t}{\Gamma; \Delta \vdash (e_1; e_2) : s \cdot t} \text{CAT-R} \qquad \frac{\Gamma(x : s; y : t) \vdash e : r}{\Gamma(z : s \cdot t) \vdash \text{let}_{\Gamma(-)}(x; y) = z \text{ in } e : r} \text{CAT-L} \\
\\
\frac{}{\Gamma \vdash \text{sink} : \varepsilon} \text{EPS-R} \qquad \frac{}{\Gamma \vdash () : 1} \text{ONE-R} \qquad \frac{}{\Gamma(x : s) \vdash (x : s @ \Gamma(-)) : s} \text{VAR} \\
\\
\frac{pf : \Gamma \leq \Gamma' \quad \Gamma' \vdash e : s}{\Gamma \vdash \text{subctx}(pf, e) : s} \text{SUBCTX}
\end{array}$$

Fig. 1. Kernel λ^{ST} syntax and typing rules

Formally, stream contexts are drawn from the grammar at the top right of Figure 1. Our typing judgment is now $\Gamma \vdash e : s$. Intuitively, this means that e is a transformer from a collection of streams structured like Γ to a single stream structured like s .

The *structural rules* for manipulating these contexts will determine which sequents are provable and hence which programs are type-correct. For example, including an exchange rule for the comma connective—allowing Γ, Δ to be rewritten as Δ, Γ —ensures that $s \parallel t$ is interprovable with $t \parallel s$; a proof of this fact will be a transformer that swaps the parallel components of the stream. On the other hand, $s \cdot t$ is manifestly different from $t \cdot s$: in a stream of the first type, s arrives before t ; in the second, t before s . In other words, concatenation is an *ordered* product. The consequence of this is that our type system does not include an exchange rule for semicolon contexts. A contraction rule for semicolon—converting from Γ to $\Gamma; \Gamma$ —is similarly undesirable, as every use of this structural rule would require the runtime system to save and then “replay” a potentially unbounded incoming stream.

In summary, our type system is substructural. The semicolon context former is ordered (no exchange) and affine (no contraction), while the comma context former is fully structural. Both context formers are associative, with the empty context serving as a unit for both.

3.1 Kernel Typing Rules

The typing rules for Kernel λ^{ST} are collected in Figure 1. For technical reasons, the rules are presented in sequent-calculus style, rather than the more familiar natural deduction style. The main difference is that sequent calculi have left and right rules—describing how to eliminate a connective when it’s in the context or the result type, respectively—as opposed to natural deduction’s introduction and elimination rules. (In the presence of a CUT rule, the two presentations are equivalent. The semantics of CUT itself is a bit complex, though, and we choose the sequent-calculus presentation to isolate this complexity in one place.)

The most straightforward typing rule is the right rule for parallel (PAR-R). From a context Γ , we can produce a stream of type $s \parallel t$ by producing s and t independently from Γ , using transformers e_1 and e_2 . We write the combined transformer as a “parallel pair”, (e_1, e_2) . Semantically, this transformer operates by copying the inputs arriving on Γ , passing the copies to e_1 and e_2 , and pairing up the outputs into a parallel stream. Similarly, the CAT-R rule is used to produce a stream of type $s \cdot t$.

It uses a similar pairing syntax—if term e_1 has type s and e_2 has type t , then the “sequential pair” $(e_1; e_2)$ has type $s \cdot t$ —but the context in the conclusion differs. Since e_1 needs to be run before e_2 , the part of the input stream that e_1 depends on (Γ) must arrive before the part that e_2 depends on (Δ). Semantically, this term will operate by accepting data from the Γ part of the context and running e_1 ; once the Γ part is finished and the Δ part starts to arrive, it will switch to running e_2 .

These right rules describe how to *produce* a stream of parallel or concatenation type; the corresponding left rules describe how to *use* a variable of one of these types appearing somewhere in the context. Syntactically, the terms take the form of let-bindings: we can deconstruct variables of type $s \cdot t$ or $s||t$ as pairs of variables of type s and t , connected by either $;$ or $,$ as appropriate. We use the standard BI notation $\Gamma(-)$ for a context with a hole, and $\Gamma(\Delta)$ when this hole has been filled with the context Δ . In particular, $\Gamma(x : s)$ is a context with a distinguished variable x at some leaf.

The PAR-L rule says that if z is a variable of type $s||t$ somewhere in the context, we can let-bind it as a pair of variables x and y of type s and t and use these in a continuation term e of type r . When typing e , the variables x and y appear in the same position as the original variable z , but separated by a comma—i.e., x and y are assumed to arrive in parallel. Similarly, the parallel rule CAT-L says that, if a variable z of type $s \cdot t$ appears in the context, it can be let-bound to a pair of variables x and y of types s and t , which are again used in the continuation e . This time, though, x and y are separated by a semicolon—i.e., the substream bound to x will arrive first, followed by the substream bound to y .

EPS-R and ONE-R are the right rules for the two base types, witnessed by the terms `sink` and `()`. Semantically, `sink` does nothing: it accepts inputs on Γ and produces no output. On the other hand, `()` emits a unit value as soon as it receives its first input, and then never emits anything else.

The variable rule (VAR) says that, if $x : s$ is a variable somewhere in the context, then we can simply return it. Semantically, it works by dropping everything in the context except for the s -typed data for x , which it forwards along. The corresponding proof term is $(x : s@ \Gamma(-))$, where x is the variable itself, s is its type, and $\Gamma(-)$ is surrounding context. The annotations s and $\Gamma(-)$ are included because they will be needed by the operational semantics. Because the input to a term will itself be tree-structured, the context location $\Gamma(-)$ is required to look up the data corresponding to x in the input tree. And the output type s is required when this variable lookup comes back empty handed because the prefix of the context that has arrived so far stops before it gets to any data corresponding to this variable; in this case, the semantics of the VAR rule needs to construct an “empty” value, whose shape depends on the type s . (Although these annotations are required at runtime, we expect they will be inferrable in a high-level language based on λ^{ST} —see Section 8.)

Finally, the rule SUBCTX encodes all of the structural rules as a subtyping relation on the context. For example, the weakening rule for semicolon contexts is written $\Gamma; \Delta \leq \Gamma$, and the exchange rule is $\Gamma, \Delta \leq \Delta, \Gamma$. Each of these structural rules has semantic content. For example, choosing to use the comma contraction rule $\Gamma \leq \Gamma, \Gamma$ requires duplicating the input at runtime. Because of this, the structural rules are marked in the syntax explicitly as `subctx` (pf, e), where pf is a proof term describing which structural rules structural rules were used. See Appendix B.7 for details.

Examples. To show the typing rules in action, here are two small examples of transformers written in Kernel λ^{ST} .² The first is a simple “parallel-swap” transformer, which accepts a stream z of type $s||t$, and outputs a stream of type $t||s$, swapping the “positions” of the parallel substreams: $z : s||t \vdash \text{let } (x, y) = z \text{ in } (y, x) : t||s$. It works by splitting the variable $z : s||t$ into variables $x : s$ and $y : t$ and yielding a parallel pair with the order reversed. The second example is a “prepend”

²We elide the annotations on variables and let-bindings, as well as terms for structural rules, following the style used for the examples in Section 5. Fully explicit versions can be found in Appendix C, and elaboration from a high level language to the core is discussed in Section 8.

$$\begin{array}{c}
\frac{}{\text{oneb} : \text{batch}(1)} \qquad \frac{}{\text{epsb} : \text{batch}(\varepsilon)} \\
\frac{b : \text{batch}(s) \quad b' : \text{batch}(t)}{\text{catb}(b, b') : \text{batch}(s \cdot t)} \qquad \frac{b : \text{batch}(s) \quad b' : \text{batch}(t)}{\text{parb}(b, b') : \text{batch}(s||t)} \\
\frac{b : \text{batch}(s)}{b : \text{batch}((x : s))} \qquad \frac{b : \text{batch}(\Gamma) \quad b' : \text{batch}(\Gamma')}{\text{catb}(b, b') : \text{batch}(\Gamma; \Gamma')} \qquad \frac{b : \text{batch}(\Gamma) \quad b' : \text{batch}(\Gamma')}{\text{parb}(b, b') : \text{batch}(\Gamma, \Gamma')}
\end{array}$$

Fig. 2. Batches of types and contexts

transformer, which takes a variable $x : s$ and outputs a stream of type $1 \cdot s$, prepending a unit value to the front: $x : s \vdash ((; x) : 1 \cdot s$.

3.2 Batches and Batch Processing Semantics

Although our goal is to develop an incremental semantics for stream transformers, it turns out that we need a “batch semantics” as a subroutine, where the full contents of a stream is provided to a transformer and it produces its entire output all at once. A *batch*, then, is a structured value representing the full history of a completed stream.

Stream types dictate the shape of their batches. In particular, a batch of a parallel stream type $s||t$ is a pair of a batch of s and a batch of t . Crucially, this definition encodes no information about any particular interleaving of the batch of s and the batch of t . Similarly, a batch of a concatenation $s \cdot t$ is also a batch of s and a batch of t . The grammar of batches is

$$b ::= \text{oneb} \mid \text{epsb} \mid \text{catb}(b, b') \mid \text{parb}(b, b')$$

where $\text{catb}(b, b')$ and $\text{parb}(b, b')$ construct batches for parallel and concatenation types and oneb and epsb are the unique batches of types 1 and ε . Batches are given types by the judgment $b : \text{batch}(s)$, the rules of which are given in Figure 2. This relation is further lifted to contexts: Γ, Δ and $\Gamma; \Delta$ are the batches of $\text{parb}(b, b')$ and $\text{catb}(b, b')$, respectively, when b and b' are batches of Γ and Δ .

The batch semantics is defined by the judgment $b \Rightarrow e \Rightarrow b'$, pronounced “running term e on batch b produces batch b' ” (see Figure 3). The following theorem establishes the correctness of the semantics: well typed terms $\Gamma \vdash e : s$ take batches of Γ to batches of s .

THEOREM 3.1 (BATCH SAFETY). *If $\Gamma \vdash e : s$ and $b : \text{batch}(\Gamma)$, then there is a unique b' such that $b \Rightarrow e \Rightarrow b'$; moreover, $b' : \text{batch}(s)$.*

This and all of the theorems in the paper are formalized in Coq (see Appendix A) and discussed in more detail in the Appendix B.

The cases for the two pairing constructs in the batch semantics are straightforward. In B-PAR-R, we take the input batch, pass it to both e_1 and e_2 , and then combine the results as a parallel pair of batches. In B-CAT-R, the input must be a concatenation batch because the conclusion of the CAT-R rule has a semicolon context $\Gamma; \Delta$. We pass the first batch to e_1 and the second to e_2 , and we combine the results as a concatenation pair of batches. On the other hand, the batch semantics for let-bindings is trivial, since it is easy to verify that a batch of $\Gamma(z : s \cdot t)$ is exactly the same as a batch for $\Gamma(x : s; y : t)$, and similarly for parallel types and comma contexts. For this reason, both left rules can simply pass their input batches to the continuation and return the result.

The rules B-EPS-R and B-ONE-R simply return the unique batches of their respective types.

$$\begin{array}{c}
\frac{b \Rightarrow e_1 \Rightarrow b_1 \quad b \Rightarrow e_2 \Rightarrow b_2}{b \Rightarrow (e_1, e_2) \Rightarrow \text{parb}(b_1, b_2)} \text{B-PAR-R} \quad \frac{b \Rightarrow e \Rightarrow b'}{b \Rightarrow \text{let}_{\Gamma(-)}(z, x) = y \text{ in } e \Rightarrow b'} \text{B-PAR-L} \\
\frac{b_1 \Rightarrow e_1 \Rightarrow b'_1 \quad b_2 \Rightarrow e_2 \Rightarrow b'_2}{\text{catb}(b_1, b_2) \Rightarrow (e_1; e_2) \Rightarrow \text{catb}(b'_1, b'_2)} \text{B-CAT-R} \quad \frac{b \Rightarrow e \Rightarrow b'}{b \Rightarrow \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e \Rightarrow b'} \text{B-CAT-L} \\
\frac{}{b \Rightarrow \text{sink} \Rightarrow \text{epsb}} \text{B-EPS-R} \quad \frac{}{b \Rightarrow () \Rightarrow \text{oneb}} \text{B-ONE-R} \quad \frac{b \rightsquigarrow_{\Gamma(-)} b'}{b \Rightarrow (x : s @ \Gamma(-)) \Rightarrow b'} \text{B-VAR} \\
\frac{\text{coeBatch}(pf, b) \sim b' \quad b' \Rightarrow e \Rightarrow b''}{b \Rightarrow \text{subctx}(pf, e) \Rightarrow b''} \text{B-SUBCTX}
\end{array}$$

Fig. 3. Batch semantics rules (input and output batched are red, terms are blue)

The variable rule (B-VAR) takes the input batch b and uses the context location $\Gamma(-)$ to look up the part of b that corresponds to the variable x . This lookup is accomplished with an auxiliary relation called *batch projection*, written $b \rightsquigarrow_{\Gamma(-)} b'$ and defined in Appendix B.5. For example, suppose $b = \text{catb}(b_1, \text{parb}(b_2, b_3))$ is a batch for the context $(x : s; (y : t, z : r))$. Then we have $b \rightsquigarrow_{(x:s; (-, z:r))} b_2$ holds, i.e., b_2 is the batch located at the hole where y occurs.

The rule for context subtyping accepts a batch b , rearranges it by transporting it along the subtyping relation pf , and then runs it through e . This rule uses another auxiliary relation called *batch transport* and written $\text{coeBatch}(pf, b) \sim b'$. This relation transforms a batch for Γ into a (unique) batch for Γ' when $pf : \Gamma \leq \Gamma'$. See Appendix B.7 for details.

3.3 Prefixes and Incremental Semantics

We're now ready to discuss the *incremental* semantics of Kernel λ^{ST} . The natural notion of “value” for an incremental semantics is finite prefixes of streams: that is, a well-typed term $\Gamma \vdash e : s$ should accept any prefix of a stream of type Γ and produce a prefix of a stream of type s .

(Ultimately, we will be interested in running transformers one event at a time, as discussed in Section 6. Even so, the generalization to arbitrary prefixes here is needed to define the semantics compositionally, since the behavior of a composite transformer in response to a single event may involve generating and processing multiple events internally.)

The possible prefixes of a stream can be calculated from its type. In particular, there are two prefixes of a stream of type 1: the empty prefix, written onepA , and the prefix containing the single element $()$, written onepB . Similarly, the unique stream of type ε has a single prefix, the empty prefix, which we write epsb .

What about $s \parallel t$? A parallel stream of type $s \parallel t$ is conceptually a pair of independent streams of type s and t , so a prefix of a parallel stream should be a pair $\text{parp}(p_1, p_2)$, where p_1 is a prefix of a stream of type s , and p_2 is a prefix of a stream of type t . As with batches of streams of parallel type, this definition includes no information about orderings between the two parallel components: the prefix $\text{parp}(p_1, p_2)$ equally represents a situation where p_1 arrived first and the p_2 did, one where p_2 arrived before p_1 , and one where the data for p_1 and p_2 arrived in some interleaved order. In a nutshell, this definition is what guarantees deterministic processing. By representing all possible interleavings using the same prefix value, we ensure that a transformer that operates on these values cannot possibly depend on ordering information that isn't present in the type.

$$\begin{array}{c}
\overline{\text{emsp} : \text{prefix}(\varepsilon)} \\
\frac{p : \text{prefix}(s) \quad p' : \text{prefix}(t)}{\text{parp}(p, p') : \text{prefix}(s||t)} \\
\overline{\text{onepA} : \text{prefix}(1)} \\
\frac{p : \text{prefix}(s)}{\text{catpA}(p) : \text{prefix}(s \cdot t)} \\
\overline{\text{onepB} : \text{prefix}(1)} \\
\frac{b : \text{batch}(s) \quad p : \text{prefix}(t)}{\text{catpB}(b, p) : \text{prefix}(s \cdot t)}
\end{array}$$

Fig. 4. Prefixes for types

$$\begin{array}{c}
\overline{\delta_{\text{emsp}}(\varepsilon) \sim \varepsilon} \quad \overline{\delta_{\text{onepA}}(1) \sim 1} \quad \overline{\delta_{\text{onepB}}(1) \sim \varepsilon} \quad \overline{\delta_{\text{catpA}(p)}(s \cdot t) \sim s' \cdot t} \\
\frac{\delta_p(t) \sim t'}{\delta_{\text{catpB}(b, p)}(s \cdot t) \sim t'} \quad \frac{\delta_p(s) \sim s' \quad \delta_{p'}(t) \sim t'}{\delta_{\text{parp}(p, p')}(s||t) \sim s' || t'}
\end{array}$$

Fig. 5. Derivatives

Finally, let's consider the prefixes of streams of type $s \cdot t$. One case is a prefix that only includes data from s because it is cut off before reaching the point where the $s \cdot t$ stream stops being s and starts being t . We write such a prefix as $\text{catpA}(p)$, with p a prefix of type s . The other case is where the prefix does include the crossover point—i.e., it consists of a complete batch of s plus a prefix of t . We write this as $\text{catpB}(b, p)$, with b a batch of s and p a prefix of t .

We formalize all these possibilities as a judgment $p : \text{prefix}(s)$ (see Figure 4). We further lift this judgment to contexts, with prefixes of streams of comma contexts behaving like prefixes of streams of parallel type and prefixes of streams of semicolon contexts behaving like prefixes of streams of concatenation type.

Every type s has a distinguished *empty prefix*, written emp_s and defined by straightforward recursion on s . There is also a natural operation that takes a batch and turns it into a prefix, “forgetting” the fact that it was a complete batch. We write this operation as $(b)^\circ$. See Appendix B.1 for details.

Derivatives. After a well-typed transformer $\Gamma \vdash e : s$ accepts a prefix p of type Γ and produces a prefix p' of type s , it should transition to a new transformer e' that is prepared to accept the rest of the input stream to produce the rest of the output stream. In other words, e' should have type “the rest of a stream of type s , after p' ” in context “the rest of a stream of type Γ , after p ”.

The “rest” of a type or context is, intuitively, its *derivative* in the sense of standard Brzozowski derivatives of regular expressions [15]. Formally, we define a 3-place relation, $\delta_p(s) \sim s'$, pronounced “the derivative of s with respect to p is s' ” (see Figure 5). This relation is actually a partial function that is defined when $p : \text{prefix}(s)$; in this case, we write $\delta_p(s)$ for the unique s' such that $\delta_p(s) \sim s'$.

The derivative of the type 1 with respect to the empty prefix onepA is 1 (the “rest” of the stream is the entire stream), and the derivative with respect to the full prefix onepB is ε (there is no more stream left after the unit element has arrived). For parallel, the derivative is taken component-wise. The most interesting cases are those for the concatenation type. If the prefix has the form $\text{catpA}(p)$, the derivative $\delta_{\text{catpA}(p)}(s \cdot t)$ is $(\delta_p(s)) \cdot t$, i.e., some of the s has gone by but not all, and we still expect t to come after it. On the other hand, if the prefix has the form $\text{catpB}(b, p)$, the derivative

$\delta_{\text{catpB}(b,p)}(s \cdot t)$ is just $\delta_p(t)$, i.e., the s component is complete, and the rest of the stream is just the part of t after p .

Incremental Semantics. The incremental semantics is given by a judgment $p \Rightarrow e \downarrow e' \Rightarrow p'$, pronounced “running the core term e on the input prefix p yields the output prefix p' and steps to e' .” The following correctness theorem establishes correctness for this semantics: if we run a well-typed core term on a prefix from the context type, it will return a prefix with the result type and step to a term that can accept the rest of the context and produce the rest of the result.

THEOREM 3.2 (CORRECTNESS OF THE INCREMENTAL SEMANTICS). *If $\Gamma \vdash e : s$ and $p : \text{prefix}(\Gamma)$, then there are unique e' and p' such that $p \Rightarrow e \downarrow e' \Rightarrow p'$; moreover, $p' : \text{prefix}(s)$ and $\delta_p(\Gamma) \vdash e' : \delta_{p'}(s)$.*

The rules for the incremental semantics of Kernel λ^{ST} are gathered in Figure 6 and described in the remainder of this subsection; the full set of rules for all of λ^{ST} can be found in Appendix B.10.

Incremental Semantics of the Right Rules. The right rules for parallel and concatenation are the simplest to understand, so we address them first. For P-PAR-R, we accept a prefix p and use it to run the component terms e_1 and e_2 , independently producing outputs p_1 and p_2 and stepping to new terms e'_1 and e'_2 . The pair term (e_1, e_2) then steps to (e'_1, e'_2) and produces the output $\text{parp}(p_1, p_2)$.

There are two right rules for the concatenation pair $(e_1; e_2) : s \cdot t$. Because the conclusion of the corresponding typing rule ensures that the input context has the form $\Gamma; \Delta$, there are two possible shapes of input prefixes: either $\text{catpA}(p)$, with p a prefix of type Γ , or else $\text{catpB}(b, p)$, with b a batch of type Γ and p a prefix of type p . If the input is $\text{catpA}(p)$, the rule P-CAT-R-1 applies: we run e_1 with p , which returns p' as output and e'_1 as a new term. The original term then steps to $(e'_1; e_2)$ and returns $\text{catpA}(p')$. (Since the input prefix includes no data for the second half of the context, we do not need to run e_2 at all.) On the other hand, if the input is $\text{catpB}(b, p)$, the rule P-CAT-R-2 applies. Here, the data required to run e_1 has arrived as a full batch b , so we run e_1 using the batch semantics to produce an output batch b' . We then use the prefix p to run e_2 , returning p' and stepping to e'_2 . The entire term then outputs $\text{catpB}(b', p')$ and steps to e'_2 . Note that the pair is eliminated in the process: we step from $(e_1; e_2)$ to just e'_2 . To see why, note that, for the correctness theorem to hold, the resulting term must have type $\delta_{\text{catpB}(b',p')}(s \cdot t) = \delta_p(t)$ in context $\delta_{\text{catpB}(b,p)}(\Gamma; \Delta) = \delta_p(\Delta)$ —but this is the typing for e'_2 .

Incremental Semantics of Context Subtyping. The incremental semantics for context subtyping is conceptually similar to the batch semantics. Given $pf : \Gamma \leq \Gamma'$, we accept a prefix p of type Γ in the input, coerce it with a relation $\text{coePfx}(pf, p) \sim p'$ to a prefix of type Γ' , pass it to the premise term e , and return the result p'' . The difference comes in how we decide what term to step to. In the recursive call, the premise term e steps to e' , which is well typed in context $\delta_{p'}(\Gamma')$. Meanwhile, the overall conclusion must be well typed in context $\delta_p(\Gamma)$. These two contexts are related by a subtyping rule—intuitively, the one given by the *derivative* of the proof term pf with respect to p . To calculate this derivative, we use yet another relation $\delta_p(pf) \sim pf'$ (defined in Figure B.66 in Appendix B), which takes a prefix p of type Γ together with a proof term $pf : \Gamma \leq \Gamma'$ and produces a proof term $pf' : \delta_p(\Gamma) \leq \delta_{p'}(\Gamma')$. In the rule P-SUBCTX-1, we use this to update the proof term and step to the term $\text{subctx}(pf', e')$. However, it is sometimes the case that $\delta_p(\Gamma) = \delta_{p'}(\Gamma')$ on the nose, in which case the structural rule is no longer needed. We write this as $\delta_p(pf) \sim \perp$, and use it in P-SUBCTX-2, where we simply step to the term e' .

Incremental Semantics of Variables and Prefix Projection. The rules for a variable term $(x : s @ \Gamma(-))$ are a bit more involved, but the basic idea is straightforward: given a prefix p of type $\Gamma(x : s)$, we

$$\begin{array}{c}
\frac{p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)}{p \Rightarrow (x : s@{\Gamma}(-)) \downarrow (x : s@{\Gamma}'(-)) \Rightarrow \text{emp}_s} \text{P-VAR-1} \\
\\
\frac{p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)}{p \Rightarrow (x : s@{\Gamma}(-)) \downarrow (x : \delta_p(s)@{\Gamma}'(-)) \Rightarrow p'} \text{P-VAR-2} \quad \frac{p@{\Gamma}(-) \rightsquigarrow b}{p \Rightarrow (x : s@{\Gamma}(-)) \downarrow \text{sink}_b \Rightarrow (b)^\circ} \text{P-VAR-3} \\
\\
\frac{p \Rightarrow e_1 \downarrow e'_1 \Rightarrow p_1 \quad p \Rightarrow e_2 \downarrow e'_2 \Rightarrow p_2}{p \Rightarrow (e_1, e_2) \downarrow (e'_1, e'_2) \Rightarrow \text{parp}(p_1, p_2)} \text{P-PAR-R} \\
\\
\frac{(p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)) \vee (p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)) \quad p \Rightarrow e \downarrow e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e \downarrow \text{let}_{\Gamma'(-)}(x, y) = z \text{ in } e' \Rightarrow p'} \text{P-PAR-L-1} \\
\\
\frac{p@{\Gamma}(-) \rightsquigarrow _@p \Rightarrow e \downarrow e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e \downarrow e' \Rightarrow p'} \text{P-PAR-L-2} \\
\\
\frac{p \Rightarrow e_1 \downarrow e'_1 \Rightarrow p'}{\text{catpA}(p) \Rightarrow (e_1; e_2) \downarrow (e'_1; e_2) \Rightarrow \text{catpA}(p')} \text{P-CAT-R-1} \\
\\
\frac{b \Rightarrow e_1 \Rightarrow b' \quad p \Rightarrow e_2 \downarrow e'_2 \Rightarrow p'}{\text{catpB}(b, p) \Rightarrow (e_1; e_2) \downarrow e'_2 \Rightarrow \text{catpB}(b', p')} \text{P-CAT-R-2} \\
\\
\frac{(p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)) \vee (p@{\Gamma}(-) \rightsquigarrow \text{catpA}(_)@{\Gamma}'(-)) \quad p \Rightarrow e \downarrow e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e \downarrow \text{let}_{\Gamma'(-)}(z; x) = y \text{ in } e' \Rightarrow p'} \text{P-CAT-L-1} \\
\\
\frac{(p@{\Gamma}(-) \rightsquigarrow \text{catpB}(_, _)@_) \vee (p@{\Gamma}(-) \rightsquigarrow b) \quad p \Rightarrow e \downarrow e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e \downarrow e' \Rightarrow p'} \text{P-CAT-L-2} \\
\\
\frac{}{p \Rightarrow \text{sink} \downarrow \text{sink} \Rightarrow \text{eps}_p} \text{P-Eps-R} \quad \frac{}{p \Rightarrow () \downarrow \text{sink} \Rightarrow \text{onep}_B} \text{P-ONE-R} \\
\\
\frac{\text{coePfx}(pf, p) \sim p' \quad p' \Rightarrow e \downarrow e' \Rightarrow p'' \quad \delta_p(pf) \sim pf'}{p \Rightarrow \text{subctx}(pf, e) \downarrow \text{subctx}(pf', e') \Rightarrow p''} \text{P-SUBCTX-1} \\
\\
\frac{\text{coePfx}(pf, p) \sim p' \quad p' \Rightarrow e \downarrow e' \Rightarrow p'' \quad \delta_p(pf) \sim \perp}{p \Rightarrow \text{subctx}(pf, e) \downarrow e' \Rightarrow p''} \text{P-SUBCTX-2}
\end{array}$$

Fig. 6. Incremental semantics of Kernel λ^{ST}

use the location $\Gamma(-)$ of x within the context to traverse p and locate the sub-prefix corresponding to x . A slight complication arises from the fact that traversing p to look for the data corresponding to x might lead to three different results. First, the input prefix might include *none* of the data for the variable x , as in the prefix $p = \text{catpA}(p')$ for the context $\Gamma_0 = (y : t; (x : s; z : r))$, using the one-hole context $\Gamma(-) = (y : t; (-; z : r))$. Second, the input prefix might include a *prefix* of the data for the variable x , as in the prefix $\text{catpB}(b, \text{catpA}(p))$ for the same Γ_0 . In this case, p is the result of projecting out the data for x . And third, the input might include an entire *batch* corresponding to the variable x , as in the prefix $\text{catpB}(b, \text{catpB}(b', p))$ for Γ_0 . Here, the batch b' is the result of projecting out x .

To account for these three situations, we define a trio of relations. The four-place relation $p@Γ(-) \rightsquigarrow \perp@Γ'(-)$ holds in the first situation, when data for x has not arrived yet. Next, the relation $p@Γ(-) \rightsquigarrow p'@Γ'(-)$ holds in the second situation, where p' is the prefix corresponding to x . Last, the relation $p@Γ(-) \rightsquigarrow b$ holds when there is a batch b located at $Γ(-)$ in p . In the first two relations, $Γ'(-)$ is the one-hole context corresponding the location of x in the derivative of $Γ(x : s)$ with respect to p . These take the place of the original $Γ(-)$ in the term that $(x : s@Γ(-))$ steps to. The rules defining (slightly more general versions of) these relations can be found in Appendix B.5.

Now we need three variable rules, one for each possibility. When the prefix p includes no data for x (so $p@Γ(-) \rightsquigarrow \perp@Γ'(-)$ holds), P-VAR-1 applies. We return emp_s , the empty prefix of type s , and step to the term $(x : s@Γ'(-))$ with the updated location of x in the remaining context. If p includes a prefix p' for x , P-VAR-2 applies. We return the prefix p' and step to the term $(x : \delta_{p'}(s) @Γ'(-))$ with the updated location of x and its updated type after p' has been sent out. Last, if p includes an entire batch b for x , P-VAR-3 applies. We return $(b)^\circ$ (the result of turning b into a prefix) and step to sink_b , a term that produces no output (sink_b is defined (in Appendix B.9) by recursion on b and has type $\delta_{(b)^\circ}(s)$, as required by the correctness theorem).

Incremental Semantics of Left Rules. The prefix-to-prefix parts of the left rules for concatenation and parallel are straightforward: we accept a prefix p of type $Γ(z : s \otimes t)$, where \otimes is one of the two products, and pass it on to the continuation term.

For both connectives, the term that we step to depends on the “size” of the input prefix p . This term must be well typed in context $\delta_p(Γ(z : s \otimes t))$, but, if p is large enough, the derivative context will no longer include the connective being eliminated, in which case we must remove the `let` binding. For example, given $(z : s || t; u : r) \vdash \text{let } (x, y) = z \text{ in } u : r$, if we have $p = \text{catpB}(b, p')$, then $\delta_p((z:s||t); (u:r)) = (u : \delta_{p'}(r))$. The stream has advanced past the segment which contained the connective being eliminated, so the `let`-binding itself no longer makes sense. In the case of the left rule for parallel, this happens when the prefix p includes a batch of $z : s || t$, as shown in the example above. In the case of the left rule for concatenation, it happens when p includes either a batch of $z : s \cdot t$ or a prefix $\text{catpB}(b, p)$ that includes a batch b of s ; in this case the derivative $\delta_p(Γ(z : s \cdot t))$ looks like $Γ'(z : \delta_p(t))$ for some $Γ'(-)$; the connective has again been deleted by the derivative, and the `let`-binding is no longer required.

3.4 The Homomorphism Property and Determinism

Our incremental semantics is designed to run a stream transformer on an “input chunk” of any size. In practical terms, whenever one or more input events are ready, a transformer can take a step that consumes these events and produces the appropriate output. Later, when more of the input arrives, it can take another step to produce more output, and so on. But this flexibility has a cost, as it raises the question of *coherence*—whether we are guaranteed to arrive at the same final output depending on how we carve up a transformer’s input into a series of prefixes. Fortunately, this is indeed guaranteed.

Coherence follows from our main technical result, a *homomorphism theorem* that says running a term e on a prefix p_1 and then running the resulting term e' on a prefix p'_1 of the remaining stream produces the same end result as running e on the combined prefix.

THEOREM 3.3 (HOMOMORPHISM THEOREM). *Suppose*

- (1) $Γ \vdash e : s$
- (2) $p_1 : \text{prefix}(Γ)$
- (3) $p'_1 : \text{prefix}(\delta_{p_1}(Γ))$
- (4) $p_2 : \text{prefix}(s)$

$$(5) p'_2 : \text{prefix}(\delta_{p_2}(s))$$

$$(6) p_1 \Rightarrow e \downarrow e' \Rightarrow p_2$$

$$(7) p'_1 \Rightarrow e' \downarrow e'' \Rightarrow p'_2.$$

Then $p_1 \cdot_{\Gamma} p'_1 \Rightarrow e \downarrow e'' \Rightarrow p_2 \cdot_s p'_2$.

The operation $p \cdot_s p'$ here is type-indexed *prefix concatenation*, which takes a prefix p of type s and a prefix p' of type $\delta_p(s)$ and produces the prefix of type s that is first p , and then p' . Formally, we define this as a 4-place partial inductive relation $p \cdot_s p' \sim p''$, which is defined when p and p' have types s and $\delta_p(s)$, respectively. The operation $p \cdot_{\Gamma} p'$ does the same for prefixes of contexts. See Appendix B.4 for details.

The homomorphism theorem not only justifies the prefixes-at-a-time perspective; it also implies deterministic processing of parallel streams. Intuitively, determinism states that the results of a stream transformer do not depend on the particular order in which parallel data arrives. We formalize this through the following scenario. Suppose $\Gamma, \Gamma' \vdash e : s$ is a term with two parallel contexts serving as its input, and suppose that p_1 is a prefix of type Γ and p_2 a prefix of type Γ' . There are two different ways of running e on this data. One is to first run e on p_1 (passing the empty prefix for the Γ' component) and then run the resulting term on p_2 (with an empty prefix for Γ). The other does the opposite, first running e on p_2 and then running the resulting term on p_1 . Determinism says that these strategies produce equal results. It is proved by observing that the homomorphism theorem guarantees that each of these options is equivalent to running e on $\text{parp}(p_1, p_2)$.

THEOREM 3.4 (DETERMINISM). *Suppose*

$$(1) \Gamma, \Gamma' \vdash e : s$$

$$(2) \text{parp}(p_1, p_2) : \text{prefix}(\Gamma, \Gamma')$$

$$(3) \text{parp}(p_1, \text{emp}_{\Gamma'}) \Rightarrow e \downarrow e_1 \Rightarrow p'_1 \text{ and } \text{parp}(\text{emp}_{\Gamma}, p_2) \Rightarrow e_1 \downarrow e_2 \Rightarrow p'_2$$

$$(4) \text{parp}(\text{emp}_{\Gamma}, p_2) \Rightarrow e \downarrow e'_1 \Rightarrow p''_1 \text{ and } \text{parp}(p_1, \text{emp}_{\Gamma'}) \Rightarrow e'_1 \downarrow e'_2 \Rightarrow p''_2.$$

Then $e_2 = e'_2$ and $p'_1 \cdot_s p'_2 = p''_1 \cdot_s p''_2$.

Two crucial observations make the proof work. First, prefixes are morally equivalence classes of sequences of stream elements, up to the possible reorderings defined by their type [57]. When presented syntactically, these equivalence classes serve as *normal forms* for all the possible interleavings of the same data. The homomorphism theorem then guarantees that these normal forms are processed compositionally, and so are independent of the actual temporal ordering of parallel data—it suffices to compute on aggregate normal forms.

4 FULL λ^{ST}

We now introduce—more briefly—the remaining types and terms of λ^{ST} that are not part of Kernel λ^{ST} . Along the way, we explore both type-theoretic and logical consequences of the concepts in λ^{ST} , as well as features that are more practically motivated.

4.1 Cut

A logically inclined reader might have wondered about the absence of a *cut rule* in Kernel λ^{ST} . In sequent calculus-style presentations of logics, a cut rule supports proofs with “lemmas”: we can prove a proposition and then use it as a premise to help prove another.³ Transported along the proofs-as-programs principle, cut corresponds to *sequencing*: running one transformer and then

³Traditionally in sequent calculi, the cut rule is introduced only to be immediately shown to be admissible. We have reason to believe that the cut rule in λ^{ST} will indeed be admissible—Frumin [32] has proven that BI plus an arbitrary set of structural rules admits cut—but we have not proven it ourselves for the specific set of rules and constructs in λ^{ST} . Because the point of

passing its output as an argument to another. Given a transformer term e_1 whose output type is used in the input of another term e_2 , we can cut them together to form a single term $\text{let } x = e_1 \text{ in } e_2$ that operates as the sequential composition of e_1 followed by e_2 .

One way to think about the term $\text{let } x = e_1 \text{ in } e_2$ is that it “internalizes” the sequential composition of stream transformers, in the following sense. In an eventual high-level stream-processing language based on λ^{ST} , a programmer might write one transformer to describe the behavior of each node in a distributed stream processing system and wire these together into a dataflow graph, connecting the output of one transformer to the inputs of another by way of links that might traverse the network. This is an *external* form of composition, occurring outside the bounds of the calculus. The term $\text{let } x = e_1 \text{ in } e_2$ *internalizes* this composition by forming a single core term that behaves like the external sequential composition of e_1 and e_2 . This step from external to internal composition of transformers corresponds to *operator fusion*, a key optimization technique in the stream processing literature.

The typing rule for cut is shown to the right. If e_1 has type s in context Δ , and e_2 has type t in a context $\Gamma(x : s)$ with a variable of type s , we can form the cut term $\text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2$, which has type t in context $\Gamma(\Delta)$. (The one-hole context $\Gamma(-)$ is included as an annotation on the term for the same reason that it is included in the terms for concatenation and parallel let-bindings: it is needed at runtime.)

Running a cut incrementally intuitively takes the following steps: (1) accept an input prefix p and project out the component p_Δ corresponding to the subcontext Δ , (2) run p_Δ through e_1 to get a prefix p_s , (3) bind p_s to x by substituting it back into p where p_Δ was to get a prefix of type $\Gamma(x : s)$, and (4) run that through e_2 , producing the final output. (Full details can be found in Appendix B.10.) An important point to note is that this semantics is non-blocking: even if no data for Δ has arrived, we still run e_2 , potentially producing output.

4.2 Sums

The first type constructor that appears in λ^{ST} but not in Kernel λ^{ST} is the sum type, written $s + t$. Sums in λ^{ST} are *tagged unions*: a stream of type $s + t$ is either a stream of type s or a stream of type t , and a consumer can distinguish between the two. Streams of type s are not the same as streams of type $s + s$, and streams of type $s + t$ are equivalent to but not the same as streams of type $t + s$. Operationally, a producer of a sum stream sends a tag bit prior to sending the rest of the stream, to say which side it is choosing. Conversely, a consumer of $s + t$ first reads the bit to learn if it should expect a stream of type s or one of type t and then gets a stream of that type.

Batches of $s + t$ are either a tagged batch of s , written $\text{sumbA}(b)$, or a tagged batch of t , written $\text{sumbB}(b)$. A prefix of $s + t$ can be a prefix of one of s or one of t , written $\text{sumpA}(p)$ or $\text{sumpB}(p)$. A third possible prefix is sumpEmp , the empty prefix of type $s + t$, which does not even say which way the rest of the stream is going to go. The derivatives with respect to these prefixes are defined by (a) the empty prefix taking nothing off the type ($\delta_{\text{sumpEmp}}(s + t) = s + t$) and (b) the two injections reducing to taking the derivative of the corresponding branch of the sum ($\delta_{\text{sumpA}(p)}(s + t) = \delta_p(s)$ and $\delta_{\text{sumpB}(p)}(s + t) = \delta_p(t)$).

The typing rules for sums are the normal injections on the right (SUM-R-1 and a symmetric rule SUM-R-2) and a case analysis principle on the left

$$\frac{\Gamma \vdash e : s}{\Gamma \vdash \text{inl}(e) : s + t} \text{SUM-R-1} \quad \frac{\Gamma(x : s) \vdash e_1 : r \quad \Gamma(y : t) \vdash e_2 : r}{\Gamma(z : s + t) \vdash \text{case}_{\Gamma(-),s,t,r}(z, x.e_1, y.e_2) : r} \text{SUM-L}$$

cut elimination is to enable effective proof search, whereas we are most interested in the calculus from a programming perspective, we will not dwell on this point here.

(SUM-L). The right rules operate by prepending their respective tag and then running the embedded term. The left rule performs case analysis: if the incoming stream is a left, then it is processed with e_1 ; if a right, then e_2 .

The batch semantics for the PLUS-L and PLUS-R rules are straightforward, but the incremental semantics is a bit more subtle. When a prefix arrives at a sum case term, the incremental semantics must dispatch on the tag that says if the stream is a left or a right. But that prefix might not include a tag (if it is `sumEmp`) or it might not include any value for the variable at all (as in the rule P-VAR-1 discussed earlier); in these cases, we have no way of determining which branch to run. The solution to this conundrum is to run neither! Instead, we hold on to the prefix, adding incoming data to the buffer until the tag arrives. Once we get a prefix that includes the tag, we continue by running the corresponding branch. Note that this buffering is necessarily a blocking operation.⁴ Even if data arrives on part of the input that is in parallel with the variable being scrutinized, the case analysis must block until the tag arrives. This requires a slightly generalized typing rule that includes a *buffer* $p : \text{prefix}(\Gamma(z : s + t))$ of the context type in the term. As prefixes arrive, we append to this buffer until we get the tag. When the buffer is empty ($p = \text{emp}_{\Gamma(z:s+t)}$, e.g., before the transformer starts running), the generalized rule simplifies to the SUM-L rule in shown above. Full details on the generalized rule and the incremental semantics for case analysis can be found in Appendix B.10.

4.3 Star

Full λ^{ST} also includes a type constructor for unbounded streams,⁵ written s^* ; it describes a stream that contains zero or more sub-streams of type s , in sequence. This formulation is inspired by the Kleene star from the theory of regular languages: zero or more concatenations of a stream type s with itself. In regular languages, r^* is equal to $\varepsilon + r \cdot r^*$. Interpreted in the language of stream types, this equation says that streams of type s^* are either empty (ε) or a stream of type s , followed by another stream of type s^* —i.e., s^* can be understood as the least fixpoint of the stream type operator $x \mapsto \varepsilon + s \cdot x$. It is a list, whose elements are separated in time. The definitions of batches, prefixes, and typing rules for star all follow from this perspective.

For example, since a batch of $s + t$ is either a batch of s or a batch of t , a batch of $\varepsilon + s \cdot s^*$ is either a batch of ε (the unit value), or a batch of $s \cdot s^*$, which is a pair of a batch of s and another batch of s^* . This gives us the possible batches of type s^* : the empty batch $[]$ and the “cons batch” $b :: b'$.

Similarly, we cook our definition of the prefixes of type s^* so that $\text{prefix}(s^*) = \text{prefix}(\varepsilon + s \cdot s^*)$. The empty prefix of type s^* , written `stpEmp`, is effectively the empty prefix of the sum that makes up s^* . The second form of prefix—the “done” prefix of type s^* —is written `stpDone`. It corresponds to the left injection of the sum, and receiving it means that the stream has ended. Note that, despite containing no data, this prefix is not empty: it conveys the information that the stream is complete. The final two cases correspond to the right injection of the sum, i.e., a prefix of type $s \cdot s^*$. This is either `stpA(p)`, with p a prefix of s , or `stpB(b, p)`, with b a batch of s and p another prefix of s^* .

For derivatives, the empty prefix leaves the type as is ($\delta_{\text{stpEmp}}(s^*) = s^*$). Because no data will arrive after the done prefix, the derivative of s^* with respect to `stpDone` is ε , the type of the empty stream. In the case for `stpA(p)`, after some of an s has been received, the remainder of s^* looks like the remainder of the first s followed by some more s^* , so the derivative is defined as $\delta_{\text{stpA}(p)}(s^*) = (\delta_p(s)) \cdot s^*$. Finally, $\delta_{\text{stpB}(b,p)}(s^*) = \delta_p(s^*)$.

The typing rules for star are again motivated by the analogy with lists. There are right rules for `nil` and `cons`, and a case analysis principle for the left rule. The “nil” rule STAR-R-1 corresponds to

⁴Depending on the rest of the context, it could also require unbounded memory! In Section 7, we discuss how the type system might allow us to statically detect such space leaks.

⁵We do not need to distinguish between unbounded finite streams and “truly infinite” ones, because our incremental operational semantics cannot distinguish between a stream that goes on forever and a finite but as yet incomplete stream.

the left injection of the sum $s^\star = \varepsilon + s \cdot s^\star$. From any context, we can produce s^\star by simply ending the stream. The “cons” rule STAR-R-2 is the right injection. From a context $\Gamma; \Delta$, we can produce an s^\star by producing one s from Γ and the remaining s^\star from Δ . Operationally, this should run the same way as the CAT-R rule: by first running e_1 using the Γ part of the context to produce the first s , and then switching over to running e_2 to produce the tail s^\star once we start to receive Δ -data.

$$\frac{}{\Gamma \vdash \text{nil} : s^\star} \text{STAR-R-1}$$

$$\frac{\Gamma \vdash e_1 : s \quad \Delta \vdash e_2 : s^\star}{\Gamma; \Delta \vdash e_1 \cdot e_2 : s^\star} \text{STAR-R-2}$$

The STAR-L rule is a case analysis principle for star streams: either a stream is empty, or it includes one s followed by an s^\star . The fact that the head s will come first and the tail s^\star later tells us that the variables $x : s$ and $xs : s^\star$ should be separated by a semicolon.

$$\frac{\Gamma(\cdot) \vdash e_1 : r \quad \Gamma(x : s; xs : s^\star) \vdash e_2 : r}{\Gamma(z : s^\star) \vdash \text{case}_{\Gamma(-),s,r}(p; z, e_1, x.xs.e_2) : r} \text{STAR-L}$$

The incremental and batch semantics for the right rules are straightforward: the rules for STAR-R-1 are like those for EPS-R, while the rules for STAR-R-2 are like those for CAT-R. But like PLUS-L rule, the STAR-L rule needs to wait until it knows if the stream will be empty or include at least one s before it can pick a continuation to run. For this reason, STAR-L must also be generalized to include a prefix buffer, and its semantics are also blocking in the same way as PLUS-L. For full details on the batch and incremental semantics for star, see Appendix B.10.

4.4 Recursion

At this point, although the λ^{ST} type system now includes the type s^\star , we cannot write very interesting transformers over s^\star streams: we need a way to define transformers recursively. Because λ^{ST} has no function types, we accomplish this by adding explicit term-level recursion and recursive call operators.

The term $\text{fix}(e)$ denotes a recursive transformer with body e . Recursive calls are made with a constant term rec , which refers back to e . This back-reference works in the same way that uses of the variable x in the body of a traditional fix point $\text{fix}(x.e)$ refer to the term $\text{fix}(x.e)$ itself. This function-free approach is inspired by the concept of *cyclic proofs* [14, 20, 24] from proof theory, where derivations may refer back to themselves. Alternatively, one can think of this construction as defining our terms and proof trees as infinite *coinductive* trees; then the term-level fix operator defines terms as *cofixpoints*. Section 5 shows how these recursion operators can be used to define recursive transformers that process star streams like `map` and `filter`.

To type these new terms, we extend the Kernel λ^{ST} typing judgment to include a *recursion signature* $\Sigma ::= \text{NoRec} \mid \Gamma \rightarrow s$, written on the turnstile: $\Gamma \vdash_\Sigma e : s$. The point of this addition is to ensure that recursive calls are well typed: if a recursive transformer $\text{fix}(e)$ is typed in context Γ with type s , then all instances of rec in e must also occur in a position where a program from Γ to s is expected. A recursion signature of the form NoRec says that the term being typed is not in the body of a recursive function, while a recursion signature of the form $\Gamma \rightarrow s$ means that we are under a `fix`, where the recursive transformer being defined takes Γ to s . The REC rule shown on the right says that, if we are in the body of a recursive program from Γ to s , we can make a recursive call. The FIX rule defines a recursive transformer from Γ to s by typing the body e with a recursion signature $\Gamma \rightarrow s$. (Note that FIX clobbers any existing recursion signature, so using the REC rule inside a multiply-nested recursive program refers to the innermost one. This means that λ^{ST} does not support mutually recursive definitions. (We have yet to run up against

$$\frac{}{\Gamma \vdash_{\Gamma \rightarrow s} \text{rec} : s} \text{REC}$$

$$\frac{\Gamma \vdash_{\Gamma \rightarrow s} e : s}{\Gamma \vdash_{\Sigma} \text{fix}(e) : s} \text{FIX}$$

this restriction in practice, but generalizing to allow nested recursions is straightforward, though notationally heavy.)

In both the incremental and batch semantics, we only run transformers that are well typed with a closed (NoRec) recursion signature. To run $\text{fix}(e)$ in either semantics, we simply unfold the recursion one step, substituting $\text{fix}(e)$ itself for instances of use in the body e . Naturally, this can

$$\frac{p \Rightarrow e [\text{fix}(e)/\text{use}] \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{fix}(e) \Downarrow^{n+1} e' \Rightarrow p'} \text{P-Fix}$$

Fig. 7. Fix Incremental Semantics

lead to non-termination—for example, there are no finite derivations of $p \Rightarrow \text{fix}(\text{rec}) \Downarrow e \Rightarrow p'$ for any p, e, p' , since $\text{fix}(\text{rec})$ unfolds to itself.⁶ To bound the depth of evaluation, we *step index* both semantic judgments by adding a fuel parameter that decreases when we unfold a use of Fix. The incremental semantic judgment then looks like $p \Rightarrow e \Downarrow^n e' \Rightarrow p'$ and means that when we run e on p , it steps to e' producing p' and unfolding at most n uses of fix along the way. The incremental semantics rule for Fix is presented in Figure 7; all the other rules in the incremental and batch semantics for λ^{ST} need to be updated to propagate the step index compositionally. The only place the step index decreases is in P-Fix, and all base cases run in any amount of gas to ensure that the step index is monotone.

The inclusion of a step index now means that there are well-typed terms about which the λ^{ST} semantics say nothing at all. In particular, an “infinite generator” term $\cdot \vdash_{\text{NoRec}} \text{rec}(\cdot) : \text{use} : 1^\star$, which runs forever and should produce an infinite sequence of unit values, has no meaning in λ^{ST} . Semanticists may find this behavior odd, but it mimics the incremental semantics of present-day stream processing systems, which wait for a step of computation to terminate before sending out any of its results.

The inclusion of recursion also requires an update to the correctness theorem: if a well typed term takes a step on a well-typed input *using some amount of gas*, then the output and resulting term are also well typed.

THEOREM 4.1 (INCREMENTAL SEMANTICS CORRECTNESS). *If $\cdot \mid \Gamma \vdash_\emptyset e : s$, and $p : \text{prefix}(\Gamma)$, and $p \Rightarrow e \Downarrow^n e' \Rightarrow p'$, then $p' : \text{prefix}(s)$ and $\cdot \mid \delta_p(\Gamma) \vdash_\emptyset e' : \delta_{p'}(s)$*

A similarly updated theorem statement for the homomorphism theorem—and many other theorems and lemmas used in the proofs of the correctness and homomorphism theorems—can be found in Appendix B.10.

4.5 Stateful Transformers

In the λ^{ST} typing judgment $\Gamma \vdash e : s$, the variables in Γ range over *future values* that have yet to arrive at the transformer e . The ordered nature of semicolon contexts means that variables further to the right in Γ correspond to data that will arrive further in the future. This imposes a strong restriction on programming: if earlier values in the stream are to be used at all, they must be used *before* later values, and once a value in the stream has “gone past” there is no way to refer to it again.

This limitation can be seen as arising from an asymmetry in the design of Kernel λ^{ST} . By using variables from the Γ context, a term e can refer to values that will arrive in the future; but it has no way of referring to values that have arrived in the *past*.

From a programming perspective, referring to variables from the past requires *state*. Many important streaming functions (e.g., `map` and `filter`) are stateless, but many others (e.g., “running

⁶Viewed logically, this means that λ^{ST} is unsound as a proof system: $\text{rec}(\text{use})$ is a proof of any sequent. Cyclic proof systems usually ensure soundness by imposing a guardedness condition [14] which requires certain rules be applied before a back-edge can be inserted in the derivation tree. Because we are not primarily concerned with λ^{ST} as a logic at the moment, we leave a guardedness condition to future work.

sums”) do require state. Moreover, state isn’t *just* about referring to values that have arrived previously: many stateful transformers must maintain accumulators and other scratch values that were never actually received on an input stream.

What types do streams from the past have? We saw in Section 3 that values of completed streams are batches. And the batches of type $s||t$ and of type $s \cdot t$ are essentially the same, since the partial order structure of the data no longer matters. For example, once a complete stream of type $(\text{Int}^*||\text{Int}^*) \cdot \text{Int}^*$ has been received, we can regard the data that has been received as a value of the conventional type $(\text{list}(\text{Int}) \times \text{list}(\text{Int})) \times \text{list}(\text{Int})$ from the simply typed lambda-calculus (STLC). This suggests that, while parts of streams that will arrive in the future have stream types, parts of streams that have arrived in the past should have standard STLC types.

Accordingly, we define an operation called *flatten*, written $\langle s \rangle$, that transforms stream types into STLC types. The important cases are $\langle s \cdot t \rangle = \langle s||t \rangle = \langle s \rangle \times \langle t \rangle$, and $\langle s^* \rangle = \text{list}(\langle s \rangle)$.

We then extend the typing judgment of λ^{ST} to include a second context, Ω , called the *historical context*, writing $\Omega \mid \Gamma \vdash_{\Sigma} e : s$. The historical context is fully structural: $\Omega ::= \cdot \mid \Omega, x : A$, where the types A are drawn from some set of STLC types including at least products, sums, a unit, and a list type. Operationally, the historical context behaves like a standard context in a functional programming language: at the top level, terms to be run with the batch or incremental semantics must be typed in an empty historical context; at runtime, historical variables get their values by way of substitution.

Rather than giving a set of ad-hoc rules for manipulating values from the historical context, we parameterize the λ^{ST} calculus over an arbitrary language with terms M , typing judgment $\Omega \vdash M : A$, and terminating big-step semantics $M \downarrow v$. We call any such fixed choice of language the *history language*. Programs from the history language can be embedded in λ^{ST} programs using the HISTPGM rule, which says that a historical program $M : \langle s \rangle$ that accesses the historical context Ω can be used in place of a λ^{ST} term of type s . Operationally, as soon as any prefix of the input arrives, we simply run the embedded historical program to completion and yield the result as its stream output (after converting it into a batch of type s).

How does information get added to the historical context? Intuitively, a variable in Γ (a stream that will arrive in the future) can be moved to Ω , where streams that have arrived in the past are saved, by waiting long enough for the future to become the past! Formally, we define an operation called “wait,” which allows us to specify part of the incoming context and *block* this subcomputation until that part of the input stream has arrived in full. Once it has, we can bind it to the variables in the historical context and continue by running e .

The WAIT rule encodes the typing content of this behavior. It allows us to specify a subcontext Δ of the input, and then flatten and move it to the historical context so that the continuation e can refer to its variables as historical variables. Semantically, this works by *waiting* for Δ : buffering in prefixes of the context until a whole batch of Δ has arrived. Once we have a batch, we substitute it into e and continue running the resulting term.⁷ This buffering is implemented the same way as in the left rules for plus and star, by generalizing the typing rule to include an explicit prefix buffer. The generalized typing rule for WAIT as well as the incremental and batch semantics for both constructs can be found in Appendix B.8 and Appendix B.10.

⁷The semantics of the WAIT rule is reminiscent of the “blocking reads” of Kahn Process Networks, where every read from a parallel stream blocks all other reads to ensure determinism. Here, we choose a subcontext and block the rest of the program until it is complete and in memory. For a subcontext Δ that is a single variable, this is essentially the same.

Most of the remaining typing rules in λ^{ST} change only by adding an Ω to the typing judgment everywhere; the historical context just goes along for the ride. The only meaningful change is in the `FIX` and `REC` rules. First, recursion signatures change to include a historical context: $\Sigma ::= \text{NoRec} \mid (\Omega \mid \Gamma \rightarrow s)$. Then, when defining a recursive transformer, the programmer must set up the Ω context with any accumulators or auxiliary data that the program needs and provide initial values for each. Then, when making recursive calls with `REC`, the programmer decides how those values those values will be updated before continuing with the next iteration of the program.

The syntax for recursive calls thus becomes `rec @ [M1, ..., Mn]`, where the M_i are historical programs computing updated values of the accumulators (i.e., all of the variables in the historical context in the recursion signature). Similarly, the recursive definition syntax gets extended to `fix (x1 : A1, ..., xn : An) @ [M1, ..., Mn] .e`, where the $x_i : A_i$ define the shape of the historical context in the body of e and the M_i are its initial values. See Appendix B.8 for details.

5 EXAMPLES

In this section, we show how λ^{ST} addresses the problems of type-safe programming with temporal patterns and deterministic processing of parallel data from Section 2 and how some other characteristic streaming idioms are expressed in λ^{ST} .

The process of encoding any of these common streaming idioms in λ^{ST} begins by determining the input and output types. Because these idioms are from the homogeneously typed world, we must work to “extract the stream types” from the situation. Then, we build a term of that type. In many cases, we will see that the high-level style of λ^{ST} means that the streaming programs we derive strongly resemble well-known standard functional programming idioms. But while the programs look like high-level functions computing over entire streams at once, the incremental semantics allows them to be run event by event. We’ve tested all of the incremental semantics on all of the examples in this section with a prototype interpreter, written in Haskell: see Appendix A and in the supplementary materials for details about our implementation.

Many of the idioms we’ll discuss are conceptually higher-order functions, like `map`, `filter`, and `fold`. But λ^{ST} is a first-order language: the calculus has no function types. We treat higher-order functions as macros: transformers like `map`, `filter`, and `fold` are all parameterized by a term which represents the function argument. In Section 8, we’ll discuss a high level language design which could support this.

We present all examples in this section with a syntax which is a sugared version of our core terms. First, type and context annotations on variables, left rules, cut and wait are omitted, and structural rules are written silently (without their explicit terms `subctx (pf, e)`). We also use a syntax `wait x then e end` for the `waitΔ(e)` operator when Δ is a single variable x and e the continuation. Embedded historical programs are written surrounded with angle brackets: $\langle M \rangle$. The syntax for defining recursive programs is `fix with (y1 = M1, ..., yn = Mn) . e`. This defines a recursive transformer with accumulator variables y_1, \dots, y_n which take initial values M_1, \dots, M_n . For recursive calls, we use an argument-passing syntax `rec(xs)` to make clear which subcontext a recursive call is being performed on, and write square brackets `rec(x1, ..., xM)[acc1, ..., accN]` to update the values of any accumulators in the historical context. Last, we sometimes use a natural deduction-style elimination form (like `let (x, y) = e1 in e2`) instead of a left rule: this can be de-sugared into a cut followed by a use of the corresponding left rule.

We present the full core terms for all these examples in Appendix C. In Section 8, we’ll discuss potential methods for elaborating a higher-level syntax like this one down into λ^{ST} .

Map. Given a transformer from s to t , we can lift it to a transformer from s^* to t^* . Note that this type is more general than the standard `map` function on homogeneous streams, which has type

$(a \rightarrow b) \rightarrow (\text{Stream } a \rightarrow \text{Stream } b)$: our types s and t can be arbitrary stream types: they need not be singletons.

The map transformer is parameterized by a transformer e from s to t . The code for `map` is essentially identical to the traditional functional code.

```
map (e : s → t) (xs : s*) : t* =
  fix. case xs of
    | nil => nil
    | y :: ys => e(y) :: rec(ys)
```

Filter. Similarly, given a “predicate” transformer e from s to $1 + 1$ (i.e., Booleans), we can filter an incoming stream of type s^* to include the elements of type s (and which may consist of many events) for which e outputs `inl()`.

The code is identical to the standard functional `filter`, yet it runs incrementally: as soon as it sees that y passes the filter, the function will start to forward it along. See Appendix C.3 for details.

Singletons, Head, Tail. In the homogeneous model, stream types are always conceptually unbounded. But in practical situations, some streams are only expected to contain a single element—a constraint that cannot be expressed with homogeneous streams. Using stream types, we can write stream transformers which are statically known to only produce a single output. For example, the “head” function is trivially expressible in the same manner as `head` on lists. (Exercise: try writing the term for `tail` on star streams. This requires a use of `wait` and an accumulator argument like in `fold`, as shown below.)

```
head (xs : s*) : ε + s =
  fix. case xs of
    | nil => inl(sink)
    | y :: ys => inr(y)
```

Using stream types, we can write stream transformers which are statically known to only produce a single output. For example, the “head” function is trivially expressible in the same manner as `head` on lists. (Exercise: try writing the term for `tail` on star streams. This requires a use of `wait` and an accumulator argument like in `fold`, as shown below.)

Fold. In λ^{ST} , we can express a fold which returns only the final accumulator, as opposed to traditional streaming folds which are *running* folds, outputting a stream of all intermediate results (we can also do the running fold; see Appendix C.4). `Fold` maintains an accumulator of type $\langle t \rangle$ in the historical context, which gets updated by a streaming step function $e : \langle t \rangle \mid s \rightarrow t$ that takes a stream argument of type s and a historical accumulator argument of type $\langle t \rangle$, and produces a t . Then, the whole fold takes a stream xs of type s^* and an initial accumulator `init` of type $\langle t \rangle$ and produces a stream of accumulator values t^* .

```
fold (e : ⟨t⟩ | s → t) (xs : s*) : t =
  fix with (y = m0).
  case xs of
    | nil => y
    | x' :: xs' =>
      let y' = e(x')[y] in
      wait y' then rec(xs')[y'] end
```

The code for `fold` cases on the incoming stream. If the stream is empty, we return the accumulator y . If there is a head, it uses the term e to compute the next accumulator y' from the head x' and the previous accumulator value. Then, we pull y' into the historical context with a `wait` and recurse to handle the tail of the input stream with the updated accumulator y' .

Brightness Levels. The “input protocol” from the brightness-levels example in Section 2 can be encoded as the type $(\text{Int} \cdot \text{Int}^*)^*$: a stream of nonempty sequences of `Int`s, which represent the nonempty runs of light levels greater than some threshold. Writing a transformer that assumes the input invariant is easy. For example, to average the brightness-levels in each run, we can simply `map` an average operation—which takes $\text{Int} \cdot \text{Int}^*$ to Int —over the incoming stream to produce a stream Int^* of averages: $\cdot \mid xs : (\text{int} \cdot \text{int}^*)^* \vdash \text{map}(\text{avg})(xs) : \text{int}^*$.

By contrast, with a homogeneous stream type like $(\text{Start} + \text{Int} + \text{End})^*$, this operation would need to be written in a low-level and stateful manner, remembering the current run of `Int`s until an `End` event arrives, averaging, and handling the divide-by-zero error which could in principle (although not in practice) occur if no `Int`s arrived between a `Start` and an `End`. In λ^{ST} the program is instead high-level and functional, describing what happens to each run of `Int`s above the threshold.

The actual per-run average operation can be defined by simply waiting on the whole run to arrive, then computing the average with an embedded historical program, as shown to the right.

The final thresholding operation—which takes Int^* and produces the runs of elements above the threshold $(\text{Int} \cdot \text{Int}^*)^*$ —is somewhat complex because it requires statefully parsing the input stream, depending on if we're currently on a run of elements above the threshold or not. A certain amount of stateful logic for constructing the thresholded stream is unavoidable: similar code would be required in a language based on homogeneous streams. But in λ^{ST} , type safety guarantees that (1) the transformer *does in fact* output a stream which adheres to the protocol, and (2) the downstream transformer does not have to replicate this parsing logic. The code can be found in Appendix C.5.

```
avg (x : int ; xs : int*) : int =
  wait x,xs then
    ⟨(x + sum xs) / (length xs + 1)⟩
  end
```

Side Outputs & Error Handling. A common streaming idiom is the use of “side outputs” for reporting errors. Operations include extra output streams where error messages are sent as they arise at runtime. In most frameworks, side outputs are a second-class mechanism: the error streams cannot be transformed or used in a manner other than dumping them to a log somewhere. λ^{ST} provides a first-class account of side outputs, by encoding them as a parallel output type. A function $s \rightarrow t$ that may produce errors of type e can have type $s \rightarrow t \parallel e^*$. Alternately, errors can be handled inline in the traditional typed functional manner, using a sum type $s + e$.

Partitioning. Partitioning is a crucial streaming idiom where a homogeneous stream of data is split into two or more parallel streams, which are then routed to different downstream nodes in the dataflow graph, thus exposing parallelism and increasing potential throughput. Appendix C.6 shows how two different partitioning strategies can be implemented in λ^{ST} : a *round-robin partitioner*, which fairly distributes an incoming stream of type s^* into a parallel pair of streams $s^* \parallel s^*$ by by sending the first element to the left branch, the second to the right, and so on, and a *hash-based partitioner*, which routes stream elements based on a hash of their data modulo the number of downstream processing nodes.

Windowing and Punctuation. *Windowing* is a core concept in stream processing systems, where common aggregation operations like moving averages or sums are defined over “windows”—groupings of consecutive events, gathered together into a set. In λ^{ST} , these transformers are just maps over a stream whose elements are windows. Given a per-window aggregation transformer f that takes an individual window s^* to a result type t plus a “windowing strategy” win which takes a stream r^* and turns it into a stream of windows s^{**} , we can write the windowed operation as $\text{xs} : r^* \mid\text{-} \text{map}(f)(\text{win}(\text{xs})) : t^*$. In Appendix C.7, we define both sliding and tumbling size-based window operators, as well as punctuation-based windowing, where windows are delimited by punctuation marks inserted into the stream.

6 COMPILING TO HOMOGENEOUS STREAMS

Our new logical foundation for stream processing represents a significant departure from the traditional homogeneous view. The values that λ^{ST} programs operate over are highly structured syntactic objects, not the simple sequences that modern stream processing systems are designed to handle. Given this, one might wonder whether we have thrown the baby out with the streamwater (as it were): have we given up the distributed fault tolerance and message delivery guarantees that mature implementations have worked so hard to achieve? Fortunately, we have not. We prove in this section that stream types can be encoded using traditional homogeneous streams. This result implies that we can build a distributed stream processing application by writing a λ^{ST} transformer to run at each node, equipping each with adapters that convert between the structured prefixes expected by the λ^{ST} incremental semantics and the homogeneous streams of tagged events used by an industrial-strength stream processing substrate.

The key idea is that a structured λ^{ST} stream with an arbitrarily complex type can be represented as a homogeneous, totally ordered sequence over a fixed alphabet of *events*. Sequences of events can be translated into λ^{ST} prefixes and vice versa. Then, given a λ^{ST} term e and a sequence of events xs , we can turn xs into a prefix p , run e with p , and turn the output back into an event sequence.

Translating from prefixes to event sequences is a kind of serialization: we go from the structured prefix representation to a wire format in the form of a sequence of tagged events. Formally, this translation is nondeterministic: a parallel pair prefix $\text{parp}(p_1, p_2)$ can be serialized to any tagged interleaving of the serializations of p_1 and p_2 . The other direction, going from event sequences to prefixes, consists of deserializing the unstructured event sequence representation into a structured one. Crucially, deserializing any event sequence that is a serializations of a given prefix will yield that prefix back.

The grammar of events is as follows:

$$x ::= \text{oneev} \mid \text{parevA}(x) \mid \text{parevB}(x) \mid +_{\text{puncA}} \mid +_{\text{puncB}} \mid \cdot_{\text{punc}} \mid \text{catevA}(x)$$

Like prefixes, events have types. Intuitively, an event of type s is an element that could arrive as the *first* element in a stream of type s . We write the typing relation as $x : \text{event}(s)$ (the rules are given in Appendix B.11 and explained intuitively below). Then, the derivative $\delta_x(s) \sim s'$ means that s' is the type of streams that follow event x in a stream of type s . We lift the definition of events to sequences by saying that a sequence of events xs has type s if the first event xs_0 has type s , the second event xs_1 has type $\delta_{xs_0}(s)$, and so on.

The unique event of type 1 is `oneev`, and there are no events of type ε . The events `parevA`(x) and `parevB`(x) have type $s \parallel t$ when x is an event of type s and t , respectively—i.e., an event of a parallel pair type is a tagged event of one of the two types. There are two events of type $s + t$, corresponding to the first bit that says the stream will be of type s (`+puncA`) or of type t (`+puncB`). These events are *punctuation* [60]: they carry no data and serve only to inform a transformer of the choice of left or right. The derivatives with respect to these events reduce the stream to the corresponding branch: $\delta_{+_{\text{puncA}}}(s + t) = s$, and vice versa.

Finally, an event of type $s \cdot t$ can be one of two possibilities. The first is an event `catevA`(x), where x is an event of type s . The second is `\cdotpunc`, which is a punctuation mark which says that we're done with the s part of $s \cdot t$, and are ready to start with the t part. The derivative $\delta_{\cdot_{\text{punc}}}(s \cdot t) = t$ moves the stream past the s part.

Moreover, the set of events that will ever be needed to represent a stream of a given type is *finite*: for a fixed type s , the size of the possible events that could arrive on a stream of type s is bounded.

THEOREM 6.1 (BOUNDED EVENT SIZE). *For all s , there is some N such that for any $xs : \text{events}(s)$ and any $x \in xs$, we have that $|x| \leq N$, where $|\cdot|$ denotes the size of the AST.*

Serialization and Deserialization. Serializing a prefix p of type s into a sequence of events is written with a relation $p \dagger^s xs$, where p and s are inputs and xs is an output. For a given p and s , there may be many xs for which this holds; in particular, all possible interleavings of any parallel data in p give valid serializations. To deserialize, we use a relation going the other way, written $xs \hookrightarrow^s p$. This one is deterministic but not injective: each event sequence uniquely determines a prefix, but the same prefix may result from many sequences of events. Serialization and deserialization round-trip in the expected way: serializing a prefix to an event sequence and then deserializing it yields the same prefix that we started with.

Homogeneous Stream Transformers. Now we can finally describe how we compile a λ^{ST} term to a homogeneous stream transformer—that is, an element of the greatest set H satisfying the equation $H \cong \text{Seq}(X) \rightarrow 1 + 2^{\text{Seq}(X)} \times H$, where X is the set of all events and $\text{Seq}(A)$ is the set of finite sequences of elements of A . An element h of H is a nondeterministic *event handler*.

Given an input sequence of events, h may either halt (returning the left injection of the sum), or nondeterministically produce an output sequence of events and step to a new state $h' \in H$.

We compile terms to handlers as follows, defining a map $h : \text{Ctx} \times \text{Tm} \times \text{Ty} \times \mathbb{N} \rightarrow H$ which takes a 4-tuple of a context, term, type, and amount of gas (for recursive calls) to a handler. Intuitively, $h(\Gamma, e, s, n)$ takes an event sequence xs , deserializes it to obtain a prefix p , and runs e on p (using the incremental prefix semantics) to obtain e' and p' . It then returns nondeterministically one of the potential serializations of p' and steps to the handler defined by e' , with the proper derivative type. If the handler happens to run out of gas because of too many nested recursive calls, it simply halts and produces no output.

Definition 6.2 (Term to Handler Compilation). If

- (1) $xs \hookrightarrow^\Gamma p$
- (2) $p \Rightarrow e \downarrow^n e' \Rightarrow p'$
- (3) $Y = \{ys \mid p' \dagger^s ys\}$
- (4) $h' = h(\delta_p(\Gamma), e', \delta_{p'}(s), n)$,

then $h(\Gamma, e, s, n)(xs) = \text{inl}(Y, h')$; otherwise $h(\Gamma, e, s, n)(xs) = \text{inl}()$.

7 RELATED WORK

Streams as a programming abstraction can be traced back decades, to early work in the programming languages [16, 42, 58, 59] and database [1, 2, 5–7, 19, 48] communities. While streams are commonly viewed as homogeneous sequences, more interesting treatments of streams and stream-like data have been proposed. Streams in the database literature are sometimes viewed as time-varying relations, while the PL community has produced formalisms like process calculi and functional reactive programming. To our knowledge, ours is first type system for streams capturing both (1) heterogeneous patterns of events over time and (2) combinations of parallel and sequential data.

Sequential, homogeneous streams and dataflow programs. Traditionally, streams have been viewed as coinductive sequences [16]: a stream of A has a single (co)constructor, $\text{cocons} : \text{Stream } A \rightarrow (A \times \text{Stream } A)$ and acts as a lazily evaluated infinite list. This is the setting of traditional *dataflow programming* [58]. A major challenge in reasoning about dataflow over sequential streams is the nondeterminism arising from operators whose output may depend on the order in which events arrive on multiple input streams. Kahn's seminal process networks [42] (including their restriction to synchronous networks [12, 47, 59]) avoid this problem by allowing only *blocking reads* of messages passed between processes on FIFO queues. In contrast, the semantics of λ^{ST} leverages the type structure to guarantee deterministic parallel processing *without* blocking in many cases. For example, in the context of a Cut rule, if the type system can detect statically that a transformer is using two parallel streams safely, it can read from them simultaneously.

Partitioned streams. Building on streams as homogeneous sequences, modern stream processing systems such as Flink [17, 25], Spark Streaming [28, 63], Samza [27, 51], and Storm [29] support dynamic partitioning: a stream type defines one stream with many parallel substreams (where the number of substreams and assignment of data to substreams is determined at runtime). The type $\text{Stream } t$ in these systems is implicitly a parallel composition of homogeneous streams: $t^* \parallel \dots \parallel t^*$. Unlike in λ^{ST} , these parallel substreams cannot have more general types. Dynamic parallelism is difficult to support in the context of universal calculi for streaming; for example, Brooklet [56] and the DON Calculus [21] support data parallelism only as an optimization pass in limited cases. This is because stream partitioning does not in general preserve the semantics of the source program and can introduce undesirable nondeterminism [34, 49, 55]. While λ^{ST} does not support dynamic partitioning, we hope to address it in future work; see Section 8.

Streams as time-varying relations. In the database literature, streams are often viewed as relations (sets of tuples) that vary over time. Stream management systems in the early 2000s pioneered this paradigm, including Aurora [2] and Borealis [1], TelegraphCQ [19] and CACQ [48], and STREAM [5]. A time-varying relation can be viewed as either a function from timestamps to finite relations or an infinite set of timestamped values; this correspondence was elegantly exploited by early streaming query languages such as CQL [6, 7] and remains popular today [11, 39]. Time-varying relations can be expressed in λ^{ST} using Kleene star and concatenation: a relation of tuples of type T timestamped by Time can be expressed as $(T^* \cdot \text{Time})^*$. We can also express the common pattern where parallel streams are synchronized by a single timestamp (again, modulo dynamic partitioning) with types like $((T^* || T^*) \cdot \text{Time})^*$. Each Time event is a punctuation mark containing the timestamp of the prior set of tuples [41, 61]. Traditional systems include separate APIs for operations that modify punctuation (e.g., a *delay* function that increments timestamps); whereas in our system they are ordinary stream operators and punctuation markers are ordinary events.

Streams as pomsets. A sweet spot between the homogeneous sequential and relational viewpoints is identified by prior work treating streams as *pomsets* (partially ordered multisets) [3, 43–45, 49], inspired by prior work in concurrency theory [22, 50]. In a pomset, data items may be completely ordered (a sequence), completely unordered (a bag), or somewhere in between. While existing works have proposed pomset-based types for streams [3, 49], their types do not support concatenation and do not come with type *systems*—programs must be shown to be well typed semantically, rather than statically.

Functional reactive programming (FRP) [23] treats programs as incremental, reactive state machines written using functional combinators. The fundamental abstraction is a “signal”: a time-varying value $\text{Sig}(A) = \text{Time} \rightarrow A$. Work on type systems for FRP has used modal and substructural types [8, 9, 18, 46] to guarantee properties like causality, productivity, and space leak freedom. While our type system is not *designed* to address these issues, it does incidentally have bearing on them. For one, the existence of our incremental semantics demonstrates that our type system enforces causality: since outputs that have been incrementally emitted cannot be retracted or changed, the type system must ensure that past outputs cannot depend on future inputs. Similarly, potential space leaks can be detected statically by checking if only bounded-sized types are buffered using `wait` or the buffering built into the left rules for sums and star. Our current calculus does not guarantee productivity (new inputs eventually produce new outputs), but in Section 8 we discuss how to remedy this by imposing guardedness conditions on recursive calls.

Work by Jeffrey [40] permits the type of a signal to vary over time, using dependent types inspired by Linear Temporal Logic [53]. The system includes an *until* type that behaves like our concatenation type: a signal of type $A \cup B$ is a signal of type A , followed by a signal of type B . However, unlike parallel streams in our setting, time updates in steps, discretely; i.e., there’s a synchronous ordering between elements of one signal and elements on another. Concurrently with our work, Bahr and Møgelberg [10] proposes a modal type system to weaken the synchronicity assumption; however, it still treats signals as homogeneous.

Session types and process calculi. Another large body of work with a vision similar to ours is session types for process calculi [36], where types describe complex sequential protocols between communicating processes as they evolve through time. A main difference from our work is that the session type of a process describes the *protocol* for its communications with other processes—i.e., the sequence of sends and receives on different channels—while the stream type of a λ^{ST} program describes only the data that it communicates. Indeed, a stream transformer might display many patterns of communication with downstream transformers: it can run in “batch mode”—sending exactly one output after accepting all available input—or in a sequence smaller steps, sending along

partial outputs as it receives partial inputs. Also, a single channel in a process calculus cannot carry parallel substreams of events that are internally ordered but not ordered relative to each other. Recently, Frumin et al. [33] proposed a session-types interpretation of BI that the bunched structure very differently from λ^{ST} . In particular, processes of type $A * B$ and $A \wedge B$ both behave semantically like a process of type A in parallel with a process of type B ; in λ^{ST} , $s \cdot t$ and $s||t$ describe very different streams.

Concurrent Kleene Algebras and regular expression types. Finally, stream types are also inspired by Concurrent Kleene Algebras (CKAs) [35] and related syntaxes for pomset languages [45], though we are apparently the first to use these formalisms as *types* in a programming language, rather than as a tool for reasoning about concurrency. In particular, traditional applications of Kleene algebra such as NetKAT [4] and Concurrent NetKAT [62], use KA to model *programs*, whereas in λ^{ST} we use the KA structure to describe the data that programs exchange, while the programs themselves are written in a separate language. We have also been inspired by languages for programming with XML data [13, 31, 37, etc.] using types based on regular expressions.

8 CONCLUSIONS AND FUTURE WORK

We have proposed a new static type system for stream programming, strongly connected to a novel variant of BI logic and intended to capture both complex temporal patterns and deterministic parallel processing.

In Section 5, we presented a speculative higher-level syntax for λ^{ST} terms. Two significant challenges in building a real, statically typed, higher-level language on this foundation are (a) inferring type and context annotations on variables, cut, and left rules, and (b) inferring usages of structural rules. The first, we believe, is relatively easy: given a context and a variable, it is straightforward to find the one-hole context in which the variable resides, and thus the variable's type. A promising beginning for the second is a recent “labeled” formulation of the BI sequent calculus [38].

In Section 6, we showed how to run λ^{ST} terms as sequential automata, demonstrating the feasibility in principle of compiling λ^{ST} down to programs for an existing stream processing system like Apache Storm and thus inheriting its desirable fault-tolerance and delivery guarantees. We plan to build such a compiler and use it as a platform for experimenting with type-enabled optimizations and resource usage analysis.

Finally, we plan to extend the theory of stream types by (1) investigating a denotational semantics in pomsets, (2) adding support for *bags* (unbounded parallelism, the parallel analog of Kleene star) to enable dynamic partitioning, and (3) adding a *guardedness* condition on recursive calls to ensure termination and hence productivity.

REFERENCES

- [1] Daniel J Abadi, Yanif Ahmad, Magdalena Balazinska, Uğur Çetintemel, Mitch Cherniack, Jeong-Hyon Hwang, Wolfgang Lindner, Anurag Maskey, Alex Rasin, Esther Ryvkina, Nesime Tatbul, Ying Xing, and Stanley Zdonik. 2005. The Design of the Borealis Stream Processing Engine. In *Second Biennial Conference on Innovative Data Systems Research (CIDR)*.
- [2] Daniel J Abadi, Don Carney, Uğur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. 2003. Aurora: A New Model and Architecture for Data Stream Management. *The VLDB Journal* 12, 2 (2003). <https://doi.org/10.1007/s00778-003-0095-z>
- [3] Rajeev Alur, Phillip Hilliard, Zachary G Ives, Konstantinos Kallas, Konstantinos Mamouras, Filip Niksic, Caleb Stanford, Val Tannen, and Anton Xue. 2021. Synchronization Schemas. *Invited contribution, Principles of Database Systems*.
- [4] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. 2014. NetKAT: Semantic foundations for networks. *Acm sigplan notices* 49, 1 (2014), 113–126.
- [5] Arvind Arasu, Brian Babcock, Shivnath Babu, John Cieslewicz, Mayur Datar, Keith Ito, Rajeev Motwani, Utkarsh Srivastava, and Jennifer Widom. 2004. *STREAM: The Stanford Data Stream Management System*. Technical Report 2004-20. Stanford InfoLab.

- [6] Arvind Arasu, Shivnath Babu, and Jennifer Widom. 2003. CQL: A language for continuous queries over streams and relations. In *International Workshop on Database Programming Languages*. Springer.
- [7] Arvind Arasu, Shivnath Babu, and Jennifer Widom. 2006. The CQL Continuous Query Language: Semantic Foundations and Query Execution. *The VLDB Journal* 15, 2 (2006). <https://doi.org/10.1007/s00778-004-0147-z>
- [8] Patrick Bahr, Christian Uldal Graulund, and Rasmus Ejlers Møgelberg. 2019. Simply RaTT: A Fitch-Style Modal Calculus for Reactive Programming without Space Leaks. *Proc. ACM Program. Lang.* 3, ICFP, Article 109 (jul 2019), 27 pages. <https://doi.org/10.1145/3341713>
- [9] Patrick Bahr, Christian Uldal Graulund, and Rasmus Ejlers Møgelberg. 2021. Diamonds Are Not Forever: Liveness in Reactive Programming with Guarded Recursion. *Proc. ACM Program. Lang.* 5, POPL, Article 2 (jan 2021), 28 pages. <https://doi.org/10.1145/3434283>
- [10] Patrick Bahr and Rasmus Ejlers Møgelberg. 2023. Asynchronous Modal FRP. arXiv:2303.03170 [cs.PL]
- [11] Edmon Begoli, Tyler Akidau, Fabian Hueske, Julian Hyde, Kathryn Knight, and Kenneth Knowles. 2019. One SQL to Rule Them All—an Efficient and Syntactically Idiomatic Approach to Management of Streams and Tables. In *International Conference on Management of Data (SIGMOD)*.
- [12] Albert Benveniste, Paul Caspi, Stephen A Edwards, Nicolas Halbwachs, Paul Le Guernic, and Robert De Simone. 2003. The synchronous languages 12 years later. *Proc. IEEE* 91, 1 (2003).
- [13] Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. 2003. CDuce: An XML-Centric General-Purpose Language. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*. 51–63.
- [14] James Brotherston. 2005. Cyclic Proofs for First-Order Logic with Inductive Definitions. In *Automated Reasoning with Analytic Tableaux and Related Methods (Lecture Notes in Computer Science)*, Bernhard Beckert (Ed.). Springer, Berlin, Heidelberg, 78–92. https://doi.org/10.1007/11554554_8
- [15] Janusz A Brzozowski. 1964. Derivatives of regular expressions. *J. ACM* 11, 4 (1964).
- [16] William H Burge. 1975. Stream processing functions. *IBM Journal of Research and Development* 19, 1 (1975).
- [17] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink: Stream and Batch Processing in a Single Engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 36, 4 (2015).
- [18] Andrew Cave, Francisco Ferreira, Prakash Panangaden, and Brigitte Pientka. 2014. Fair Reactive Programming. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '14). Association for Computing Machinery, New York, NY, USA, 361–372. <https://doi.org/10.1145/2535838.2535881>
- [19] Sirish Chandrasekaran, Owen Cooper, Amol Deshpande, Michael J Franklin, Joseph M Hellerstein, Wei Hong, Sailesh Krishnamurthy, Samuel R Madden, Fred Reiss, and Mehul A Shah. 2003. TelegraphCQ: continuous dataflow processing. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 668–668.
- [20] Farzaneh Derakhshan. 2021. *Session-Typed Recursive Processes and Circular Proofs*. Ph.D. Dissertation. Caregie Mellon University. https://www.andrew.cmu.edu/user/fderakhs/publications/Dissertation_Farzaneh.pdf
- [21] Philip Dexter, Yu David Liu, and Kenneth Chiu. 2022. The essence of online data processing. *Proceedings of the ACM on Programming Languages* 6, OOPSLA2 (2022), 899–928.
- [22] Volker Diekert and Grzegorz Rozenberg. 1995. *The Book of Traces*. World Scientific. <https://doi.org/10.1142/2563>
- [23] Conal Elliott and Paul Hudak. 1997. Functional reactive animation. In *Second ACM SIGPLAN International Conference on Functional Programming (ICFP)*.
- [24] Jérôme Fortier and Luigi Santocanale. 2013. Cuts for circular proofs: semantics and cut-elimination. In *Computer Science Logic 2013 (CSL 2013) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 23)*, Simona Ronchi Della Rocca (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 248–262. <https://doi.org/10.4230/LIPIcs.CSL.2013.248> ISSN: 1868-8969.
- [25] Apache Software Foundation. 2019. Apache Flink. <https://flink.apache.org/>. (Accessed July 2022.).
- [26] Apache Software Foundation. 2019. Apache Heron (originally Twitter Heron). <https://heron.incubator.apache.org/>. (Accessed July 2022.).
- [27] Apache Software Foundation. 2019. Apache Samza. <https://samza.apache.org/>. (Accessed July 2022.).
- [28] Apache Software Foundation. 2019. Apache Spark Streaming. <https://spark.apache.org/streaming/>. (Accessed July 2022.).
- [29] Apache Software Foundation. 2019. Apache Storm. <https://storm.apache.org/>. (Accessed July 2022.).
- [30] Apache Software Foundation. 2021. Apache Beam. <https://beam.apache.org/>. (Accessed July 2022.).
- [31] Alain Frisch, Giuseppe Castagna, and Veronique Benzaken. 2002. Semantic Subtyping. In *Logic in Computer Science (LICS)*.
- [32] Dan Frumin. 2022. Semantic Cut Elimination for the Logic of Bunched Implications, Formalized in Coq. In *Proceedings of the 11th ACM SIGPLAN International Conference on Certified Programs and Proofs (Philadelphia, PA, USA) (CPP 2022)*. Association for Computing Machinery, New York, NY, USA, 291–306. <https://doi.org/10.1145/3497775.3503690>

- [33] Dan Frumin, Emanuele D’Osualdo, Bas van den Heuvel, and Jorge A. Pérez. 2022. A Bunch of Sessions: A Propositions-as-Sessions Interpretation of Bunched Implications in Channel-Based Concurrency. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 155 (oct 2022), 29 pages. <https://doi.org/10.1145/3563318>
- [34] Martin Hirzel, Robert Soulé, Scott Schneider, Buğra Gedik, and Robert Grimm. 2014. A catalog of stream processing optimizations. *ACM Computing Surveys (CSUR)* 46, 4 (2014).
- [35] CAR (Tony) Hoare, Bernhard Möller, Georg Struth, and Ian Wehrman. 2009. Concurrent Kleene Algebra. In *CONCUR 2009-Concurrency Theory: 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings* 20. Springer, 399–414.
- [36] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 273–284.
- [37] Haruo Hosoya, Jérôme Vouillon, and Benjamin C. Pierce. 2005. Regular Expression Types for XML. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 27, 1 (Jan. 2005), 46–90. Preliminary version in ICFP 2000.
- [38] Zhe Hou, Alwen Tiu, and Rajeev Gore. 2015. A Labelled Sequent Calculus for BBI: Proof Theory and Proof Search. arXiv:1302.4783 [cs.LO]
- [39] Namit Jain, Shailendra Mishra, Anand Srinivasan, Johannes Gehrke, Jennifer Widom, Hari Balakrishnan, Uğur Çetintemel, Mitch Cherniack, Richard Tibbetts, and Stan Zdonik. 2008. Towards a streaming SQL standard. *Proceedings of the VLDB Endowment* 1, 2 (2008).
- [40] Alan Jeffrey. 2012. LTL Types FRP: Linear-Time Temporal Logic Propositions as Types, Proofs as Functional Reactive Programs. In *Proceedings of the Sixth Workshop on Programming Languages Meets Program Verification (Philadelphia, Pennsylvania, USA) (PLPV ’12)*. Association for Computing Machinery, New York, NY, USA, 49–60. <https://doi.org/10.1145/2103776.2103783>
- [41] Theodore Johnson, Shanmugavelayutham Muthukrishnan, Vladislav Shkapenyuk, and Oliver Spatscheck. 2005. A heartbeat mechanism and its application in Gigascope. In *31st International Conference on Very Large Data Bases (VLDB)*. VLDB Endowment.
- [42] Gilles Kahn. 1974. The semantics of a simple language for parallel programming. *Information Processing* 74 (1974).
- [43] Konstantinos Kallas, Filip Niksic, Caleb Stanford, and Rajeev Alur. 2020. DiffStream: differential output testing for stream processing programs. *Proceedings of the ACM on Programming Languages* 4, OOPSLA (2020).
- [44] Konstantinos Kallas, Filip Niksic, Caleb Stanford, and Rajeev Alur. 2022. Stream Processing With Dependency-Guided Synchronization. In *Principles and Practice of Parallel Programming (PPoPP)*.
- [45] Tobias Kappé, Paul Brunet, Bas Luttik, Alexandra Silva, and Fabio Zanasi. 2019. On series-parallel pomset languages: Rationality, context-freeness and automata. *Journal of Logical and Algebraic Methods in Programming* 103 (2019), 130–153. <https://doi.org/10.1016/j.jlamp.2018.12.001>
- [46] Neelakantan R. Krishnaswami. 2013. Higher-Order Functional Reactive Programming without Spacetime Leaks. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (Boston, Massachusetts, USA) (ICFP ’13)*. Association for Computing Machinery, New York, NY, USA, 221–232. <https://doi.org/10.1145/2500365.2500588>
- [47] Edward A Lee and David G Messerschmitt. 1987. Synchronous data flow. *Proc. IEEE* 75, 9 (1987).
- [48] Samuel Madden, Mehul Shah, Joseph M Hellerstein, and Vijayshankar Raman. 2002. Continuously Adaptive Continuous Queries over Streams. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 12 pages. <https://doi.org/10.1145/564691.564698>
- [49] Konstantinos Mamouras, Caleb Stanford, Rajeev Alur, Zachary G Ives, and Val Tannen. 2019. Data-trace types for distributed stream processing systems. In *40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.
- [50] Antoni Mazurkiewicz. 1986. Trace theory. In *Advanced course on Petri nets*. Springer.
- [51] Shadi A Noghabi, Kartik Paramasivam, Yi Pan, Navina Ramesh, Jon Bringham, Indranil Gupta, and Roy H Campbell. 2017. Samza: Stateful Scalable Stream Processing at LinkedIn. *Proceedings of the VLDB Endowment* 10, 12 (2017).
- [52] Peter W O’Hearn and David J Pym. 1999. The logic of bunched implications. *Bulletin of Symbolic Logic* 5, 2 (1999), 215–244.
- [53] Amir Pnueli. 1977. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. iee, 46–57.
- [54] John C Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*. IEEE, 55–74.
- [55] Scott Schneider, Martin Hirzel, Buğra Gedik, and Kun-Lung Wu. 2013. Safe data parallelism for general streaming. *IEEE Trans. Comput.* 64, 2 (2013).
- [56] Robert Soulé, Martin Hirzel, Robert Grimm, Buğra Gedik, Henrique Andrade, Vibhore Kumar, and Kun-Lung Wu. 2010. A universal calculus for stream processing languages. In *European Symposium on Programming (ESOP)*. Springer.
- [57] Caleb Stanford. 2022. *Safe Programming over Distributed Streams*. Ph.D. Dissertation. University of Pennsylvania.

- [58] Robert Stephens. 1997. A survey of stream processing. *Acta Informatica* 34, 7 (1997).
- [59] William Thies, Michal Karczmarek, and Saman Amarasinghe. 2002. StreamIt: A language for streaming applications. In *International Conference on Compiler Construction*. Springer.
- [60] Peter A. Tucker, David Maier, Tim Sheard, and Leonidas Fegaras. 2003. Exploiting Punctuation Semantics in Continuous Data Streams. *IEEE Trans. on Knowl. and Data Eng.* 15, 3 (mar 2003), 555–568. <https://doi.org/10.1109/TKDE.2003.1198390>
- [61] Peter A Tucker, David Maier, Tim Sheard, and Leonidas Fegaras. 2003. Exploiting punctuation semantics in continuous data streams. *IEEE Transactions on Knowledge and Data Engineering* 15, 3 (2003).
- [62] Jana Wagemaker, Nate Foster, Tobias Kappé, Dexter Kozen, Jurriaan Rot, and Alexandra Silva. 2022. Concurrent NetKAT: Modeling and analyzing stateful, concurrent networks. In *European Symposium on Programming*. Springer International Publishing Cham, 575–602.
- [63] Matei Zaharia, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. 2013. Discretized Streams: Fault-tolerant Streaming Computation at Scale. In *24th Symposium on Operating Systems Principles (SOSP)*. ACM. <https://doi.org/10.1145/2517349.2522737>

A MECHANIZATION AND REFERENCE INTERPRETER

The supplemental material contains a Coq formalization of our main results and a reference interpreter written in Haskell.

The Coq formalization includes the following:

Table 1. Overview of Coq Development

Name	Description	Location
Basics	Definitions of types, contexts, batches, and prefixes	verif/basics.v
CtxSub	Context subtyping and its semantics	verif/ctxsubty.v
PProj	Prefix and batch projection definitions and theorems	verif/pproj.v
PrefixConcat	Prefix and batch concatenation definitions and theorems	verif/prefixconcat.v
Events	Events, serialization and deserialization	verif/events.v
Language	Syntax of terms, the typing judgment, batch and incremental semantics, correctness and the homomorphism property	verif/language.v

The Haskell interpreter assumes well-typed terms, failing if anything is ill-typed; we have not translated the type system definition to Haskell. We do not have proofs that the interpreter (a function) and the Coq definition of the semantics (an inductive relation) agree, but we have tested a majority of the important theorems from Coq via QuickCheck. We also use QuickCheck to check the homomorphism theorem for fixed terms. The important directories include:

Table 2. Overview of Haskell Interpreter

Name	Description	Location
Types	Types and contexts	src/Types.hs
Interpreter	Code for the interpreter	src/Interpreter.hs
Examples	Core terms for examples	src/Examples.hs

The Coq development has been tested with Coq version 8.17.0 using CoqHammer Tactics version 1.3.2. The Haskell development has been tested with GHC version 9.2.7 and Stack version 2.9.3.

B TECHNICALITIES

This appendix collects technical definitions that did not fit in the main body of the paper.

B.1 Basics

Stream types are defined by the following grammar. The base types included are the unit type 1 which types streams that contain exactly one unit element, the type of the empty stream ε , and the type of streams consisting of a single integer, Int . Larger types include the concatenation type $s \cdot t$, the sum type $s + t$, the parallel stream type $s||t$, and the star type s^* .

$$s, t, r ::= 1 \mid \varepsilon \mid \text{Int} \mid s \cdot t \mid s + t \mid s||t \mid s^*$$

Contexts in the stream types calculus system have a bunched structure. The context former Γ , Δ corresponds to the parallel type, while the context former Γ ; Δ corresponds to the concatenation type. The two context formers share a unit, written as “.”.

$$\Gamma ::= \cdot \mid \Gamma, \Gamma \mid \Gamma; \Gamma \mid x : s$$

A stream type is *nullable* if it includes no data. Nullable types are parallel combinations of ε .

Definition B.1 (Nullable). We define a judgment s nullable as follows:

$$\frac{}{\varepsilon \text{ nullable}} \qquad \frac{s \text{ nullable} \quad t \text{ nullable}}{s||t \text{ nullable}}$$

We extend to contexts in the natural way:

$$\frac{}{\cdot \text{ nullable}} \qquad \frac{\Gamma \text{ nullable} \quad \Gamma' \text{ nullable}}{\Gamma, \Gamma' \text{ nullable}}$$

Batches are defined as in the paper, with a typing relation $b : \text{batch}(s)$ and a context typing relation $g : \text{batch}(\Gamma)$.

Definition B.2 (Batch). The grammar of batches is given by:

$$b ::= \text{oneb} \mid \text{batch}_{\text{Int}}(n) \mid \text{epsb} \mid \text{catb}(b, b') \mid \text{parb}(b, b') \mid \text{sumbA}(b) \mid \text{sumbB}(b) \mid [] \mid b :: b'$$

We define the batch has type relation as:

$$\frac{}{\text{oneb} : \text{batch}(1)} \quad \frac{}{\text{epsb} : \text{batch}(\varepsilon)} \quad \frac{b : \text{batch}(s) \quad b' : \text{batch}(t)}{\text{catb}(b, b') : \text{batch}(s \cdot t)} \\ \frac{b : \text{batch}(s) \quad b' : \text{batch}(t)}{\text{parb}(b, b') : \text{batch}(s||t)} \quad \frac{b : \text{batch}(s)}{\text{sumbA}(b) : \text{batch}(s + t)} \quad \frac{b : \text{batch}(t)}{\text{sumbB}(b) : \text{batch}(s + t)} \\ \frac{}{[] : \text{batch}(s^*)} \quad \frac{b : \text{batch}(s) \quad b' : \text{batch}(s^*)}{b :: b' : \text{batch}(s^*)}$$

And we lift this to contexts by:

$$\frac{}{\text{epsb} : \text{batch}(\cdot)} \quad \frac{b : \text{batch}(\Gamma) \quad b' : \text{batch}(\Delta)}{\text{catb}(b, b') : \text{batch}(\Gamma; \Delta)} \quad \frac{b : \text{batch}(\Gamma) \quad b' : \text{batch}(\Delta)}{\text{parb}(b, b') : \text{batch}(\Gamma, \Delta)}$$

Prefixes are also like in the main paper, a typing relation $p : \text{prefix}(s)$ and a context typing relation $p : \text{prefix}(\Gamma)$.

Definition B.3 (Prefix). The grammar of prefixes is given by:

$$\begin{aligned}
 p ::= & \text{onepA} | \text{onepB} | \text{epsp} | \text{parp}(p, p') \\
 & | \text{catpA}(p) | \text{catpB}(b, p) \\
 & | \text{sumpEmp} | \text{sumpA}(p) | \text{sumpB}(p) \\
 & | \text{stpEmp} | \text{stpDone} \\
 & | \text{stpA}(p) | \text{stpB}(b, p)
 \end{aligned}$$

$$\begin{array}{c}
 \frac{}{\text{epsp} : \text{prefix}(\varepsilon)} \quad \frac{}{\text{onepA} : \text{prefix}(1)} \quad \frac{}{\text{onepB} : \text{prefix}(1)} \\
 \frac{p : \text{prefix}(\Gamma) \quad p' : \text{prefix}(\Gamma')}{\text{parp}(p, p') : \text{prefix}(\Gamma, \Gamma')} \quad \frac{p : \text{prefix}(s)}{\text{catpA}(p) : \text{prefix}(s \cdot t)} \\
 \frac{b : \text{batch}(s) \quad p : \text{prefix}(t)}{\text{catpB}(b, p) : \text{prefix}(s \cdot t)} \quad \frac{}{\text{sumpEmp} : \text{prefix}(s + t)} \quad \frac{}{\text{sumpA}(p) : \text{prefix}(s + t)} \\
 \frac{p : \text{prefix}(t)}{\text{sumpB}(p) : \text{prefix}(s + t)} \quad \frac{}{\text{stpEmp} : \text{prefix}(s^*)} \quad \frac{}{\text{stpDone} : \text{prefix}(s^*)} \\
 \frac{p : \text{prefix}(s)}{\text{stpA}(p) : \text{prefix}(s^*)} \quad \frac{b : \text{batch}(s) \quad p : \text{prefix}(s^*)}{\text{stpB}(b, p) : \text{prefix}(s^*)}
 \end{array}$$

We lift this to contexts by:

$$\frac{p : \text{prefix}(\Gamma) \quad p' : \text{prefix}(\Gamma')}{\text{parp}(p, p') : \text{prefix}(\Gamma, \Gamma')} \quad \frac{p : \text{prefix}(\Gamma)}{\text{catpA}(p) : \text{prefix}(\Gamma; \Gamma')} \\
 \frac{b : \text{batch}(\Gamma) \quad p : \text{prefix}(\Gamma')}{\text{catpB}(b, p) : \text{prefix}(\Gamma; \Gamma')}$$

For each type s , we define the “empty” prefix emp_s inductively on the structure of s . In principle, one could define a single empty prefix shared by all types, but this would over-represent in many cases: the pointed empty prefix emp would not be the same as $\text{parp}(\text{emp}, \text{emp})$, for example. While type-indexing the definition requires us to carry around more types at runtime (to compute the empty prefix of the right type), we view this as a preferable choice.

Definition B.4 (Empty Prefix). The empty prefix is defined as follows:

$$\begin{aligned}
 (\varepsilon) \quad \text{emp}_\varepsilon &= \text{epsp} \\
 (1) \quad \text{emp}_1 &= \text{onepA} \\
 (s || t) \quad \text{emp}_{s || t} &= \text{parp}(\text{emp}_s, \text{emp}_t) \\
 (s + t) \quad \text{emp}_{s+t} &= \text{sumpEmp} \\
 (s \cdot t) \quad \text{emp}_{s \cdot t} &= \text{catpA}(\text{emp}_s) \\
 (s^*) \quad \text{emp}_{s^*} &= \text{stpEmp}
 \end{aligned}$$

We lift this to contexts in the natural way, with $\text{emp}_\cdot = \text{epsp}$, and $\text{emp}_{\Gamma; \Delta} = \text{catpA}(\text{emp}_\Gamma)$, and $\text{emp}_{\Gamma, \Delta} = \text{parp}(\text{emp}_\Gamma, \text{emp}_\Delta)$.

Given a batch, we can forget about its batch structure and demote it to a prefix.

Definition B.5 (Batch to Prefix Lift).

$$\begin{aligned}
(\text{oneb})^\circ &= \text{onepB} \\
(\text{epsb})^\circ &= \text{epsp} \\
(\text{catb}(b, b'))^\circ &= \text{catpB}(b, (b')^\circ) \\
(\text{parb}(b, b'))^\circ &= \text{parp}((b)^\circ, (b')^\circ) \\
(\text{sumbA}(b))^\circ &= \text{sumpA}((b)^\circ) \\
(\text{sumbB}(b))^\circ &= \text{sumpB}((b)^\circ) \\
([\])^\circ &= \text{stpDone} \\
(b :: b')^\circ &= \text{stpB}(b, (b')^\circ)
\end{aligned}$$

This function takes well-typed batches to well-typed prefixes.

THEOREM B.6 (BATCH TO PREFIX LIFT WELL-TYPED). *If $b : \text{batch}(s)$, then $(b)^\circ : \text{prefix}(s)$*

PROOF. `basics.v:prefixOf_correct` □

If a batch converted to a prefix is well-typed, then the batch was also well-typed.

THEOREM B.7 (BATCH TO PREFIX LIFT WELL-TYPED CONVERSE). *If $(b)^\circ : \text{prefix}(s)$, then $b : \text{batch}(s)$.*

PROOF. `basics.v:prefixOf_inv` □

THEOREM B.8 (BATCH TO PREFIX LIFT WELL-TYPED (CONTEXT)). *If $b : \text{batch}(\Gamma)$, then $(b)^\circ : \text{prefix}(\Gamma)$*

PROOF. `basics.v:prefixOf_correct_ctx` □

B.2 Derivatives

Definition B.9 (Derivatives). We define a 3-place relation $\delta_p(s) \sim s'$ between a prefix and two types.

$$\begin{array}{c}
\frac{}{\delta_{\text{epsp}}(\varepsilon) \sim \varepsilon} \quad \frac{}{\delta_{\text{onepA}}(1) \sim 1} \quad \frac{}{\delta_{\text{onepB}}(1) \sim \varepsilon} \quad \frac{\delta_p(s) \sim s'}{\delta_{\text{catpA}(p)}(s \cdot t) \sim s' \cdot t} \\
\frac{\delta_p(t) \sim t'}{\delta_{\text{catpB}(b,p)}(s \cdot t) \sim t'} \quad \frac{\delta_p(s) \sim s' \quad \delta_{p'}(t) \sim t'}{\delta_{\text{parp}(p,p')}(s || t) \sim s' || t'} \quad \frac{}{\delta_{\text{sumpEmp}}(s + t) \sim s + t} \\
\frac{}{\delta_{\text{sumpA}(p)}(s + t) \sim s'} \quad \frac{\delta_p(s) \sim s'}{\delta_{\text{sumpB}(p)}(s + t) \sim t'} \quad \frac{}{\delta_{\text{stpEmp}}(s^*) \sim s^*} \quad \frac{}{\delta_{\text{stpDone}}(s^*) \sim \varepsilon} \\
\frac{}{\delta_{\text{stpA}(p)}(s^*) \sim s' \cdot s^*} \quad \frac{}{\delta_{\text{stpB}(_,p)}(s^*) \sim s'}
\end{array}$$

We lift this to contexts in the natural way:

$$\frac{}{\delta_{\text{catpA}(p)}(\Gamma; \Delta) \sim \Gamma'; \Delta} \quad \frac{}{\delta_{\text{catpB}(b,p)}(\Gamma; \Delta) \sim \Delta'} \quad \frac{\delta_p(\Gamma) \sim \Gamma' \quad \delta_{p'}(\Delta) \sim \Delta'}{\delta_{\text{parp}(p,p')}(\Gamma, \Delta) \sim \Gamma', \Delta'}$$

Derivatives are functions defined when the prefix input is well-typed.

THEOREM B.10 (DERIVATIVE FUNCTION). *For any p and s , there is at most one s' such that $\delta_p(s) \sim s'$. If $p : \text{prefix}(s)$, then such an s' exists.*

PROOF. Existence by `basics.v:derivrel_fun`, uniqueness by `basics.v:derivrel_det`. \square

When it's guaranteed to exist, we write this s' simply as $\delta_p(s)$. The empty prefix is the identity for the derivative operator.

THEOREM B.11 (EMPTY PREFIX DERIVATIVE). $\delta_{\text{emp}_s}(s) = s$.

PROOF. `basics.v:derivrel_emp` \square

THEOREM B.12 (DERIVATIVE FUNCTION (CONTEXTS)). *For any p and Γ , there is at most one Γ' such that $\delta_p(\Gamma) \sim \Gamma'$. If $p : \text{prefix}(\Gamma)$, then such an Γ' exists.*

PROOF. Existence by `basics.v:derivrelCtx_fun`, uniqueness by `basics.v:derivrelCtx_det`. \square

THEOREM B.13 (EMPTY PREFIX DERIVATIVE (CONTEXTS)). $\delta_{\text{emp}_s}(\Gamma) \sim \Gamma$.

PROOF. `basics.v:derivrelCtx_emp` \square

If we take a batch, convert it to a prefix, and then take a derivative by that prefix, the resulting type is nullable. Intuitively, this is because “the rest” of a stream after a batch should be empty.

THEOREM B.14 (NULLABLE BATCH DERIVATIVE). *If $b : \text{batch}(s)$ then $\delta_{(b)^\circ}(s)$ nullable*

PROOF. `basics.v:batch_lift_derivrel` \square

THEOREM B.15 (NULLABLE BATCH DERIVATIVE). *If $b : \text{batch}(\Gamma)$ then $\delta_{(b)^\circ}(\Gamma)$ nullable*

PROOF. `basics.v:batch_lift_derivrelCtx` \square

B.3 Done Batches

For a nullable type s , there is a unique “empty batch”, which we refer to as the done batch.

Definition B.16 (Done Batch).

$$\frac{}{\varepsilon \text{ done epsb}} \qquad \frac{s \text{ done } b \quad t \text{ done } b'}{s \parallel t \text{ done parb}(b, b')}$$

THEOREM B.17 (DONE BATCH CORRECTNESS). *For any s , there is at most one b such that $s \text{ done } b$. Moreover, such a b exists exactly when s nullable.*

PROOF. Uniqueness by `basics.v:doneBatch_det`, existence by `basics.v:doneBatchTheorem` \square

B.4 Concatenation

Prefix-Batch Concatenation. Given a prefix p of s , and a batch b of $\delta_p(s)$, we often want to consider the batch of s given by p , followed by b . The following functional relation gives a way of computing this. We also lift this definition to contexts.

Definition B.18 (Prefix-Batch Concatenation). We define a relation $p \cdot_s b \sim b'$.

$$\begin{array}{c}
\frac{}{\text{epsp} \cdot_{\varepsilon} \text{epsb} \sim \text{epsb}} \quad \frac{}{\text{onepA} \cdot_1 \text{oneb} \sim \text{oneb}} \quad \frac{}{\text{onepB} \cdot_1 \text{epsb} \sim \text{oneb}} \\
\frac{p_1 \cdot_s p'_1 \sim b'_1 \quad p_2 \cdot_t p'_2 \sim b'_2}{\text{parp}(p_1, p_2) \cdot_{s||t} \text{parb}(b_1, b_2) \sim \text{parb}(b'_1, b'_2)} \quad \frac{}{\text{sumpEmp} \cdot_{s+t} b \sim b} \\
\frac{p \cdot_s b \sim b'}{\text{sumpA}(p) \cdot_{s+t} b \sim \text{sumbA}(b')} \quad \frac{p \cdot_t b \sim b'}{\text{sumpB}(p) \cdot_{s+t} b \sim \text{sumbB}(b')} \\
\frac{}{\text{catpA}(p) \cdot_{s,t} \text{catb}(b_1, b_2) \sim \text{catb}(b'_1, b'_2)} \quad \frac{p \cdot_t b' \sim b''}{\text{catpB}(b, p) \cdot_{s,t} b' \sim \text{catb}(b, b'')} \\
\frac{}{\text{stpEmp} \cdot_{s^*} b \sim b} \quad \frac{}{\text{stpDone} \cdot_{s^*} \text{epsb} \sim []} \quad \frac{}{\text{stpA}(p) \cdot_{s^*} \text{catb}(b_1, b_2) \sim b'_1 :: b_2} \\
\frac{p \cdot_s b' \sim b''}{\text{stpB}(b, p) \cdot_{s^*} b' \sim b :: b''}
\end{array}$$

We lift this to contexts in the natural way

$$\frac{p_1 \cdot_{\Gamma} p'_1 \sim b'_1 \quad p_2 \cdot_{\Delta} p'_2 \sim b'_2}{\text{parp}(p_1, p_2) \cdot_{\Gamma, \Delta} \text{parb}(b_1, b_2) \sim \text{parb}(b'_1, b'_2)} \quad \frac{p \cdot_{\Gamma} b_1 \sim b'_1}{\text{catpA}(p) \cdot_{\Gamma, \Delta} \text{catb}(b_1, b_2) \sim \text{catb}(b'_1, b_2)} \\
\frac{p \cdot_{\Delta} b' \sim b''}{\text{catpB}(b, p) \cdot_{\Gamma, \Delta} b' \sim \text{catb}(b, b'')}$$

Theorem B.19 establishes the functional behavior of this relation. When the preconditions are satisfied (p : prefix (s) and b : batch ($\delta_p(s)$)), we write $p \cdot_s b$ for the unique b' that it guarantees.

THEOREM B.19 (BATCH PREFIX CONCATENATION FUNCTION). *For all s, p, b , there is at most one b' such that $p \cdot_s b \sim b'$. If p : prefix (s) and b : batch ($\delta_p(s)$), then such a b' exists, and it satisfies $b' : \text{batch}(s)$.*

PROOF. Existence by `prefixconcat.v:prefixBatchConcatRel_fun`, uniqueness by `prefixconcat.v:prefixBatchConcatRel_det`, type-correctness by `prefixconcat.v:prefixBatchConcatRel_correct`. \square

If we take a batch b and demote it to a prefix with $(b)^\circ$, then concatenating any batch b' to it just gives b back. Intuitively, this is because if b : batch (s), then the type $\delta_{(b)^\circ}(s)$ is nullable, and so any batch b' of this type contains no data.

THEOREM B.20 (BATCH PREFIX CONCATENATION OF PREFIX LIFT). *If b : batch (s) and $\delta_{(b)^\circ}(s) \sim s'$ and $b' : \text{batch}(s')$, then $(b)^\circ \cdot_s b' = b$.*

PROOF. `prefixconcat.v:prefixBatchConcatRel_prefixOf` \square

Conversely, concatenating a batch b onto the empty prefix `emps` gives the whole batch b . This theorem relies on Theorem B.13, which ensures that $\delta_{\text{emp}_s}(s) = s$.

THEOREM B.21 (BATCH PREFIX CONCATENATION OF EMPTY). *If b : batch (s), then `emps` ·_s $b = b$.*

PROOF. `prefixconcat.v:prefixBatchConcatRel_emp` \square

All of the three previous theorems all have identical versions for batches and prefixes of contexts.

THEOREM B.22 (BATCH PREFIX CONCATENATION FUNCTION (CONTEXTS)). *For all Γ, p, b , there is at most one b' such that $p \cdot_{\Gamma} b \sim b'$. If $p : \text{prefix}(\Gamma)$ and $b : \text{batch}(\delta_p(\Gamma))$, then such a b' exists, and it satisfies $b' : \text{batch}(\Gamma)$.*

PROOF. Existence by `prefixconcat.v:prefixBatchConcatCtxRel_fun`,
 uniqueness by `prefixconcat.v:prefixBatchConcatCtxRel_det`,
 type-correctness by `prefixconcat.v:prefixBatchConcatCtxRel_correct`. \square

THEOREM B.23 (BATCH PREFIX CONCATENATION OF PREFIX LIFT (CONTEXTS)). *If $b : \text{batch}(\Gamma)$ and $b' : \text{batch}(\delta_{(b)}^\circ(\Gamma))$, then $(b)^\circ \cdot_{\Gamma} b' = b$.*

PROOF. `prefixconcat.v:prefixBatchConcatCtxRel_prefixOf` \square

THEOREM B.24 (BATCH PREFIX CONCATENATION OF EMPTY (CONTEXTS)). *If $b : \text{batch}(s)$, then $\text{emp}_s \cdot_s b = b$.*

PROOF. `prefixconcat.v:prefixBatchConcatCtxRel_emp` \square

Prefix-Prefix Concatenation. More generally, we often want to concatenate a prefix p of s with a prefix p' of $\delta_p(s)$. This is defined with another 4-place, type-indexed relation.

Definition B.25 (Prefix Concatenation). We define a relation $p \cdot_s p' \sim p''$.

$$\begin{array}{c}
 \frac{}{\text{epsp} \cdot_{\varepsilon} \text{epsp} \sim \text{epsp}} \\
 \frac{p : \text{prefix}(1)}{\text{onepA} \cdot_1 p \sim p} \qquad \frac{}{\text{onepB} \cdot_1 \text{epsp} \sim \text{onepB}} \\
 \frac{p_1 \cdot_s p'_1 \sim p''_1 \quad p_2 \cdot_t p'_2 \sim p''_2}{\text{parp}(p_1, p_2) \cdot_{s||t} \text{parp}(p'_1, p'_2) \sim \text{parp}(p''_1, p''_2)} \\
 \frac{p \cdot_s p' \sim p''}{\text{catpA}(p) \cdot_{s \cdot t} \text{catpA}(p') \sim \text{catpA}(p'')} \qquad \frac{p \cdot_s b \sim b'}{\text{catpA}(p) \cdot_{s \cdot t} \text{catpB}(b, p') \sim \text{catpB}(b', p')} \\
 \frac{p \cdot_t p' \sim p''}{\text{catpB}(b, p) \cdot_{s \cdot t} p' \sim \text{catpB}(b, p'')} \\
 \frac{p \cdot_s p' \sim p''}{\text{sumpEmp} \cdot_{s+t} p \sim p} \qquad \frac{p \cdot_s p' \sim p''}{\text{sumpA}(p) \cdot_{s+t} p' \sim \text{sumpA}(p'')} \qquad \frac{p \cdot_t p' \sim p''}{\text{sumpB}(p) \cdot_{s+t} p' \sim \text{sumpB}(p'')} \\
 \frac{p \cdot_s p' \sim p''}{\text{stpEmp} \cdot_{s^*} p \sim p} \qquad \frac{p \cdot_s p' \sim p''}{\text{stpDone} \cdot_{s^*} \text{epsp} \sim \text{stpDone}} \qquad \frac{p \cdot_s p' \sim p''}{\text{stpA}(p) \cdot_{s^*} \text{catpA}(p') \sim \text{stpA}(p'')} \\
 \frac{p \cdot_s b \sim b'}{\text{stpA}(p) \cdot_{s^*} \text{catpB}(b, p') \sim \text{stpB}(b', p')} \qquad \frac{p \cdot_s p' \sim p''}{\text{stpB}(b, p) \cdot_{s^*} p' \sim \text{stpB}(b, p'')}
 \end{array}$$

We lift this to contexts.

$$\begin{array}{c}
 \frac{p_1 \cdot_{\Gamma} p'_1 \sim p''_1 \quad p_2 \cdot_{\Delta} p'_2 \sim p''_2}{\text{parp}(p_1, p_2) \cdot_{\Gamma, \Delta} \text{parp}(p'_1, p'_2) \sim \text{parp}(p''_1, p''_2)} \\
 \frac{p \cdot_{\Gamma} p' \sim p''}{\text{catpA}(p) \cdot_{\Gamma, \Delta} \text{catpA}(p') \sim \text{catpA}(p'')} \qquad \frac{p \cdot_{\Gamma} b \sim b'}{\text{catpA}(p) \cdot_{\Gamma, \Delta} \text{catpB}(b, p') \sim \text{catpB}(b', p')} \\
 \frac{p \cdot_{\Delta} p' \sim p''}{\text{catpB}(b, p) \cdot_{\Gamma, \Delta} p' \sim \text{catpB}(b, p'')}
 \end{array}$$

This relation is a function when the inputs are well-typed. Because of this, when $p : \text{prefix}(s)$ and $p' : \text{prefix}(\delta_p(s))$, we write $p \cdot_s p'$ for the unique p'' that the following theorem guarantees.

THEOREM B.26 (PREFIX CONCATENATION FUNCTION). *For all p, p' and s , there is at most one p'' such that $p \cdot_s p' \sim p''$. If $p : \text{prefix}(s)$ and $p' : \text{prefix}(\delta_p(s))$, then such a p'' exists, and satisfies:*

- (1) $p'' : \text{prefix}(s)$
- (2) $\delta_{p''}(s) = \delta_{p'}(\delta_p(s))$

PROOF. Existence by `prefixconcat.v:prefixConcatRel_fun`, uniqueness by `prefixconcat.v:prefixConcatRel_det`, well-typedness and derivative condition by `prefixconcat.v:prefixConcatRel_correct`. \square

Concatenating a prefix p to the empty prefix yields p back, and...

THEOREM B.27 (PREFIX CONCATENATION EMPTY 1). *If $p : \text{prefix}(s)$, then $\text{emp}_s \cdot_s p = p$*

PROOF. `prefixconcat.v:prefixConcatRel_emp` \square

Concatenating the empty prefix to a prefix p also gives p .

THEOREM B.28 (PREFIX CONCATENATION EMPTY 2). *If $p : \text{prefix}(s)$, then $p \cdot_s \text{emp}_{\delta_p(s)} = p$*

PROOF. `prefixconcat.v:prefixConcatRel_emp'` \square

THEOREM B.29 (PREFIX CONCATENATION BATCH LIFT 1). *If $b : \text{batch}(s)$ and $p : \text{prefix}(\delta_{(b)^\circ}(s))$, then $(b)^\circ \cdot_s p = (b)^\circ$*

PROOF. `prefixconcat.v:prefixConcatRel_prefixOf` \square

THEOREM B.30 (PREFIX CONCATENATION BATCH LIFT 2). *If $p \cdot_s b \sim b'$, then $p \cdot_s (b)^\circ \sim (b')^\circ$.*

PROOF. `prefixconcat.v:prefixConcatRel_prefixOf'` \square

Batch-to-prefix concatenation is associative.

THEOREM B.31 (BATCH PREFIX CONCATENATION ASSOCIATIVITY). $p \cdot_s (p' \cdot_{\delta_p(s)} b) = (p \cdot_s p') \cdot_s b$

PROOF. `prefixconcat.v:prefixBatchConcatRel_assoc` \square

Prefix concatenation is also associative.

THEOREM B.32 (PREFIX CONCATENATION ASSOCIATIVITY). $p \cdot_s (p' \cdot_{\delta_p(s)} p'') = (p \cdot_s p') \cdot_s p''$

PROOF. `prefixconcat.v:prefixConcatRel_assoc` \square

All of the previous theorems about prefix concatenation for types are also needed for contexts.

THEOREM B.33 (PREFIX CONCATENATION FUNCTION (CONTEXTS)). *For all p, p' and Γ , there is at most one p'' such that $p \cdot_\Gamma p' \sim p''$. If $p : \text{prefix}(\Gamma)$ and $p' : \text{prefix}(\delta_p(\Gamma))$, then such a p'' exists, and satisfies:*

- (1) $p'' : \text{prefix}(\Gamma)$
- (2) $\delta_{p''}(\Gamma) = \delta_{p'}(\delta_p(\Gamma))$

PROOF. Existence by `prefixconcat.v:prefixConcatCtxRel_fun`, uniqueness by `prefixconcat.v:prefixConcatCtxRel_det`, well-typedness and derivative condition by `prefixconcat.v:prefixConcatCtxRel_correct`. \square

THEOREM B.34 (BATCH PREFIX CONCATENATION ASSOCIATIVITY (CONTEXTS)). $p \cdot_{\Gamma} (p' \cdot_{\delta_p(\Gamma)} b) = (p \cdot_{\Gamma} p') \cdot_{\Gamma} b$

PROOF. `prefixconcat.v:prefixBatchConcatCtxRel_assoc` □

THEOREM B.35 (PREFIX CONCATENATION ASSOCIATIVITY (CONTEXTS)). $p \cdot_{\Gamma} (p' \cdot_{\delta_p(\Gamma)} p'') = (p \cdot_{\Gamma} p') \cdot_{\Gamma} p''$

PROOF. `prefixconcat.v:prefixConcatCtxRel_assoc` □

THEOREM B.36 (PREFIX CONCATENATION BATCH LIFT 1 (CONTEXTS)). *If $b : \text{batch}(\Gamma)$ and $p : \text{prefix}(\delta_{(b)^\circ}(\Gamma))$, then $(b)^\circ \cdot_{\Gamma} \Gamma = (b)^\circ$*

PROOF. `prefixconcat.v:prefixConcatCtxRel_prefixOf` □

THEOREM B.37 (PREFIX CONCATENATION BATCH LIFT 2 (CONTEXTS)). *If $p \cdot_{\Gamma} b \sim b'$, then $p \cdot_{\Gamma} (b)^\circ \sim (b')^\circ$.*

PROOF. `prefixconcat.v:prefixConcaCtxtRel_prefixOf` □

B.5 Projection Relations

Batch Projection. Given a batch or prefix of a context $\Gamma(\Delta)$, we often want to project out the part of the data corresponding to the subtree Δ . In the case of batches, this always exists: a batch of $\Gamma(\Delta)$ contains a full batch of Δ , located by following the hole in $\Gamma(-)$ down to Δ . We also often want to substitute a new batch in where the old one was — this corresponds to a sort of “re-binding” of the subcontext Δ' to a new one. The batch projection relation $b \rightsquigarrow_{\Gamma(-)} b', f$ computes both of these things: b' is the batch found at the $\Gamma(-)$ -position in b , and f is a function from batches to batches such that $f(b'')$ is the batch b , but with b'' where b' used to be.

Definition B.38 (Batch Projection). We inductively define the relation $b \rightsquigarrow_{\Gamma(-)} b'$ as follows

$$\frac{\frac{\frac{b \rightsquigarrow_{\Gamma_2(-)} b, f \quad g(b') = \text{parb}(b_1, f(b'))}{\text{parb}(b_1, b_2) \rightsquigarrow_{\Gamma_1, \Gamma_2(-)} b, g} \quad \frac{b_1 \rightsquigarrow_{\Gamma_1(-)} b, f \quad g(b') = \text{catb}(f(b'), b_2)}{\text{catb}(b_1, b_2) \rightsquigarrow_{\Gamma_1(-); \Gamma_2} b, g}}{\frac{b_1 \rightsquigarrow_{\Gamma_1(-)} b, f \quad g(b') = \text{parb}(f(b'), b_2)}{\text{parb}(b_1, b_2) \rightsquigarrow_{\Gamma_1(-), \Gamma_2} b, g}} \quad \frac{b_2 \rightsquigarrow_{\Gamma_2(-)} b, f \quad g(b') = \text{catb}(b_1, f(b'))}{\text{catb}(b_1, b_2) \rightsquigarrow_{\Gamma_1; \Gamma_2(-)} b, g}}{\frac{b_1 \rightsquigarrow_{\Gamma_1(-)} b, f \quad g(b') = \text{parb}(f(b'), b_2)}{\text{parb}(b_1, b_2) \rightsquigarrow_{\Gamma_1(-), \Gamma_2} b, g}}$$

Batch projection is deterministic, in the sense that the batch b and the context $\Gamma(-)$ determine b' and f .

THEOREM B.39 (BATCH PROJECTION DETERMINISM). *For all b , and $\Gamma(-)$, there is at most one pair (b', f) such that $b \rightsquigarrow_{\Gamma(-)} b', f$.*

PROOF. `pproj.v:bproj_det` □

The following establishes the fact that batch projection is a function: if the input is well-typed, there is an output pair which makes the relation hold (exactly one, by the previous theorem).

THEOREM B.40 (BATCH PROJECTION FUNCTION). *For all b , and $\Gamma(-)$, if $b : \text{batch}(\Gamma(\Delta))$, then there are b' and f such that $b \rightsquigarrow_{\Gamma(-)} b', f$.*

PROOF. `pproj.v:bproj_fun` □

This theorem establishes the correctness condition for batch projection. If the input is well-typed and the relation $b \rightsquigarrow_{\Gamma(-)} b', f$ holds, then b' has the expected type (Δ) , and the substitution function takes well-typed batches to well-typed batches.

THEOREM B.41 (BATCH PROJECTION CORRECTNESS). *If $b \rightsquigarrow_{\Gamma(-)} b', f$, then for all Δ , if $b : \text{batch}(\Gamma(\Delta))$, we have that $b' : \text{batch}(\Delta)$ and $f(b') = b$. Moreover, for any Δ' and b'' such that $b'' : \text{batch}(\Delta')$, we have $f(b'') : \text{batch}(\Gamma(\Delta'))$.*

PROOF. `pproj.v:brproj_correct` □

Prefix Projection. With prefixes, the projection story is a bit more complicated. If $p : \text{prefix}(\Gamma(\Delta))$, exactly one of three possibilities can occur. First, p can contain no Δ -data, like in the case $\Gamma(-) = \cdot; -,$ with $p = \text{catpA}(\text{epsp})$. Next, p may contain some prefix of Δ , like when $\Gamma(-) = -, \Gamma'$, and $p = (p_1, p_2)$. In this case, the prefix of Δ that p contains is precisely p_1 . Last, p may contain an entire batch of Δ , as in the case $\Gamma(-) = -; \cdot,$ with $p = \text{catpB}(b, \text{epsp})$. In this case, the batch in question is exactly b .

Because of this trichotomy, prefix projection actually consists of three mutually disjoint relations, exactly one of which must hold when p is well-typed. They are written as follows:

- $p @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-)$ holds when p contains none of Δ
- $p @ \Gamma(-) \rightsquigarrow p' @ \Gamma'(-) @ f$ holds when p contains a prefix p' of Δ
- $p @ \Gamma(-) \rightsquigarrow b @ f$ holds when p contains a batch b of Δ .

In some cases where we use prefix projection, we will also need to compute how the context $\Gamma(-)$ context changes when we take the derivative of $\Gamma(\Delta)$ with respect to the p . When $p \rightsquigarrow_{\Gamma(-)} \perp$, the derivative $\delta_p(\Gamma(\Delta))$ looks like $\Gamma'(\Delta)$: some of the surrounding context has been chipped away, but since p contains none of Δ , it remains untouched. When $p \rightsquigarrow_{\Gamma(-)} p'$, the derivative $\delta_p(\Gamma(\Delta))$ looks like $\Gamma'(\delta_{p'}(\Delta))$. The outer context has changed, and Δ has been reduced by p' . Lastly, when $p \rightsquigarrow_{\Gamma(-)} b$, the derivative $\delta_p(\Gamma(\Delta))$ is a hole-free context Γ' : because p contained a whole batch of Δ , the derivative context has no Δ left. In the first two cases, the semantics will need these $\Gamma'(-)$ s, and so we include computing them in the prefix projection relation.

These relations also include one feature that wasn't presented in the body of the paper: when some of Δ was found (a prefix or a batch), we also need a function for *substituting* a different prefix or a batch (respectively) back in – this is the f in the $p @ \Gamma(-) \rightsquigarrow p' @ \Gamma'(-) @ f$ and $p @ \Gamma(-) \rightsquigarrow b @ f$. For example, if $p = \text{parp}(p_1, p_2)$ and $\Gamma(-) = -, \Gamma'$, then $p @ \Gamma(-) \rightsquigarrow p_1 @ (-, \delta_{p_2}(\Gamma')) @ f$, and for all p' , we have that $f(p') = \text{parp}(p', p_2)$: the prefix p' replaces the projected out prefix.

Definition B.42 (Prefix Projection).

For a well-typed p , exactly one of these three relations must hold.

THEOREM B.43 (PREFIX PROJECTION TRICHOTOMY). *For any prefix p and $\Gamma(-)$, at most one of the following three things holds:*

- (1) $p @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-)$
- (2) $p @ \Gamma(-) \rightsquigarrow p' @ \Gamma'(-) @ f$, where f is a function from prefixes to prefixes.
- (3) $p @ \Gamma(-) \rightsquigarrow b @ f$, where f is a function from batches to prefixes.

Moreover, when $p : \text{prefix}(\Gamma(\Delta))$, exactly one of three holds.

PROOF. Existence is `pproj.v:pproj_tricot`, and uniqueness is `pproj.v:pproj_det`. □

The following theorem establishes the correctness of empty prefix projection. If there is no Δ in p , then p is actually a prefix of $\Gamma(\Delta')$ for any Δ' (because it contains only data for the surrounding $\Gamma(-)$), and all the derivatives line up.

$$\begin{array}{c}
\text{PPROJ-HERE} \\
\frac{p@- \rightsquigarrow p@ - @id}{p_1@_{\Gamma_1}(-) \rightsquigarrow p'@_{\Gamma'_1}(-)@f} \\
\text{PPROJ-,1-PFX} \\
\frac{p_1@_{\Gamma_1}(-) \rightsquigarrow p'@_{\Gamma'_1}(-)@f \quad f'(p'') = \text{parp}(f(p''), p_2) \quad \delta_{p_2}(\Gamma_2) \sim \Gamma'_2}{\text{parp}(p_1, p_2)@_{\Gamma_1}(-), \Gamma_2 \rightsquigarrow p'@_{\Gamma'_1}(-), \Gamma'_2@f'} \\
\text{PPROJ-,1-BAT} \\
\frac{p_1@_{\Gamma_1}(-) \rightsquigarrow b@f \quad f'(b') = \text{parp}(f(b'), p_2)}{\text{parp}(p_1, p_2)@_{\Gamma_1}(-), \Gamma_2 \rightsquigarrow b@f'} \\
\text{PPROJ-,2-PFX} \\
\frac{p_2@_{\Gamma_2}(-) \rightsquigarrow p'@_{\Gamma'_2}(-)@f \quad f'(p'') = \text{parp}(p_1, f(p'')) \quad \delta_{p_1}(\Gamma_1) \sim \Gamma'_1}{\text{parp}(p_1, p_2)@_{\Gamma_1}, \Gamma_2(-) \rightsquigarrow p'@_{\Gamma'_1}, \Gamma'_2(-)@f'} \\
\text{PPROJ-,2-BAT} \\
\frac{p_2@_{\Gamma_2}(-) \rightsquigarrow b@f \quad f'(p'') = \text{parp}(p_1, f(p''))}{\text{parp}(p_1, p_2)@_{\Gamma_1}(-), \Gamma_2 \rightsquigarrow b@f'} \\
\text{PPROJ-;1-A-PFX} \\
\frac{p@_{\Gamma_1}(-) \rightsquigarrow p'@_{\Gamma'_1}(-)@f \quad f'(p'') = \text{catpA}(f(p''))}{\text{catpA}(p)@_{\Gamma_1}(-); \Gamma_2 \rightsquigarrow p'@_{\Gamma'_1}(-); \Gamma_2@f'} \\
\text{PPROJ-;1-A-BAT} \\
\frac{p@_{\Gamma_1}(-) \rightsquigarrow b@f \quad f'(b') = \text{catpA}(f(b'))}{\text{catpA}(p)@_{\Gamma_1}(-); \Gamma_2 \rightsquigarrow b@f'} \\
\text{PPROJ-;1-B} \\
\frac{p@_{\Gamma_1}(-) \rightsquigarrow b'@f \quad f'(b'') = \text{catpB}(b, f(b'')) \quad \delta_p(\Gamma_2) \sim \Gamma'_2}{\text{catpB}(b, p)@_{\Gamma_1}(-); \Gamma_2 \rightsquigarrow b'@f'} \\
\text{PPROJ-;2-A} \\
\frac{\delta_p(\Gamma_1) \sim \Gamma'_1}{\text{catpA}(p)@_{\Gamma_1}; \Gamma_2(-) \rightsquigarrow \perp@_{\Gamma'_1}; \Gamma'_2(-)} \\
\text{PPROJ-;2-B-PFX} \\
\frac{p@_{\Gamma_2}(-) \rightsquigarrow p'@_{\Gamma'_2}(-)@f \quad f'(p'') = \text{catpB}(b, f(p''))}{\text{catpB}(b, p)@_{\Gamma_1}; \Gamma_2(-) \rightsquigarrow p'@_{\Gamma'_2}(-)@f'} \\
\text{PPROJ-;2-B-PFX} \\
\frac{p@_{\Gamma_2}(-) \rightsquigarrow b'@_{\Gamma'_2}(-)@f \quad f'(b'') = \text{catpB}(b, f(b''))}{\text{catpB}(b, p)@_{\Gamma_1}; \Gamma_2(-) \rightsquigarrow b'@_{\Gamma'_2}(-)@f'} \\
\text{PPROJ-;2-B-EMP} \\
\frac{p@_{\Gamma_2}(-) \rightsquigarrow \perp@_{\Gamma'_2}(-)}{\text{catpB}(b, p)@_{\Gamma_1}; \Gamma_2(-) \rightsquigarrow \perp@_{\Gamma'_2}(-)}
\end{array}$$

Fig. 8. Prefix Projection Rules

THEOREM B.44 (PREFIX PROJECTION EMPTY COHERENCE). *If $p : \text{prefix}(\Gamma(\Delta))$ and $p@_{\Gamma}(-) \rightsquigarrow \perp@_{\Gamma'}(-)$ then for all Δ' , we have $p : \text{prefix}(\Gamma(\Delta'))$ and $\delta_p(\Gamma(\Delta')) = \Gamma'(\Delta')$*

PROOF. pproj.v:pproj_emp □

The following theorem says that if we get a prefix p' out of prefix projection, then (1) it actually is a prefix of Δ , (2) substituting it back into p just gives p again, and (3) for any prefix p'' of a context Δ' , substituting it into p (i.e. $f(p'')$) gives a prefix of $\Gamma(\Delta')$ such that $\delta_{f(p'')}(\Gamma(\Delta'))$ is of the form $\Gamma'(\delta_{p''}(\Delta'))$, where Γ' does not depend on Δ' or p'' .

THEOREM B.45 (PREFIX PROJECTION PREFIX COHERENCE). *If $p : \text{prefix}(\Gamma(\Delta))$ and $p@_{\Gamma(-)} \rightsquigarrow p'@_{\Gamma'(-)}@f$, then*

- (1) $p' : \text{prefix}(\Delta)$
- (2) $f(p') = p$
- (3) *For any Δ' and any $p'' : \text{prefix}(\Delta')$, we have $f(p'') : \text{prefix}(\Gamma(\Delta'))$ and $\delta_{f(p'')}(\Gamma(\Delta')) = \Gamma'(\delta_{p''}(\Delta'))$*

PROOF. `pproj.v:pproj_pfx` □

The following theorem is basically the same story as Theorem B.45. If we look for the Δ bit of $p : \text{prefix}(\Gamma(\Delta))$ and get out a whole batch b , then (1) b actually has type Δ , (2) substituting b back into p just yields p again, and (3) for all contexts Δ' and all batches b' of Δ' , substituting b' into p gives a prefix of $\Gamma(\Delta')$, and all the derivatives of such substitutions are equal to a fixed context (without a hole) Γ' , which is given by the relation.

THEOREM B.46 (PREFIX PROJECTION BATCH COHERENCE). *If $p : \text{prefix}(\Gamma(\Delta))$ and $p@_{\Gamma(-)} \rightsquigarrow b@f$, then*

- (1) $b : \text{batch}(\Delta)$
- (2) $f(b) = p$
- (3) *For any Δ' and any $b' : \text{batch}(\Delta')$, we have $f(b') : \text{prefix}(\Gamma(\Delta'))$*

PROOF. `pproj.v:pproj_bat` □

Projection Concatenation Theorems. This section includes a handful of theorems about how prefix and batch projection relate to prefix concatenation. Given a prefix $p : \text{prefix}(\Gamma(\Delta))$ and a prefix $p' : \text{prefix}(\delta_p(\Gamma(\Delta)))$, and prefix projections for p and p' , how can we characterize the projection for the concatenation of p with p' ? This section answers all such questions. Note that if the prefix projection for p is empty or a prefix and gives back the new context $\Gamma'(-)$, then p' has type $\Gamma'(\dots)$, and so its projection happens in context $\Gamma'(-)$. All of the following theorems are required to prove subcases of the homomorphism theorem.

If a prefix p contains none of Δ and is followed by a batch b , then the resulting batch $p \cdot b$ has exactly the same batch projection behavior as b .

THEOREM B.47 (BATCH PROJECTION EMPTY CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma(-)} \rightsquigarrow \perp@_{\Gamma'(-)}$
- (3) $b : \text{batch}(\Gamma'(\Delta))$
- (4) $b \rightsquigarrow_{\Gamma'(-)} b', f$

Then, there is an f such that:

- (1) $p \cdot_{\Gamma(\Delta)} b \rightsquigarrow_{\Gamma'(-)} b', f'$
- (2) *For any b'' and Δ' with $b'' : \text{batch}(\Delta')$, we have $f'(b'') = p \cdot_{\Gamma(\Delta')} f(b'')$.*

PROOF. `pproj.v:bproj_hom_emp` □

If a prefix p contains a prefix of Δ and is followed by a batch b , then the batch projection of $p \cdot b$ is the concatenation of the projections of p and b , respectively, and the substitution functions also line up correctly.

THEOREM B.48 (BATCH PROJECTION PREFIX CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma(-)} \rightsquigarrow p_d@_{\Gamma'(-)}@f$

- (3) $b : \text{batch}(\Gamma'(\delta_{p_d}(\Delta)))$
- (4) $b \rightsquigarrow_{\Gamma'(-)} b_d, f'$

Then, there is some f'' such that

- (1) $p \cdot_{\Gamma(\Delta)} b \rightsquigarrow_{\Gamma(-)} p_d \cdot_{\Delta} b_s, f''$
- (2) For any $p'' b''$ and Δ' with $p'' : \text{prefix}(\Delta') b'' : \text{batch}(\delta_{p''}(\Delta'))$, we have $f''(b'') = f(p'') \cdot_{\Gamma(\Delta')} f'(b'')$.

PROOF. `pproj.v:bproj_hom_pfx` □

If a prefix p contains a batch of Δ and is followed by a batch b , then the batch projection of $p \cdot b$ is just the original batch of Δ .

THEOREM B.49 (BATCH PROJECTION BATCH CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p @ \Gamma(-) \rightsquigarrow b_d @ f$
- (3) $b : \text{batch}(\delta_p(\Gamma(\Delta)))$

Then, there is some f' such that

- (1) $p \cdot_{\Gamma(\Delta)} b \rightsquigarrow_{\Gamma(-)} b_d, f'$
- (2) For any $p' b'$ and Δ' with $p' : \text{prefix}(\Delta') b' : \text{batch}(\delta_p(\Gamma(\Delta')))$, we have $f'(b') = f(p') \cdot_{\Gamma(\Delta')} b'$.

PROOF. `pproj.v:bproj_hom_bat` □

If two subsequent prefixes p and p' both have empty projections, then their concatenation does as well.

THEOREM B.50 (PREFIX PROJECTION EMPTY EMPTY CONCATENATION). *Suppose the following*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-)$
- (3) $p' : \text{prefix}(\Gamma'(\Delta'))$
- (4) $p' @ \Gamma'(-) \rightsquigarrow \perp @ \Gamma''(-)$

Then, for any Δ'' , we have $p \cdot_{\Gamma(\Delta')} p' @ \Gamma(-) \rightsquigarrow \perp @ \Gamma''(-)$.

PROOF. `pproj.v:pproj_hom_emp_emp` □

If subsequent prefixes p and p' are such that p has an empty projection and p' has a prefix projection, then the projection of $p \cdot p'$ is that of p' .

THEOREM B.51 (PREFIX PROJECTION EMPTY PREFIX CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-)$
- (3) $p' : \text{prefix}(\Gamma'(\Delta'))$
- (4) $p' @ \Gamma'(-) \rightsquigarrow p_d @ \Gamma''(-) @ f$

Then, there exists f' such that

- (1) $p \cdot_{\Gamma(\Delta)} p' @ \Gamma(-) \rightsquigarrow p_d @ \Gamma''(-) @ f'$
- (2) For any Δ' and p'' such that $p'' : \text{prefix}(\Delta')$, we have $f'(p'') = p \cdot_{\Gamma(\Delta')} f(p'')$

PROOF. `pproj.v:pproj_hom_emp_pfx` □

If subsequent prefixes p and p' are such that p has an empty projection and p' has a batch projection, then the projection of $p \cdot p'$ is again that of p' .

THEOREM B.52 (PREFIX PROJECTION EMPTY BATCH CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma}(-) \rightsquigarrow \perp@_{\Gamma'}(-)$
- (3) $p' : \text{prefix}(\Gamma'(\Delta))$
- (4) $p'@_{\Gamma'}(-) \rightsquigarrow b_d@_f$

Then, there exists f' such that

- (1) $p \cdot_{\Gamma(\Delta)} p'@_{\Gamma}(-) \rightsquigarrow b_d@_{f'}$
- (2) *For any Δ' and b such that $b : \text{batch}(\Delta')$, we have $f'(b) = p \cdot_{\Gamma(\Delta')} f(b)$*

PROOF. `pproj.v:pproj_hom_emp_bat` □

It is impossible for subsequent prefixes p and p' to be such that p has a prefix projection and p' has an empty projection. Once a prefix has reached the hole part of the context, all subsequent prefixes have projection which is at least a prefix, if not a batch.

THEOREM B.53 (PREFIX PROJECTION PREFIX NOT EMPTY CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma}(-) \rightsquigarrow p_d@_{\Gamma'}(-)@_f$
- (3) $p' : \text{prefix}(\Gamma'(\delta_{p_d}(\Delta)))$

Then, it is never the case that $p'@_{\Gamma'}(-) \rightsquigarrow \perp@_{\Gamma''}(-)$, for any $\Gamma''(-)$

PROOF. `pproj.v:pproj_hom_pfx_not_emp` □

If subsequent prefixes p and p' both have prefix projections, then then the projection of $p \cdot p'$ is the concatenation of those of p and p' . The substitution function for $p \cdot p'$ is just the pointwise concatenation of the substitution functions for p and p' .

THEOREM B.54 (PREFIX PROJECTION PREFIX PREFIX CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma}(-) \rightsquigarrow p_d@_{\Gamma'}(-)@_f$
- (3) $p' : \text{prefix}(\Gamma'(\delta_{p_d}(\Delta)))$
- (4) $p'@_{\Gamma'}(-) \rightsquigarrow p'_d@_{\Gamma''}(-)@_{f'}$

Then, there exists f'' such that:

- (1) $p \cdot_{\Gamma(\Delta)} p'@_{\Gamma}(-) \rightsquigarrow p_d \cdot_{\Delta} p'_d@_{\Gamma''}(-)@_{f'}$
- (2) *For any p_0, p'_0 and Δ_0 , if $p_0 : \text{prefix}(\Delta_0)$ and $p'_0 : \text{prefix}(\delta_{p_0}(\Delta_0))$, then $f''(p_0 \cdot_{\Delta_0} p'_0) = f(p_0) \cdot_{\Gamma(\Delta_0)} f'(p'_0)$.*

PROOF. `pproj.v:pproj_hom_pfx_pfx` □

THEOREM B.55 (PREFIX PROJECTION PREFIX BATCH CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p@_{\Gamma}(-) \rightsquigarrow p_d@_{\Gamma'}(-)@_f$
- (3) $p' : \text{prefix}(\Gamma'(\delta_{p_d}(\Delta)))$
- (4) $p'@_{\Gamma'}(-) \rightsquigarrow b_d@_{f'}$

Then, there exists f'' such that:

- (1) $p \cdot_{\Gamma(\Delta)} p'@_{\Gamma}(-) \rightsquigarrow p_d \cdot_{\Delta} b_d@_{f'}$
- (2) *For any p_0, b_0 and Δ_0 , if $p_0 : \text{prefix}(\Delta_0)$ and $b_0 : \text{batch}(\delta_{p_0}(\Delta_0))$, then $f''(p_0 \cdot_{\Delta_0} b_0) = f(p_0) \cdot_{\Gamma(\Delta_0)} f'(b_0)$.*

PROOF. `pproj.v:pproj_hom_pfx_bat` □

THEOREM B.56 (PREFIX PROJECTION BATCH CONCATENATION). *Suppose the following:*

- (1) $p : \text{prefix}(\Gamma(\Delta))$
- (2) $p @ \Gamma(-) \rightsquigarrow b_d @ f$
- (3) $p' : \text{prefix}(\delta_p(\Gamma(\Delta)))$

Then, there exists f' such that:

- (1) $p \text{proj} \text{Batch} p \cdot_{\Gamma(\Delta)} p' \Gamma(-) b_d \delta_{p'}(\delta_p(\Gamma(\Delta))) f'$
- (2) *For all b' and Δ' such that $b' : \text{batch}(\Delta')$, we have that $f'(b') = f(b') \cdot_{\Gamma(\Delta')} p'$.*

PROOF. `pproj.v:pproj_hom_batch` □

B.6 Historical Contexts

Definition B.57 (Historical Context). Contexts $\Omega := \cdot | \Omega, x : A$ are fully structural contexts, where the A are STLC types.

A stream type is “flattened” into an STLC type by turning concatenations and parallels into products, and stars into lists.

Definition B.58 (Type and Context Flatten). For s a stream type, we define its flattening into an STLC type, denoted $\langle s \rangle$, inductively:

- $\langle 1 \rangle = 1$
- $\langle \varepsilon \rangle = 1$
- $\langle s \cdot t \rangle = \langle s \rangle \times \langle t \rangle$
- $\langle s || t \rangle = \langle s \rangle \times \langle t \rangle$
- $\langle s + t \rangle = \langle s \rangle + \langle t \rangle$
- $\langle s^* \rangle = \text{list}(\langle s \rangle)$

For Γ a bunched context, we define its flattening to a standard context, $\langle \Gamma \rangle$ inductively:

- $\langle \cdot \rangle = \cdot$
- $\langle x : s \rangle = x : \langle s \rangle$
- $\langle \Gamma; \Gamma' \rangle = \langle \Gamma \rangle, \langle \Gamma' \rangle$
- $\langle \Gamma, \Gamma' \rangle = \langle \Gamma \rangle, \langle \Gamma' \rangle$

For an STLC value $v : \langle s \rangle$, we write $\text{toBatch}_s(v)$ for the batch of type s that it corresponds to. Dually, for $b : \text{batch}(s)$, we write $\langle b \rangle$ for the STLC value of type $\langle s \rangle$ it corresponds to.

Definition B.59 (Historical Programs and Substitutions). Fix a language of terms M , with type system $\Omega \vdash M : A$. Write its semantics as $M \downarrow v$. We assume that this relation is a decidable partial function, in the sense that M evaluates to at most one v , and it is decidable whether or not such a v exists. We write substitutions $\theta : \Omega' \rightarrow \Omega$. Substitutions have a contravariant action on terms, written $M[\theta]$: if $\Omega \vdash M : A$, then $\Omega' \vdash M[\theta] : A$.

B.7 Context Subtyping

The following is a full listing of subtyping rules. This includes:

- The commutative monoid equations for comma contexts
- The non-commutative monoid equations for semicolon contexts.
- Weakenings for comma and semicolon. The semicolon weakenings need to record the context *not* being weakened away in the term for semantic reasons: transporting along these subtyping relations will require emitting empty prefixes, which require the context to construct.
- A rule `SubtyNullable(Δ)` which says that every context is a subtype of a context Δ if Δ is nullable. This is not intended to be used in terms (or needed in any we have written), but is needed to make the subtyping derivative operation work.

- A rule `SubtyWknAll` which is a global weakening: every context is a subtype of the empty context.
- A rule `SubtyDistr` which says that semicolon distributes over comma: $\Gamma; (\Delta_1, \Delta_2) \leq (\Gamma; \Delta_1), (\Gamma; \Delta_2)$. Semantically, we can turn a prefix of $\Gamma; (\Delta_1, \Delta_2)$ into $(\Gamma; \Delta_1), (\Gamma; \Delta_2)$ by copying the Γ to both parallel components, and then forwarding the Δ components. A program that has the same batch behavior as this is implementable without the structural rule (using `wait`), but it is less incremental, instead waiting until the entire Γ arrives, and then sending Γ all at once and forwarding the Δ s.
- A congruence rule, allowing subtyping rules to be applied in a subcontext.

Definition B.60 (Context Subtyping).

$$\begin{array}{c}
\frac{pf : \Delta \leq \Delta'}{\text{SubtyCong}(\Gamma(-), pf) : \Gamma(\Delta) \leq \Gamma(\Delta')} \\
\frac{}{\text{SubtyUnitSemicR} : \Gamma \leq \Gamma; \cdot} \\
\frac{}{\text{SubtyUnitCommaR} : \Gamma \leq \Gamma, \cdot} \\
\frac{}{\text{SubtyCommaComm} : \Gamma_1, \Gamma_2 \leq \Gamma_2, \Gamma_1} \\
\frac{}{\text{SubtyCommaWkn2} : \Gamma', \Gamma \leq \Gamma} \\
\frac{}{\text{SubtySemicWkn2}(\Gamma) : \Gamma'; \Gamma \leq \Gamma} \\
\frac{}{\text{SubtySemicAssoc1} : (\Gamma_1; \Gamma_2); \Gamma_3 \leq \Gamma_1; (\Gamma_2; \Gamma_3)} \\
\frac{}{\text{SubtyCommaAssoc1} : (\Gamma_1, \Gamma_2), \Gamma_3 \leq \Gamma_1, (\Gamma_2, \Gamma_3)} \\
\frac{}{\text{SubtyDistr} : \Gamma; \Delta, \Delta' \leq \Gamma; \Delta, \Gamma; \Delta'} \\
\frac{}{\text{SubtyUnitSemicL} : \Gamma \leq \cdot; \Gamma} \\
\frac{}{\text{SubtyUnitCommaL} : \Gamma \leq \cdot, \Gamma} \\
\frac{}{\text{SubtyCommaContr} : \Gamma \leq \Gamma, \Gamma} \\
\frac{}{\text{SubtyCommaWkn1} : \Gamma, \Gamma' \leq \Gamma} \\
\frac{}{\text{SubtySemicWkn1}(\Gamma) : \Gamma; \Gamma' \leq \Gamma} \\
\frac{}{\Gamma \text{ nullable}} \\
\frac{}{\text{SubtyWknAll} : \Gamma \leq \cdot} \\
\frac{}{\text{SubtyNullable}(\Gamma) : \Gamma' \leq \Gamma} \\
\frac{}{\text{SubtySemicAssoc2} : \Gamma_1; (\Gamma_2; \Gamma_3) \leq (\Gamma_1; \Gamma_2); \Gamma_3} \\
\frac{}{\text{SubtyCommaAssoc2} : \Gamma_1, (\Gamma_2, \Gamma_3) \leq (\Gamma_1, \Gamma_2), \Gamma_3}
\end{array}$$

Subtyping is deterministic in the sense that a choice of a subtyping rule and a left context Γ determine the context Δ for which $pf : \Gamma \leq \Delta$.

THEOREM B.61 (SUBTYPING PROOF AND LEFT CONTEXT DETERMINE RIGHT CONTEXT). *If $pf : \Gamma \leq \Delta$ and $pf : \Gamma \leq \Delta'$, then $\Delta = \Delta'$.*

PROOF. `ctxsubty.v:tckCtxSubtyPf_det`

□

Subtyping relations can be interpreted as functions which cast batches of the left context to those of the right context. In other words, a subtyping $pf : \Gamma \leq \Delta$ should denote a function $\text{batch}(\Gamma) \rightarrow \text{batch}(\Delta)$. We operationalize this with the following definition: the judgment $\text{coeBatch}(pf, b) \sim b'$ means that the subtyping proof term turns b (of type Γ) into b' (of type Δ).

Definition B.62 (Context Subtyping Batch Transport).

$$\begin{array}{c}
\frac{b \rightsquigarrow_{\Gamma(-)} b', f \quad \text{coeBatch}(pf, b') \sim b''}{\text{coeBatch}(\text{SubtyCong}(\Gamma(-), pf), b) \sim f(b'')} \\
\hline
\text{coeBatch}(\text{SubtyUnitSemicL}, \text{catb}(\text{epsb}, b)) \sim b \quad \text{coeBatch}(\text{SubtyUnitSemicR}, b) \sim \text{catb}(b, \text{epsb}) \\
\hline
\text{coeBatch}(\text{SubtyUnitCommaL}, \text{parb}(\text{epsb}, b)) \sim b \quad \text{coeBatch}(\text{SubtyUnitCommaR}, \text{parb}(b, \text{epsb})) \sim b \\
\hline
\text{coeBatch}(\text{SubtyCommaContr}, b) \sim \text{parb}(b, b) \quad \text{coeBatch}(\text{SubtyCommaComm}, \text{parb}(b_1, b_2)) \sim \text{parb}(b_2, b_1) \\
\hline
\text{coeBatch}(\text{SubtyCommaWkn1}, \text{parb}(b, b')) \sim b' \quad \text{coeBatch}(\text{SubtyCommaWkn2}, \text{parb}(b, b')) \sim b \\
\hline
\text{coeBatch}(\text{SubtySemicWkn1}(_), \text{catb}(b, b')) \sim b' \quad \text{coeBatch}(\text{SubtySemicWkn2}(_), \text{catb}(b, b')) \sim b \\
\hline
\frac{\text{coeBatch}(\text{SubtyWknAll}, b) \sim \text{epsb}}{\Gamma \text{ done } b'} \\
\hline
\text{coeBatch}(\text{SubtyNullable}(\Gamma), b) \sim b' \\
\hline
\text{coeBatch}(\text{SubtySemicAssoc1}, \text{catb}(\text{catb}(b_1, b_2), b_3)) \sim \text{catb}(b_1, \text{catb}(b_2, b_3)) \\
\hline
\text{coeBatch}(\text{SubtySemicAssoc2}, \text{catb}(b_1, \text{catb}(b_2, b_3))) \sim \text{catb}(\text{catb}(b_1, b_2), b_3) \\
\hline
\text{coeBatch}(\text{SubtyCommaAssoc1}, \text{parb}(\text{parb}(b_1, b_2), b_3)) \sim \text{parb}(b_1, \text{parb}(b_2, b_3)) \\
\hline
\text{coeBatch}(\text{SubtyCommaAssoc2}, \text{parb}(b_1, \text{parb}(b_2, b_3))) \sim \text{parb}(\text{parb}(b_1, b_2), b_3) \\
\hline
\text{coeBatch}(\text{SubtyDistr}, \text{catb}(b, \text{parb}(b_1, b_2))) \sim \text{parb}(\text{catb}(b, b_1), \text{catb}(b, b_2))
\end{array}$$

Naturally, this relation is functional: when the inputs are well-typed, there is a unique well-typed output.

THEOREM B.63 (CONTEXT SUBTYPING BATCH TRANSPORT FUNCTION). *For any pf, b , there is at most one b' such that $\text{coeBatch}(pf, b) \sim b'$. If $b : \text{batch}(\Gamma)$ and $pf : \Gamma \leq \Gamma'$, such a b' must exist, and it satisfies $b' : \text{batch}(\Gamma')$.*

PROOF. Existence by `ctxsubty.v:ctxSubtyPfBatchTransport_fun`, uniqueness by `ctxsubty.v:ctxSubtyPfBatchTransport_det`, well-typedness of b' by `ctxsubty.v:ctxSubtyPfBatchTransport_correct`. \square

Definition B.64 (Context Subtyping Prefix Transport).

$$\begin{array}{c}
\frac{p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)}{\text{coePfx}(\text{SubtyCong}(\Gamma(-), pf), p) \sim p} \\
\frac{p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)@f \quad \text{coePfx}(pf, p') \sim p''}{\text{coePfx}(\text{SubtyCong}(\Gamma(-), pf), p) \sim f(p'')} \\
\frac{p@{\Gamma}(-) \rightsquigarrow b@f \quad \text{coeBatch}(pf, b) \sim b'}{\text{coePfx}(\text{SubtyCong}(\Gamma(-), pf), p) \sim f(b')} \\
\text{coePfx}(\text{SubtyUnitSemicL}, p) \sim \text{catpB}(\text{epsb}, p) \\
\text{coePfx}(\text{SubtyUnitSemicR}, p) \sim \text{catpA}(p) \\
\text{coePfx}(\text{SubtyUnitCommaL}, p) \sim \text{parp}(\text{eps}, p) \\
\text{coePfx}(\text{SubtyUnitCommaR}, p) \sim \text{parp}(p, \text{eps}) \\
\text{coePfx}(\text{SubtyCommaContr}, p) \sim \text{parp}(p, p) \\
\text{coePfx}(\text{SubtyCommaComm}, \text{parp}(p, p')) \sim \text{parp}(p, p') \\
\text{coePfx}(\text{SubtyCommaWkn1}, \text{parp}(p, p')) \sim p' \\
\text{coePfx}(\text{SubtyCommaWkn2}, \text{parp}(p, p')) \sim p \\
\text{coePfx}(\text{SubtySemicWkn1}(\Gamma), \text{catpA}(p)) \sim \text{emp}_{\Gamma} \\
\text{coePfx}(\text{SubtySemicWkn1}(_), \text{catpB}(b, p)) \sim p \\
\text{coePfx}(\text{SubtySemicWkn2}(_), \text{catpA}(p)) \sim p \\
\text{coePfx}(\text{SubtySemicWkn2}(_), \text{catpB}(b, p)) \sim (b)^{\circ} \\
\text{coePfx}(\text{SubtyWknAll}, p) \sim \text{eps} \\
\text{coePfx}(\text{SubtyNullable}(\Gamma), p) \sim \text{emp}_{\Gamma} \\
\text{coePfx}(\text{SubtySemicAssoc1}, \text{catpA}(\text{catpA}(p))) \sim \text{catpA}(p) \\
\text{coePfx}(\text{SubtySemicAssoc1}, \text{catpA}(\text{catpB}(b, p))) \sim \text{catpB}(b, \text{catpA}(p)) \\
\text{coePfx}(\text{SubtySemicAssoc1}, \text{catpB}(\text{catb}(b, b'), p)) \sim \text{catpB}(b, \text{catpB}(b', p)) \\
\text{coePfx}(\text{SubtySemicAssoc2}, \text{catpA}(p)) \sim \text{catpA}(\text{catpA}(p)) \\
\text{coePfx}(\text{SubtySemicAssoc2}, \text{catpB}(b, \text{catpA}(p))) \sim \text{catpA}(\text{catpB}(b, p)) \\
\text{coePfx}(\text{SubtySemicAssoc2}, \text{catpB}(b, \text{catpB}(b', p))) \sim \text{catpB}(\text{catb}(b, b'), p) \\
\text{coePfx}(\text{SubtyCommaAssoc1}, \text{parp}(\text{parp}(p_1, p_2), p_3)) \sim \text{parp}(p_1, \text{parp}(p_2, p_3)) \\
\text{coePfx}(\text{SubtyCommaAssoc2}, \text{parp}(p_1, \text{parp}(p_2, p_3))) \sim \text{parp}(\text{parp}(p_1, p_2), p_3) \\
\text{coePfx}(\text{SubtyDistr}, \text{catpA}(p)) \sim \text{parp}(\text{catpA}(p), \text{catpA}(p)) \\
\text{coePfx}(\text{SubtyDistr}, \text{catpB}(b, \text{parp}(p_1, p_2))) \sim \text{parp}(\text{catpB}(b, p_1), \text{catpB}(b, p_2))
\end{array}$$

THEOREM B.65 (CONTEXT SUBTYPING PREFIX TRANSPORT FUNCTION). *For any pf and p , there is at most one p' such that $\text{coePfx}(pf, p) \sim p'$. If $p : \text{prefix}(\Gamma)$ and $pf : \Gamma \leq \Gamma'$, then such a p' exists, and $p' : \text{prefix}(\Gamma')$*

PROOF. Existence by `ctxsubty.v:ctxSubtyPfPrefixTransport_fun`,
 uniqueness by `ctxsubty.v:ctxSubtyPfPrefixTransport_det`.
 well-typedness of p' by `ctxsubty.v:ctxSubtyPfPrefixTransport_correct` □

The proof terms themselves also need to update with prefixes: if $p : \text{prefix}(\Gamma)$ and $pf : \Gamma \leq \Gamma'$, then we should be able to find $\delta_p(pf)$ so that $\delta_p(pf) : \delta_p(\Gamma) \leq \delta_{\text{coePfx}(pf,p)}\Gamma'$. However, it is sometimes the case that $\delta_p(\Gamma) = \delta_{\text{coePfx}(pf,p)}\Gamma'$, in which case the subtyping is no longer needed.

To this end, we define a pair of relations, written $\delta_p(pf) \sim \perp$ and $\delta_p(pf) \sim pf'$. The first defines the situation where after p , the contexts that pf relates agree and the subtyping is no longer needed. The second gives pf' , the updated proof term relating the derivative contexts, when necessary. We call these two relations the “none” and “some” subtyping proof term derivative relations because of the use of an option type in the Coq formalism.

Definition B.66 (Context Subtyping Proof Term Derivatives).

$$\begin{array}{c}
\text{CSPD-CONG-EMP} \\
\frac{p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)}{\delta_p(\text{SubtyCong}(\Gamma(-), pf)) \sim \text{SubtyCong}(\Gamma'(-), pf)} \\
\text{CSPD-CONG-PFX0} \\
\frac{p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)@f \quad \delta_{p'}(pf) \sim \perp}{\delta_p(\text{SubtyCong}(\Gamma(-), pf)) \sim \perp} \\
\text{CSPD-CONG-PFX1} \quad \frac{p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)@f \quad \delta_{p'}(pf) \sim pf'}{\delta_p(\text{SubtyCong}(\Gamma(-), pf)) \sim \text{SubtyCong}(\Gamma'(-), pf')} \quad \text{CSPD-CONG-BAT} \quad \frac{p@{\Gamma}(-) \rightsquigarrow b@f}{\delta_p(\text{SubtyCong}(\Gamma(-), pf)) \sim \perp} \\
\text{CSPD-UNITSEMICL} \quad \text{CSPD-UNITSEMICR} \\
\frac{\delta_p(\text{SubtyUnitSemicL}) \sim \perp \quad \delta_p(\text{SubtyUnitSemicR}) \sim \text{SubtyUnitSemicR}}{\delta_p(\text{SubtyUnitCommaL}) \sim \text{SubtyUnitCommaL} \quad \delta_p(\text{SubtyUnitCommaR}) \sim \text{SubtyUnitCommaR}} \\
\text{CSPD-COMMACONTR} \quad \text{CSPD-COMMACOMM} \\
\frac{\delta_p(\text{SubtyCommaContr}) \sim \text{SubtyCommaContr} \quad \delta_p(\text{SubtyCommaComm}) \sim \text{SubtyCommaComm}}{\delta_p(\text{SubtyCommaWkn1}) \sim \text{SubtyCommaWkn1} \quad \delta_p(\text{SubtyCommaWkn2}) \sim \text{SubtyCommaWkn2}} \\
\text{CSPD-SEMICWKNB-A} \quad \text{CSPD-SEMICWKNB-B} \\
\frac{\delta_{\text{catpA}(p)}(\text{SubtySemicWkn1}(\Delta)) \sim \text{SubtySemicWkn1}(\Delta) \quad \delta_{\text{catpB}(b,p)}(\text{SubtySemicWkn1}(\Delta)) \sim \perp}{\delta_p(\Delta) \sim \Delta'} \\
\frac{\delta_{\text{catpA}(p)}(\text{SubtySemicWkn2}(\Delta)) \sim \text{SubtySemicWkn2}(\Delta')}{\delta_{(b)^\circ}(\Delta) \sim \Delta'} \\
\text{CSPD-WKNALL} \quad \text{CSPD-NULLABLE} \\
\frac{\delta_{\text{catpB}(b,p)}(\text{SubtySemicWkn2}(\Delta)) \sim \text{SubtyNullable}(\Delta')}{\delta_p(\text{SubtyWknAll}) \sim \text{SubtyWknAll} \quad \delta_p(\text{SubtyNullable}(\Delta)) \sim \text{SubtyNullable}(\Delta)} \\
\text{CSPD-SEMICASSOCA-A-A} \\
\frac{\delta_{\text{catpA}(\text{catpA}(p))}(\text{SubtySemicAssoc1}) \sim \text{SubtySemicAssoc1}}{\delta_{\text{catpA}(\text{catpB}(b,p))}(\text{SubtySemicAssoc1}) \sim \perp \quad \delta_{\text{catpB}(b,p)}(\text{SubtySemicAssoc1}) \sim \perp} \\
\text{CSPD-SEMICASSOCB-B-A} \quad \text{CSPD-SEMICASSOCB-B-B} \\
\frac{\delta_{\text{catpA}(p)}(\text{SubtySemicAssoc2}) \sim \text{SubtySemicAssoc2}}{\delta_{\text{catpB}(b,\text{catpA}(p))}(\text{SubtySemicAssoc2}) \sim \perp \quad \delta_{\text{catpB}(b,\text{catpB}(b',p))}(\text{SubtySemicAssoc2}) \sim \perp} \\
\text{CSPD-COMMAASSOCA} \quad \text{CSPD-COMMAASSOCB} \\
\frac{\delta_p(\text{SubtyCommaAssoc1}) \sim \text{SubtyCommaAssoc1} \quad \delta_p(\text{SubtyCommaAssoc2}) \sim \text{SubtyCommaAssoc2}}{\delta_{\text{catpA}(p)}(\text{SubtyDistr}) \sim \text{SubtyDistr} \quad \delta_{\text{catpB}(b,p)}(\text{SubtyDistr}) \sim \perp} \\
\text{CSPD-DISTR-A} \quad \text{CSPD-DISTR-B}
\end{array}$$

THEOREM B.67 (CONTEXT SUBTYPING PROOF DERIVATIVE FUNCTION). *For any p and pf , there is at most one x (either some pf' or \perp) such that $\delta_p(pf) \sim x$.*

PROOF. Existence by `ctxsubty.v:ctxSubtyPfDeriv_fun`,
uniqueness by `ctxsubty.v:ctxSubtyPfDeriv_det`. □

THEOREM B.68 (CONTEXT SUBTYPING PROOF DERIVATIVE CORRECTNESS (SOME)). *If*

- (1) $p : \text{prefix } (\Gamma)$
- (2) $pf : \Gamma \leq \Gamma'$
- (3) $\text{coePfx } (pf, p) \sim p'$
- (4) $\delta_p(pf) \sim pf'$

Then $pf : \delta_p(\Gamma) \leq \delta_{p'}(\Gamma')$.

PROOF. `ctxsubty.v:ctxSubtyPfDeriv_correct_some` □

THEOREM B.69 (CONTEXT SUBTYPING PROOF DERIVATIVE CORRECTNESS (NONE)). *If*

- (1) $p : \text{prefix } (\Gamma)$
- (2) $pf : \Gamma \leq \Gamma'$
- (3) $\text{coePfx } (pf, p) \sim p'$
- (4) $\delta_p(pf) \sim \perp$

Then $\delta_p(\Gamma) = \delta_{p'}(\Gamma')$.

PROOF. `ctxsubty.v:ctxSubtyPfDeriv_correct_some` □

The following four theorems are required as lemmas for the homomorphism theorem. They describe what happens when you coerce two prefixes (or a prefix then a batch) along a subtyping relation one after the other. In short, it's equivalent to coercing the concatenated prefix (or batch). The exact theorem statements depend on the derivative of the proof term pf with respect to the first prefix p . If the result is none, the second batch or prefix does not need to be coerced, as it already has the right type on both sides of the subtyping. If the result is some pf' , the second batch or prefix gets transported across it.

THEOREM B.70 (PREFIX AND BATCH TRANSPORT CONCATENATION (NONE)). *Suppose the following:*

- (1) $pf : \Gamma \leq \Delta$
- (2) $p : \text{prefix } (\Gamma)$
- (3) $b : \text{batch } (\delta_p(\Gamma))$
- (4) $\text{coePfx } (pf, p) \sim p'$
- (5) $\delta_p(pf) \sim \perp$

Then, $\text{coeBatch } (pf, p \cdot_{\Gamma} b) \sim p' \cdot_{\Delta} b$

PROOF. `ctxsubty.v:ctxSubtyPfBatchTransport_hom_none` □

THEOREM B.71 (PREFIX AND BATCH TRANSPORT CONCATENATION (SOME)). *Suppose the following:*

- (1) $pf : \Gamma \leq \Delta$
- (2) $p : \text{prefix } (\Gamma)$
- (3) $b : \text{batch } (\delta_p(\Gamma))$
- (4) $\text{coePfx } (pf, p) \sim p'$
- (5) $\delta_p(pf) \sim pf'$
- (6) $\text{coeBatch } (pf', b) \sim b'$

Then, $\text{coeBatch } (pf, p \cdot_{\Gamma} b) \sim p' \cdot_{\Delta} b'$

PROOF. `ctxsubty.v:ctxSubtyPfBatchTransport_hom_some` □

THEOREM B.72 (PREFIX TRANSPORT CONCATENATION (NONE)). *Suppose:*

- $pf : \Gamma \leq \Delta$
- $p_1 : \text{prefix}(\Gamma)$
- $p_2 : \text{prefix}(\delta_{p_1}(\Gamma))$
- $\text{coePfx}(pf, p_1) \sim p'_1$
- $\delta_{p_1}(pf) \sim \perp$

Then, $\text{coePfx}(pf, p_1 \cdot_{\top} p_2) \sim p'_1 \cdot_{\Delta} p_2$

PROOF. `ctxsubty.v:ctxSubyPfPrefixTransport_hom_none` □

THEOREM B.73 (PREFIX TRANSPORT CONCATENATION (SOME)). *Suppose:*

- $pf : \Gamma \leq \Delta$
- $p_1 : \text{prefix}(\Gamma)$
- $p_2 : \text{prefix}(\delta_{p_1}(\Gamma))$
- $\text{coePfx}(pf, p_1) \sim p'_1$
- $\delta_{p_1}(pf) \sim pf'$
- $\text{coePfx}(pf', p_2) \sim p'_2$

Then, $\text{coePfx}(pf, p_1 \cdot_{\top} p_2) \sim p'_1 \cdot_{\Delta} p'_2$

PROOF. `ctxsubty.v:ctxSubyPfPrefixTransport_hom_some` □

THEOREM B.74 (SUBTYPING PROOF TERM DERIVATIVE CONCATENATION (NONE)). *Suppose:*

- (1) $pf : \Gamma \leq \Delta$
- (2) $p : \text{prefix}(\Gamma)$
- (3) $p' : \text{prefix}(\delta_p(\Gamma))$
- (4) $\delta_p(pf) \sim \perp$

Then, $\delta_{p \cdot_{\top} p'}(pf) \sim \perp$

PROOF. `ctxsubty.v:ctxSubyPfDeriv_hom_none` □

THEOREM B.75 (SUBTYPING PROOF TERM DERIVATIVE CONCATENATION (SOME)). *Suppose:*

- (1) $pf : \Gamma \leq \Delta$
- (2) $p : \text{prefix}(\Gamma)$
- (3) $p' : \text{prefix}(\delta_p(\Gamma))$
- (4) $\delta_p(pf) \sim pf'$
- (5) $\delta_{p'}(pf') \sim x$

Where x could either be some pf'' or \perp . Then, $\delta_{p \cdot_{\top} p'}(pf) \sim x$

PROOF. `ctxsubty.v:ctxSubyPfDeriv_hom_some` □

B.8 Type System

Definition B.76 (Recursion Signature). A recursion signature Σ is either empty (signaling that typechecking is not in the body of a recursive function), or the signature $\Omega \mid \Gamma \rightarrow s$ of a sequent which defines the recursive function we are currently checking the body of. $\Sigma ::= \emptyset \mid (\Omega \mid \Gamma \rightarrow s)$

Definition B.77 (Type System).

$$\begin{array}{c}
\text{VAR} \\
\hline
\Omega \mid \Gamma(x : s) \vdash_{\Sigma} (x : s @ \Gamma(-)) : s \\
\text{PAR-L} \\
\hline
\Omega \mid \Gamma(x : s, y : t) \vdash_{\Sigma} e : r \\
\text{CAT-L} \\
\hline
\Omega \mid \Gamma(z : s \parallel t) \vdash_{\Sigma} \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e : r \\
\text{CAT-R} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} e_1 : s \quad \Omega \mid \Delta \vdash_{\Sigma} e_2 : t \\
\hline
\Omega \mid \Gamma; \Delta \vdash_{\Sigma} (e_1; e_2) : s \cdot t \\
\text{PLUS-R-1} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} e : s \\
\text{EPS-R} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} \text{sink} : \varepsilon \\
\text{PLUS-R-2} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} e : t \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} \text{inr}(e) : s + t \\
\text{PLUS-L} \\
\hline
p : \text{prefix}(\Gamma(z : s + t)) \quad \Omega \mid \Gamma(x : s) \vdash_{\Sigma} e_1 : r \quad \Omega \mid \Gamma(y : t) \vdash_{\Sigma} e_2 : r \\
\hline
\Omega \mid \delta_p(\Gamma(z : s + t)) \vdash_{\Sigma} \text{case}_{\Gamma(-),s,t,r}(p; z, x.e_1, y.e_2) : r \\
\text{STAR-R-1} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} \text{nil} : s^* \\
\text{STAR-R-2} \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} e_1 : s \quad \Omega \mid \Delta \vdash_{\Sigma} e_2 : s^* \\
\hline
\Omega \mid \Gamma; \Delta \vdash_{\Sigma} e_1 :: e_2 : s^* \\
\text{STAR-L} \\
\hline
p : \text{prefix}(\Gamma(z : s^*)) \quad \Omega \mid \Gamma(\cdot) \vdash_{\Sigma} e_1 : r \quad \Omega \mid \Gamma(x : s; xs : s^*) \vdash_{\Sigma} e_2 : r \\
\hline
\Omega \mid \delta_p(\Gamma(z : s^*)) \vdash_{\Sigma} \text{case}_{\Gamma(-),s,r}(p; z, e_1, x.xs.e_2) : r \\
\text{WAIT} \\
\hline
p : \text{prefix}(\Gamma(\Delta)) \quad \Omega, \langle \Delta \rangle \mid \Gamma(\cdot) \vdash_{\Sigma} e : s \\
\hline
\Omega \mid \delta_p(\Gamma(\Delta)) \vdash_{\Sigma} \text{wait}_{\Gamma(-),\Delta,s}(p; e) : s \\
\text{FIX} \\
\hline
\Omega' = x_1 : A_1, \dots, x_n : A_n \quad \Omega' \mid \Gamma \vdash_{\Omega' \mid \Gamma \rightarrow s} e : s \quad \Omega \vdash M_i : A_i \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} \text{fix}(x_1 : A_1, \dots, x_n : A_n) @ [M_1, \dots, M_n].e : s \\
\text{REC} \\
\hline
\Omega' = x_1 : A_1, \dots, x_n : A_n \quad \Omega \vdash M_i : A_i \\
\hline
\Omega \mid \Gamma \vdash_{(\Omega' \mid \Gamma \rightarrow s)} \text{rec} @ [M_1 : A_1, \dots, M_n : A_n] : s \\
\text{CUT} \\
\hline
\Omega \mid \Delta \vdash_{\Sigma} e_1 : s \quad \Omega \mid \Gamma(x : s) \vdash_{\Sigma} e_2 : r \\
\hline
\Omega \mid \Gamma(\Delta) \vdash_{\Sigma} \text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2 : r \\
\text{SUBCTX} \\
\hline
pf : \Gamma \leq \Gamma' \Omega \mid \Gamma' \vdash_{\Sigma} e : s \\
\hline
\Omega \mid \Gamma \vdash_{\Sigma} \text{subctx}(pf, e) : s
\end{array}$$

Buffering Rules. Some typing rules are different from those presented in the body of the paper. In particular, the left rules for star and sums, as well as Wait, include a *buffer* in the term: a prefix of the input context, where we store inputs until we have received enough to run the term. For example, the WAIT rule has this buffer p , which we gather until it includes a whole batch of Δ .

$$\begin{array}{c}
\text{WAIT} \\
\hline
p : \text{prefix}(\Gamma(\Delta)) \quad \Omega, \langle \Delta \rangle \mid \Gamma(\cdot) \vdash_{\Sigma} e : s \\
\hline
\Omega \mid \delta_p(\Gamma(\Delta)) \vdash_{\Sigma} \text{wait}_{\Gamma(-),\Delta,s}(p; e) : s
\end{array}$$

The buffer is included in the syntax of the term. Additionally, the context in the conclusion is $\delta_p(\Gamma(\Delta))$. If we've buffered p of the input, the term is expecting the rest of the context. Users of the calculus need not worry about this detail: when writing programs and when the program starts

running, the buffer is empty: $p = \text{emp}_{\Gamma(\Delta)}$. By Theorem B.13, $\delta_p(\Gamma(\Delta)) = \Gamma(\Delta)$, which returns WAIT to the expected rule presented in the body of the paper. The other rules that include buffers are PLUS-L and STAR-L.

Generalized Fixpoints. To account for the inclusion of the historical context, the rules for defining recursive programs by fixpoint and making recursive calls are generalized from the ones in the body of the paper to the rules shown below. The main idea is the same: `fix` registers a term as a recursive function, and `rec` can be used inside the body of the term to refer back to the function being defined. Since the full λ^{ST} calculus includes memory by means of the historical context, we need a mechanism to (a) kick off a recursive function with initial values in bound to historical context, and (b) change values as we make recursive calls.

$$\frac{\text{FIX} \quad \frac{\Omega' = x_1 : A_1, \dots, x_n : A_n \quad \Omega' \mid \Gamma \vdash_{\Omega' \mid \Gamma \rightarrow s} e : s \quad \Omega \vdash M_i : A_i}{\Omega \mid \Gamma \vdash_{\Sigma} \text{fix}(x_1 : A_1, \dots, x_n : A_n) @ [M_1, \dots, M_n] . e : s}}{\text{REC} \quad \frac{\Omega' = x_1 : A_1, \dots, x_n : A_n \quad \Omega \vdash M_i : A_i}{\Omega \mid \Gamma \vdash_{(\Omega' \mid \Gamma \rightarrow s)} \text{rec} @ [M_1 : A_1, \dots, M_n : A_n] : s}}$$

When defining a recursive transformer with `FIX`, the programmer can choose a context Ω' , different from the current context Ω for the recursive function to run in. This context includes any accumulators that the transformer will maintain. At definition time, the programmer must provide initial values for each of the accumulator values. Formally, this looks like a list of terms M_i , one for each position in the context Ω' , typed in the defining context Ω . The context Ω' is also included in the updated recursion signature, $\Omega' \mid \Gamma$ tos. Then, when making recursive calls, the programmer decides how these values should be updated before continuing with the next iteration. As we will see in Appendix C, accumulator values are often updated with data computed from the stream that has arrived since the last recursive call was made. This new recursion term is written with `rec @ [M1, ..., Mn]`, where the terms M_i compute define the new accumulator values from the current historical context Ω .

B.9 Sink Terms

Once we have produced an entire batch $b : \text{batch}(s)$, a program e of type s needs to transition to a program emitting nothing: we compute this term from b with `sinkb`.

Definition B.78 (Sink Terms). We define sink terms by recursion on b , as follows:

- (1) $\text{sink}_{\text{epsb}} = \text{sink}_{\text{oneb}} = \text{sink}_{[]} = \text{sink}_{b_1 :: b_2} = \text{sink}$
- (2) $\text{sink}_{\text{parb}(b_1, b_2)} = (\text{sink}_{b_1}, \text{sink}_{b_2})$
- (3) $\text{sink}_{\text{catb}(b_1, b_2)} = \text{sink}_{b_2}$
- (4) $\text{sink}_{\text{sumbA}(b)} = \text{sink}_{\text{sumbB}(b)} = \text{sink}_b$

This program is closed, and has type we expect for a stream transformer that has just emitted an entire batch b of type s .

THEOREM B.79 (SINK TERM CORRECTNESS). *When $b : \text{batch}(s)$, we have $\cdot \mid \vdash_{\emptyset} \text{sink}_b : \delta_{(b)}^\circ(s)$.*

PROOF. `language.v:sinkexpr_tck` □

The relevant concatenation property of sink terms is that they only depend on the the shape of the type s after the batch has been emitted, so adding more to the beginning does not change anything.

THEOREM B.80 (SINK TERM CONCATENATION). *If $p : \text{prefix}(s)$ and $b : \text{batch}(\delta_p(s))$, then $\text{sink}_b = \text{sink}_{p \cdot s} b$.*

PROOF. `language.v:sinkepr'_hom` □

THEOREM B.81 (FIXPOINT SUBSTITUTION). *For e, e_f terms, we define $e[e_f/\text{use}]$ compositionally over the structure of e , with $(\text{rec } @ \left[\overline{M : A} \right]) [e_f/\text{use}] = \text{fix} \left(\overline{x : A} \right) @ \left[\overline{M} \right].e_f$*

Then, if $\Omega \mid \Gamma \vdash_{\Sigma \cup \{f:\Omega' \mid \Delta \rightarrow t\}} e : s$ and $\Omega' \mid \Delta \vdash_{(f:\Omega' \mid \Delta \rightarrow t)} e_f : t$, we have $\Omega \mid \Gamma \vdash_{\Sigma} e[e_f/\text{use}] : s$.

PROOF. `language.v:fixSubst_correct` □

B.10 Semantics

Batch Semantics. We define a whole-batch semantics of the language: given an input batch and a program e , we say that $b \Rightarrow e \Rightarrow^n b'$ if e produces output b' after doing at most n unfoldings of fixpoints.

Definition B.82 (Batch Semantics). We define a 4-place relation $b \Rightarrow e \Rightarrow^n b'$ as follows.

$$\begin{array}{c}
 \text{B-VAR} \\
 \frac{b \rightsquigarrow_{\Gamma(-)} b', _}{b \Rightarrow (x : s @ \Gamma(-)) \Rightarrow^n b'} \\
 \text{B-PAR-R} \\
 \frac{b \Rightarrow e_1 \Rightarrow^{n_1} b_1 \quad b \Rightarrow e_2 \Rightarrow^{n_2} b_2}{b \Rightarrow (e_1, e_2) \Rightarrow^{n_1+n_2} \text{parb}(b_1, b_2)} \\
 \text{B-CAT-R} \\
 \frac{b_1 \Rightarrow e_1 \Rightarrow^{n_1} b'_1 \quad b_2 \Rightarrow e_2 \Rightarrow^{n_2} b'_2}{\text{catb}(b_1, b_2) \Rightarrow (e_1; e_2) \Rightarrow^{n_1+n_2} \text{catb}(b'_1, b'_2)} \\
 \text{B-EPS-R} \\
 \frac{}{b \Rightarrow \text{sink} \Rightarrow^n \text{epsb}} \\
 \text{B-PLUS-R-1} \\
 \frac{}{b \Rightarrow e \Rightarrow^n b'} \\
 \text{B-PLUS-R-2} \\
 \frac{}{b \Rightarrow e \Rightarrow^n b'} \\
 \text{B-INL} \\
 \frac{}{b \Rightarrow \text{inl}(e) \Rightarrow^n \text{sumbA}(b')} \\
 \text{B-INR} \\
 \frac{}{b \Rightarrow \text{inr}(e) \Rightarrow^n \text{sumbB}(b')} \\
 \text{B-CASE-L-1} \\
 \frac{p \cdot \Gamma(z:s+t) \ b \sim b' \quad b' \rightsquigarrow_{\Gamma(-)} \text{sumbA}(b''), f \quad f(b'') \Rightarrow e_1 \Rightarrow^n b'''}{b \Rightarrow \text{case}_{\Gamma(-),s,t,r}(p; z, x.e_1, y.e_2) \Rightarrow^n b'''} \\
 \text{B-CASE-L-2} \\
 \frac{p \cdot \Gamma(z:s+t) \ b \sim b' \quad b' \rightsquigarrow_{\Gamma(-)} \text{sumbB}(b''), f \quad f(b'') \Rightarrow e_2 \Rightarrow^n b'''}{b \Rightarrow \text{case}_{\Gamma(-),s,t,r}(p; z, x.e_1, y.e_2) \Rightarrow^n b'''} \\
 \text{B-STAR-R-1} \\
 \frac{}{b \Rightarrow \text{nil} \Rightarrow^n []} \\
 \text{B-STAR-R-2} \\
 \frac{b_1 \Rightarrow e_1 \Rightarrow^{n_1} b'_1 \quad b_2 \Rightarrow e_2 \Rightarrow^{n_2} b'_2}{\text{catb}(b_1, b_2) \Rightarrow e_1 :: e_2 \Rightarrow^{n_1+n_2} b'_1 :: b'_2}
 \end{array}$$

$$\begin{array}{c}
\text{B-STAR-L-1} \\
\frac{p \cdot_{\Gamma(z:s^*)} b \sim b' \quad b' \rightsquigarrow_{\Gamma(z:s^*)} [], f \quad f(\text{epsb}) \Rightarrow e_1 \Rightarrow^n b''}{b \Rightarrow \text{case}_{\Gamma(-),s,r}(p; z, e_1, x.xs.e_2) \Rightarrow^n b''} \\
\text{B-STAR-L-2} \\
\frac{p \cdot_{\Gamma(z:s^*)} b \sim b' \quad b' \rightsquigarrow_{\Gamma(z:s^*)} b_1 :: b_2, f \quad f(\text{catb}(b_1, b_2)) \Rightarrow e_2 \Rightarrow^n b''}{b \Rightarrow \text{case}_{\Gamma(-),s,r}(p; z, e_1, x.xs.e_2) \Rightarrow^n b''} \\
\text{B-WAIT} \\
\frac{p \cdot_{\Gamma(\Delta)} b \sim b' \quad b' \rightsquigarrow_{\Gamma(-)} b_{\Delta}, f \quad f(\text{epsb}) \Rightarrow e [\langle b_{\Delta} \rangle / \langle \Delta \rangle] \Rightarrow^n b''}{b \Rightarrow \text{wait}_{\Gamma(-),\Delta,s}(p; e) \Rightarrow^n b''} \\
\text{B-SUBCTX} \quad \text{B-FIX} \\
\frac{\text{coeBatch}(pf, b) \sim b' \quad b' \Rightarrow e \Rightarrow^n b''}{b \Rightarrow \text{subctx}(pf, e) \Rightarrow^n b''} \quad \frac{M_i \downarrow v_i \quad b \Rightarrow e_f [e_f/\text{use}] \left[\frac{v_i}{x_i} \right] \Rightarrow^n b'}{b \Rightarrow \text{fix} \left(\frac{x_i : A_i}{x_i} \right) @ \left[\frac{M_i}{M_i} \right]. e_f \Rightarrow^{n+1} b'} \\
\text{B-CUT} \\
\frac{b \rightsquigarrow_{\Gamma(-)} b', f \quad b' \Rightarrow e_1 \Rightarrow^{n_1} b'' \quad f(b'') \Rightarrow e_2 \Rightarrow^{n_2} b'''}{b \Rightarrow \text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2 \Rightarrow^{n_1+n_2} b'''}
\end{array}$$

THEOREM B.83 (BATCH SEMANTICS GAS MONOTONICITY). *If $b \Rightarrow e \Rightarrow^n b'$ and $n' \geq n$, then $b \Rightarrow e \Rightarrow^{n'} b'$.*

PROOF. `language.v:batch_sem_mono` □

The batch semantics relation is deterministic in the sense that for any term e and any input batch b , it will only ever terminate with at most one value. Note that this is independent of the gas provided: any terminating evaluation will terminate with the same value.

THEOREM B.84 (BATCH SEMANTICS INPUTS DETERMINE OUTPUTS). *If $b \Rightarrow e \Rightarrow^{n_1} b_1$ and $b \Rightarrow e \Rightarrow^{n_2} b_2$, then $b_1 = b_2$.*

PROOF. `language.v:batch_sem_det` □

THEOREM B.85 (BATCH STEP CORRECTNESS). *Suppose the following:*

- (1) $\cdot \mid \Gamma \vdash_{\emptyset} e : s$
- (2) $b \Rightarrow e \Rightarrow^n b'$
- (3) $b : \text{batch}(\Gamma)$

Then, $b' : \text{batch}(s)$.

PROOF. `language.v:batch_pres` □

Incremental Semantics.

Definition B.86 (Incremental Semantics). We define a relation $p \Rightarrow e \Downarrow^n e' \Rightarrow p'$.

$$\begin{array}{c}
\text{P-VAR-1} \\
\frac{p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)}{p \Rightarrow (x : s@{\Gamma}(-)) \Downarrow^n (x : s@{\Gamma}'(-)) \Rightarrow \text{emp}_s} \\
\text{P-VAR-2} \\
\frac{p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)@f \quad \delta_{p'}(s) \sim s'}{p \Rightarrow (x : s@{\Gamma}(-)) \Downarrow^n (x : s'@{\Gamma}'(-)) \Rightarrow p'} \\
\text{P-VAR-3} \\
\frac{p@{\Gamma}(-) \rightsquigarrow b@{\Gamma}'@f \quad \delta_{(b)^\circ}(s) \sim s'}{p \Rightarrow (x : s@{\Gamma}(-)) \Downarrow^n \text{subctx}(\text{SubtyWknAll}, \text{sink}_b) \Rightarrow (b)^\circ} \\
\text{P-HISTPGM} \\
\frac{M \Downarrow v \quad b = \text{toBatch}_s(v)}{p \Rightarrow \langle M : s \rangle \Downarrow^n \text{subctx}(\text{SubtyWknAll}, \text{sink}_b) \Rightarrow (b)^\circ} \\
\text{P-PAR-R} \\
\frac{p \Rightarrow e_1 \Downarrow^{n_1} e'_1 \Rightarrow p_1 \quad p \Rightarrow e_2 \Downarrow^{n_2} e'_2 \Rightarrow p_2}{p \Rightarrow (e_1, e_2) \Downarrow^{n_1+n_2} (e'_1, e'_2) \Rightarrow \text{parp}(p_1, p_2)} \\
\text{P-PAR-L-1} \\
\frac{(p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)) \vee (p@{\Gamma}(-) \rightsquigarrow p'@{\Gamma}'(-)@_) \quad p \Rightarrow e \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e \Downarrow^n \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e' \Rightarrow p'} \\
\text{P-PAR-L-2} \\
\frac{p@{\Gamma}(-) \rightsquigarrow _@_p \Rightarrow e \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(x, y) = z \text{ in } e \Downarrow^n e' \Rightarrow p'} \\
\text{P-CAT-R-1} \\
\frac{p \Rightarrow e_1 \Downarrow^n e'_1 \Rightarrow p'}{\text{catpA}(p) \Rightarrow (e_1; e_2) \Downarrow^n (e'_1; e_2) \Rightarrow \text{catpA}(p')} \\
\text{P-CAT-R-2} \\
\frac{b \Rightarrow e_1 \Rightarrow^{n_1} b' \quad p \Rightarrow e_2 \Downarrow^{n_2} e'_2 \Rightarrow p'}{\text{catpB}(b, p) \Rightarrow (e_1; e_2) \Downarrow^{n_1+n_2} e'_2 \Rightarrow \text{catpB}(b', p')} \\
\text{P-CAT-L-1} \\
\frac{(p@{\Gamma}(-) \rightsquigarrow \perp@{\Gamma}'(-)) \vee (p@{\Gamma}(-) \rightsquigarrow \text{catpA}(_)@{\Gamma}'(-)@_) \quad p \Rightarrow e \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e \Downarrow^n \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e' \Rightarrow p'} \\
\text{P-CAT-L-2} \\
\frac{(p@{\Gamma}(-) \rightsquigarrow \text{catpB}(_, _)@_@_) \vee (p@{\Gamma}(-) \rightsquigarrow b@_)}{p \Rightarrow \text{let}_{\Gamma(-)}(z; x) = y \text{ in } e \Downarrow^n e' \Rightarrow p'} \\
\text{P-EPS-R} \\
\frac{}{p \Rightarrow \text{sink} \Downarrow^n \text{sink} \Rightarrow \text{eps}} \\
\text{P-PLUS-R-1} \\
\frac{}{p \Rightarrow e \Downarrow^n e' \Rightarrow p'} \\
\text{P-PLUS-R-2} \\
\frac{p \Rightarrow e \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{inl}(e) \Downarrow^n e' \Rightarrow \text{sumpA}(p')} \\
\frac{p \Rightarrow e \Downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{inr}(e) \Downarrow^n e' \Rightarrow \text{sumpB}(p')}
\end{array}$$

$$\begin{array}{c}
\text{P-PLUS-L-1} \\
\frac{p' \cdot \Gamma(z:s+t) \ p \sim p'' \quad (p'' @ \Gamma(-) \rightsquigarrow \perp @ _) \vee (p @ \Gamma(-) \rightsquigarrow \text{sumpEmp} @ _ @ _)}{p \Rightarrow \text{case}_{\Gamma(-),s,t,r} (p'; z, x.e_1, y.e_2) \downarrow^n \text{case}_{\Gamma(-),s,t,r} (p''; z, x.e_1, y.e_2) \Rightarrow \text{emp}_r} \\
\text{P-PLUS-L-2-1} \\
\frac{p' \cdot \Gamma(z:s+t) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{sumpA}(p_s) @ \Gamma'(-) @ f \quad f(p_s) \Rightarrow e_1 \downarrow^n e'_1 \Rightarrow p_r}{p \Rightarrow \text{case}_{\Gamma(-),s,t,r} (p'; z, x.e_1, y.e_2) \downarrow^n e'_1 \Rightarrow p_r} \\
\text{P-PLUS-L-2-2} \\
\frac{p' \cdot \Gamma(z:s+t) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{sumpB}(p_t) @ \Gamma'(-) @ f \quad f(p_t) \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p_r}{p \Rightarrow \text{case}_{\Gamma(-),s,t,r} (p'; z, x.e_1, y.e_2) \downarrow^n e'_2 \Rightarrow p_r} \\
\text{P-PLUS-L-3-1} \\
\frac{p' \cdot \Gamma(z:s+t) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{sumbA}(b) @ f \quad f(b) \Rightarrow e_1 \downarrow^n e'_1 \Rightarrow p_r}{p \Rightarrow \text{case}_{\Gamma(-),s,t,r} (p'; z, x.e_1, y.e_2) \downarrow^n e'_1 \Rightarrow p_r} \\
\text{P-PLUS-L-3-2} \\
\frac{p' \cdot \Gamma(z:s+t) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{sumbB}(b) @ f \quad f(b) \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p_r}{p \Rightarrow \text{case}_{\Gamma(-),s,t,r} (p'; z, x.e_1, y.e_2) \downarrow^n e'_2 \Rightarrow p_r} \\
\text{P-STAR-R-1} \\
\frac{p \Rightarrow \text{nil} \downarrow^n \text{sink} \Rightarrow \text{stpDone} \quad \text{catpA}(p) \Rightarrow e_1 :: e_2 \downarrow^n (e'_1; e_2) \Rightarrow \text{stpA}(p')}{p \Rightarrow e_1 \downarrow^n e'_1 \Rightarrow p'} \\
\text{P-STAR-R-2-2} \\
\frac{b \Rightarrow e_1 \Rightarrow^n b' \quad p \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p'}{\text{catpB}(b, p) \Rightarrow e_1 :: e_2 \downarrow^n e'_2 \Rightarrow \text{stpB}(b', p')} \\
\text{P-STAR-L-EMP} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad (p'' @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-)) \vee (p'' @ \Gamma(-) \rightsquigarrow \text{stpEmp} @ \Gamma'(-) @ f)}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n \text{case}_{\Gamma(-),s,t} (p''; z, e_1, x.xs.e_2) \Rightarrow \text{emp}_t} \\
\text{P-STAR-L-DONE-1} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{stpDone} @ \Gamma'(-) @ f \quad f(\text{epsp}) \Rightarrow e_1 \downarrow^n e'_1 \Rightarrow p'''}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n \text{subctx}(\text{SubtyCong}(\Gamma'(-), \text{SubtyWknAll}), e'_1) \Rightarrow p'''} \\
\text{P-STAR-L-DONE-2} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow [] @ f \quad f(\text{epsb}) \Rightarrow e_1 \downarrow^n e'_1 \Rightarrow p'''}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n e'_1 \Rightarrow p'''} \\
\text{P-STAR-L-CONS-1} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{stpA}(p_s) @ \Gamma'(-) @ f \quad f(p_s) \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p'''}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n \text{let}_{\Gamma(-)}(x; xs) = z \text{ in } e'_2 \Rightarrow p'''} \\
\text{P-STAR-L-CONS-2} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow \text{stpB}(b, p_s) @ \Gamma'(-) @ f \quad f(\text{catpB}(b, p_s)) \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p'''}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n e'_2 \Rightarrow p'''} \\
\text{P-STAR-L-CONS-3} \\
\frac{p' \cdot \Gamma(x:s^*) \ p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow b_1 :: b_2 @ f \quad f(\text{catb}(b_1, b_2)) \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p'''}{p \Rightarrow \text{case}_{\Gamma(-),s,t} (p'; z, e_1, x.xs.e_2) \downarrow^n e'_2 \Rightarrow p'''}
\end{array}$$

$$\begin{array}{c}
\text{P-WAIT-1} \\
\frac{p' \cdot_{\Gamma(\Delta)} p \sim p'' \quad (p'' @ \Gamma(-) \rightsquigarrow \perp @ _) \vee (p'' @ \Gamma(-) \rightsquigarrow _ @ _ @ _)}{p \Rightarrow \text{wait}_{\Gamma(-), \Delta, s}(p'; e) \downarrow^n \text{wait}_{\Gamma(-), \Delta, s}(p''; e) \Rightarrow \text{emp}_s} \\
\text{P-WAIT-2} \\
\frac{p' \cdot_{\Gamma(\Delta; \Delta')} p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow b_\Delta @ f \quad f(\text{epsb}) \Rightarrow e [\langle b_\Delta \rangle / \langle \Delta \rangle] \downarrow^n e' \Rightarrow p_s}{p \Rightarrow \text{wait}_{\Gamma(-), \Delta, s}(p; e) \downarrow^n e' \Rightarrow p_s} \\
\text{P-SUBCTX-SOME} \\
\frac{\text{coePfx}(pf, p) \sim p' \quad p' \Rightarrow e \downarrow^n e' \Rightarrow p'' \delta_p(pf) \sim pf'}{p \Rightarrow \text{subctx}(pf, e) \downarrow^n \text{subctx}(p', e') \Rightarrow p''} \\
\text{P-SUBCTX-NONE} \\
\frac{\text{coePfx}(pf, p) \sim p' \quad p' \Rightarrow e \downarrow^n e' \Rightarrow p'' \delta_p(pf) \sim \perp}{p \Rightarrow \text{subctx}(pf, e) \downarrow^n e' \Rightarrow p''} \\
\text{P-CUT-EMP} \\
\frac{p @ \Gamma(-) \rightsquigarrow \perp @ \Gamma'(-) \quad p \Rightarrow e_2 \downarrow^n e'_2 \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2 \downarrow^n \text{let}_{\Gamma'(-)} x = e_1 \text{ in } e'_2 \Rightarrow p'} \\
\text{P-CUT-PFX} \\
\frac{p @ \Gamma(-) \rightsquigarrow p' @ \Gamma'(-) @ f \quad p' \Rightarrow e_1 \downarrow^{n_1} e'_1 \Rightarrow p'' \quad f(p'') \Rightarrow e_2 \downarrow^{n_2} e'_2 \Rightarrow p'''}{p \Rightarrow \text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2 \downarrow^{n_1+n_2} \text{let}_{\Gamma'(-)} x = e'_1 \text{ in } e'_2 \Rightarrow p'''} \\
\text{P-CUT-BAT} \\
\frac{p @ \Gamma(-) \rightsquigarrow b @ f \quad b \Rightarrow e_1 \Rightarrow^{n_1} b' \quad f(b') \Rightarrow e_2 \downarrow^{n_2} e'_2 \Rightarrow p'}{p \Rightarrow \text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2 \downarrow^{n_1+n_2} e'_2 \Rightarrow p'} \\
\text{P-FIX} \\
\frac{M_i \downarrow v_i \quad p \Rightarrow e_f [e_f / \text{use}] \left[\frac{v_i}{x_i} \right] \downarrow^n e' \Rightarrow p'}{p \Rightarrow \text{fix}(x_i : A_i) @ [M_i].e_f \downarrow^{n+1} e' \Rightarrow p'}
\end{array}$$

Incremental Semantics of Buffering. The incremental semantics for PLUS-L and STAR-L and WAIT buffer in their inputs until enough of the input has arrived to run the term, where the particular value of “enough” depends on the rule in question.

To illustrate, consider the rules for Wait.

$$\begin{array}{c}
\text{P-WAIT-1} \\
\frac{p' \cdot_{\Gamma(\Delta)} p \sim p'' \quad (p'' @ \Gamma(-) \rightsquigarrow \perp @ _) \vee (p'' @ \Gamma(-) \rightsquigarrow _ @ _ @ _)}{p \Rightarrow \text{wait}_{\Gamma(-), \Delta, s}(p'; e) \downarrow^n \text{wait}_{\Gamma(-), \Delta, s}(p''; e) \Rightarrow \text{emp}_s} \\
\text{P-WAIT-2} \\
\frac{p' \cdot_{\Gamma(\Delta; \Delta')} p \sim p'' \quad p'' @ \Gamma(-) \rightsquigarrow b_\Delta @ f \quad f(\text{epsb}) \Rightarrow e [\langle b_\Delta \rangle / \langle \Delta \rangle] \downarrow^n e' \Rightarrow p_s}{p \Rightarrow \text{wait}_{\Gamma(-), \Delta, s}(p; e) \downarrow^n e' \Rightarrow p_s}
\end{array}$$

In both cases, we take the incoming prefix p , and concatenate it onto the buffer p' at the context type $\Gamma(\Delta)$, to get the combined prefix p'' . We then dispatch on whether p'' is enough input to run the continuation e . In this case, “enough” means that p'' contains a whole batch b of Δ . If it does (P-WAIT-2), we run the continuation, substituting the batch in. If it does not – p'' either gives an empty projection or just a prefix – we simply save p'' as the new buffer in the resulting wait term, and return the empty prefix in P-WAIT-1.

The semantics for PLUS-L and STAR-L are similar: in all cases, we add the incoming prefix to the buffer, and then project from the buffer. If not enough data has arrived, we return the empty prefix and step to the same term but with an updated buffer.

Incremental Semantics of Cut. To run a cut term $\text{let}_{\Gamma(-)} x = e_1 \text{ in } e_2$ incrementally, we use the prefix projection operation to pull out the part of the input prefix in the $\Gamma(-)$ hole, run it through e_1 , substitute the result back into the input prefix, and then run e_2 . This makes critical use of our prefix projection operation (Appendix B.5), and we have three incremental semantics rules corresponding to its three possibilities.

The first rule (P-CUT-1) covers the case where the input prefix p contains no Δ -data, and includes only data relevant to e_2 . So, we bypass e_1 entirely and run p through e_2 , stepping to e'_2 and producing output p' . The prefix p' is returned as the final output, and resulting term is the cut of e_1 with the updated term e'_2 . In the second case (P-CUT-2), the input prefix p contains a prefix p' of type Δ . We run this through e_1 , which steps to e'_1 and produces p'' . We then substitute p'' into p with the computed substitution function f , binding p'' to x in the context, and then run it through e_2 , which steps to e'_2 and produces p''' . The entire term then returns p''' , and steps to the let term with both e'_1 and e'_2 replacing e_1 and e_2 . In the third and final case (P-CUT-3), the input prefix contains a batch b of Δ . We then use the batch semantics to run this through e_1 , producing a batch b' . We then substitute b' into p where b was with f and run the resulting prefix through e_2 , which steps to e'_2 and produces p' . Because the part of the stream (Δ) used to produce the $x : s$ that the cut served to bind has completed, the cut is no longer needed, and so the whole term transitions to e'_2 , returning p' as output.

Incremental Semantics Theorems.

THEOREM B.87 (INCREMENTAL SEMANTICS INPUTS DETERMINE OUTPUTS). *If $p \Rightarrow e \downarrow^n e' \Rightarrow p'$ and $p \Rightarrow e \downarrow^{n'} e'' \Rightarrow p''$, then $e' = e''$, and $p' = p''$.*

PROOF. `language.v:prefix_sem_det` □

THEOREM B.88 (INCREMENTAL SEMANTICS MONOTONICITY). *If $p \Rightarrow e \downarrow^n e' \Rightarrow p'$ and $n' \geq n$, then $p \Rightarrow e \downarrow^{n'} e' \Rightarrow p'$.*

PROOF. `language.v:prefix_sem_mono` □

THEOREM B.89 (PREFIX CORRECTNESS). *Suppose*

- (1) $\cdot \mid \Gamma \vdash_0 e : s$
- (2) $p \Rightarrow e \downarrow^n e' \Rightarrow p'$
- (3) $p : \text{prefix}(\Gamma)$

Then, $p' : \text{prefix}(s)$ and $\cdot \mid \delta_p(\Gamma) \vdash_0 e' : \delta_{p'}(s)$.

PROOF. `language.v:prefix_pres` □

The following theorem proves that sink terms live up to their names. Given any input prefix p , the program sink_b will output an empty prefix of the appropriate type, and then step to itself.

THEOREM B.90 (SINK TERM SEMANTICS CHARACTERIZATION). *If $b : \text{batch}(s)$, then for all n and p , we have $p \Rightarrow \text{sink}_b \downarrow^n \text{sink}_b \Rightarrow \text{emp}_{\delta_{(b)^\circ}(s)}$.*

The following theorem is a lemma on the way to the homomorphism theorem. If e is run on a prefix p_1 , outputting p_2 and stepping to e' , and then e' is evaluated on an *batch* b_1 outputting a batch b_2 , then running e on the concatenated input $p_1 \cdot b_1$ yields the concatenated output $p_2 \cdot b_2$. The proof proceeds by induction on the derivation of e running on p_1 .

THEOREM B.91 (PREFIX/BATCH SEMANTICS CONCATENATION). *Suppose the following:*

- (1) $p_1 : \text{prefix}(\Gamma)$
- (2) $b_1 : \text{batch}(\delta_{p_1}(\Gamma))$

- (3) $p_2 : \text{prefix}(s)$
- (4) $\cdot \mid \Gamma \vdash_{\emptyset} e : s$
- (5) $p_1 \Rightarrow e \downarrow^{n_1} e' \Rightarrow p_2$
- (6) $b_1 \Rightarrow e' \Rightarrow^{n_2} b_2$

Then, $p_1 \cdot_{\top} b_1 \Rightarrow e \Rightarrow^{n_1+n_2} p_2 \cdot_s b_2$

PROOF. `language.v:prefix_batch_hom` □

THEOREM B.92 (HOMOMORPHISM THEOREM). *Suppose:*

- $\cdot \mid \Gamma \vdash_{\emptyset} e : s$
- $p_1 : \text{prefix}(\Gamma)$
- $p_2 : \text{prefix}(\delta_{p_1}(\Gamma))$
- $p_1 \Rightarrow e \downarrow^{n_1} e' \Rightarrow p'_1$
- $p_2 \Rightarrow e' \downarrow^{n_2} e'' \Rightarrow p'_2$

Then, $p_1 \cdot_{\top} p_2 \Rightarrow e \downarrow^{n_1+n_2} e'' \Rightarrow p'_1 \cdot_s p'_2$

PROOF. `language.v:prefix_hom` □

B.11 Events

Events (see Section 6) allow us to represent a data stream as a sequence of totally ordered items, while retaining information needed to infer the rich structure of the batch or prefix representations. to *Events*. The grammar of events is:

$$x ::= \text{oneev} \mid \text{parevA}(x) \mid \text{parevB}(x) \mid +_{\text{puncA}} \mid +_{\text{puncB}} \mid \cdot_{\text{punc}} \mid \text{catevA}(x)$$

Definition B.93 (Event Typing Relation). We define a binary relation $x : \text{event}(s)$ when $x \in \Sigma$ and s is a stream type.

$$\begin{array}{c}
 \frac{}{\text{oneev} : \text{event}(1)} \qquad \frac{x : \text{event}(s)}{\text{parevA}(x) : \text{event}(s \parallel t)} \qquad \frac{x : \text{event}(t)}{\text{parevB}(x) : \text{event}(s \parallel t)} \\
 \frac{}{+_{\text{puncA}} : \text{event}(s+t)} \qquad \frac{}{+_{\text{puncB}} : \text{event}(s+t)} \qquad \frac{s \text{ nullable}}{\cdot_{\text{punc}} : \text{event}(s \cdot t)} \\
 \frac{}{\text{catevA}(x) : \text{event}(s \cdot t)} \qquad \frac{}{+_{\text{puncA}} : \text{event}(s^*)} \qquad \frac{}{+_{\text{puncB}} : \text{event}(s^*)}
 \end{array}$$

Note that $s+t$ and s^* share the same punctuation events. Intuitively, this is because s^* can be unrolled as $\varepsilon + (s \cdot s^*)$.

Definition B.94 (Event Derivative Relation). We define a ternary relation $\delta_x(s) \sim s'$.

$$\begin{array}{c}
 \frac{}{\delta_{\text{oneev}}(1) \sim \varepsilon} \qquad \frac{\delta_x(s) \sim s'}{\delta_{\text{parevA}(x)}(s \parallel t) \sim s' \parallel t} \qquad \frac{\delta_x(t) \sim t'}{\delta_{\text{parevB}(x)}(s \parallel t) \sim s \parallel t'} \qquad \frac{}{\delta_{+_{\text{puncA}}}(s+t) \sim s} \\
 \frac{}{\delta_{+_{\text{puncB}}}(s+t) \sim t} \qquad \frac{s \text{ nullable}}{\delta_{\cdot_{\text{punc}}}(s \cdot t) \sim t} \qquad \frac{\delta_x(s) \sim s'}{\delta_{\text{catevA}(x)}(s \cdot t) \sim s' \cdot t} \qquad \frac{}{\delta_{+_{\text{puncA}}}(s^*) \sim \varepsilon} \\
 \frac{}{\delta_{+_{\text{puncB}}}(s^*) \sim s \cdot s^*}
 \end{array}$$

THEOREM B.95 (EVENT DERIVATIVE FUNCTION). *If $x : \text{event}(s)$, there is a unique s' such that $\delta_x(s) \sim s'$.*

PROOF. events.v:derivrelEv_fun □

Because of Theorem B.95, if we know $x : \text{event}(s)$, we may write the unique s' such that $\delta_x(s) \sim s'$ as $\delta_x(s)$.

Definition B.96 (Events Typing and Derivatives Relations). We lift event typing to lists by derivatives.

$$\frac{}{[] : \text{events}(s)} \qquad \frac{x : \text{event}(s) \quad \delta_x(s) \sim s' \quad xs : \text{events}(s')}{x :: xs : \text{events}(s)}$$

We also lift derivatives to lists of events in the natural way.

$$\frac{}{\delta_{[]} (s) \sim s} \qquad \frac{\delta_x(s) \sim s' \quad \delta_{xs}(s') \sim s''}{\delta_{x::xs}(s) \sim s''}$$

THEOREM B.97 (EVENTS DERIVATIVE FUNCTION). *If $xs : \text{events}(s)$, there is a unique s' such that $\delta_{xn}(s) \sim s'$.*

PROOF. events.v:derivrelEv_list_fun □

Because of Theorem B.97, if we know $xs : \text{events}(s)$, we may write the unique s' such that $\delta_{xs}(s) \sim s'$ as $\delta_{xs}(s)$.

THEOREM B.98 (EMPTY LIST OF EVENTS). *For all s , we have $[] : \text{events}(s)$, and $\delta_{[]} (s) \sim s$*

PROOF. Immediate. □

THEOREM B.99 (EVENTS CONCATENATION). *If*

- (1) $xs : \text{events}(s)$
- (2) $\delta_{xs}(s) \sim s'$
- (3) $ys : \text{events}(s')$
- (4) $\delta_{ys}(s') \sim s''$

Then, $xs++ys : \text{events}(s)$, and $\delta_{xs++ys}(s) \sim s''$.

In other words, if $xs : \text{events}(s)$ and $ys : \text{events}(\delta_{xs}(s))$, then $xs++ys : \text{events}(s)$ and $\delta_{xs++ys}(s) = \delta_{ys}(\delta_{xs}(s))$.

PROOF. events.v:Ev_list_concat □

Events

Events to Prefix

Definition B.100 (Event(s) to Prefix).

$$\frac{}{\text{onev} \hookrightarrow^1 \text{onepB}} \qquad \frac{x \hookrightarrow^s p}{\text{parevA}(x) \hookrightarrow^{s||t} \text{parp}(p, \text{emp}_t)} \qquad \frac{x \hookrightarrow^t p}{\text{parevB}(x) \hookrightarrow^{s||t} \text{parp}(\text{emp}_s, p)}$$

$$\frac{x \hookrightarrow^s p}{\text{catevA}(x) \hookrightarrow^{s \cdot t} \text{catpA}(p)} \qquad \frac{s \text{ nullable} \quad s \text{ done } b}{\text{punc} \hookrightarrow^{s \cdot t} \text{catpB}(b, \text{emp}_t)} \qquad \frac{}{\text{tpuncA} \hookrightarrow^{s \cdot t} \text{sumpA}(\text{emp}_s)}$$

$$\frac{}{\text{tpuncB} \hookrightarrow^{s \cdot t} \text{sumpB}(\text{emp}_t)} \qquad \frac{}{\text{tpuncA} \hookrightarrow^{s^*} \text{stpDone}} \qquad \frac{}{\text{tpuncB} \hookrightarrow^{s^*} \text{stpA}(\text{emp}_s)}$$

$$\frac{}{[] \hookrightarrow^s \text{emp}_s} \qquad \frac{t \hookrightarrow^s p \quad \delta_p s \sim s' \quad ts \hookrightarrow^{s'} p' \quad p' \cdot p \sim p''}{t :: ts \hookrightarrow^s p''}$$

THEOREM B.101 (EVENT TO PREFIX FUNCTION). *If $e \in \text{event}(s)$ then there is a unique $p \in \text{prefix}(s)$ such that $x \hookrightarrow^s p$.*

PROOF. `events.v:evToPfx_fun` □

THEOREM B.102 (EVENTS TO PREFIX FUNCTION). *If $xs : \text{events}(s)$, then there is a unique $p \in \text{prefix}(s)$ such that $xs \hookrightarrow^s p$.*

PROOF. `events.v:evToPfx_list_fun` □

THEOREM B.103 (ETOP RELATION). *If $x : \text{event}(s)$ and $x \hookrightarrow^s p$ then $p \in \text{prefix}(s)$.*

PROOF. `events.v:evToPfx_correct` □

THEOREM B.104 (ESTOP RELATION). *If $xs : \text{events}(s)$ and $xs \hookrightarrow^s p$ then $p \in \text{prefix}(s)$.*

PROOF. `events.v:evToPfx_list_correct` □

Prefix to Events

Definition B.105 (PToES). In this definition, we write occasionally lift event constructors to lists, writing $\widehat{f}(xs)$ for $[f(x) \mid x \in xs]$. Also, we write $xs||ys$ for the set of *shuffles* of the lists xs and ys .

$$\begin{array}{c}
\begin{array}{ccc}
\text{PToES-Eps} & \text{PToES-ONE-A} & \text{PToES-ONE-B} \\
\hline
\text{epsp } \dagger^\epsilon \ [] & \text{onepA } \dagger^1 \ [] & \text{onepB } \dagger^1 \ [\text{oneev}] \\
\text{PToES-PAR} & & \text{PToES-CAT-A} \\
\frac{p \dagger^s xs \quad p' \dagger^t ys \quad zs \in \widehat{\text{parevA}}(xs) \parallel \widehat{\text{parevB}}(ys)}{\text{parp}(p, p') \dagger^{s||t} zs} & & \frac{p \dagger^s xs}{\text{catpA}(p) \dagger^{s \cdot t} \widehat{\text{catevA}}(xs)} \\
\text{PToES-CAT-B} & & \text{PToES-SUM-A} \\
\frac{b^\circ \dagger^s xs \quad p \dagger^t ys}{\text{catpB}(b, p) \dagger^{s \cdot t} \widehat{\text{catevA}}(xs) ++ \text{punc} :: ys} & \frac{\text{PToES-SUM-EMP}}{\text{sumpEmp } \dagger^{s+t} \ []} & \frac{p \dagger^s xs}{\text{sumpA}(p) \dagger^{s+t} \text{puncA} :: xs} \\
\text{PToES-SUM-B} & & \text{PToES-STAR-EMP} \\
\frac{p \dagger^t xs}{\text{sumpB}(p) \dagger^{s+t} \text{puncB} :: xs} & \frac{\text{PToES-STAR-EMP}}{\text{stpEmp } \dagger^{s^*} \ []} & \frac{\text{PToES-STAR-DONE}}{\text{stpDone } \dagger^{s^*} \text{puncA}} \\
\text{PToES-STAR-A} & \text{PToES-STAR-B} & \\
\frac{p \dagger^s xs}{\text{stpA}(p) \dagger^{s^*} \text{puncB} :: \widehat{\text{catevA}}(xs)} & \frac{b^\circ \dagger^s xs \quad p \dagger^{s^*} ys}{\text{stpB}(b, p) \dagger^{s^*} \text{puncB} :: \widehat{\text{catevA}}(xs) ++ \text{punc} :: ys} &
\end{array}
\end{array}$$

THEOREM B.106 (PTOES EMPTY). *If $p \dagger^s xs$, then $xs = []$ iff $p = \text{emp}_s$*

PROOF. `events.v:PfxToEvs_Empty` and `events.v:pfxToEvs_empty'` □

THEOREM B.107 (PTOES LEFT TOTAL). *If $p \in \text{prefix}(s)$ then there exists (not necessarily unique) xs such that $p \dagger^s xs$*

PROOF. `events.v:PfxToEvs_left_total` □

LEMMA B.108 (PTOES RELATION DERIVATIVE AGREEMENT). *If*

(1) $p : \text{prefix}(s)$

(2) $p \dagger^s xs$

(3) $\delta_p(s) \sim s'$

then $xs : \text{events}(s)$ and $\delta_{xs}(s) \sim s'$

PROOF. `events.v:pfxToEvs_derivrelEv_correct_and_conf` □

The preceding lemma provides the following:

THEOREM B.109 (PTOES RELATION). *If $p \in \text{prefix}(s)$ and $p \dagger^s xs$ then $xs : \text{events}(s)$*

PROOF. `events.v:pfxToEvs_correct` □

Event Size

Each event carries tag information about where it appears within a structured stream; this is necessary for us to recover the rich prefix structure. Importantly, for a given stream type there is an upper bound on the amount of tag information to be included on any event in any stream of that type.

Definition B.110 (Event size). We define the *size* of an event recursively:

$$\begin{aligned}
 \text{size}(\text{oneev}) &= 1 \\
 \text{size}(\text{parevA}(x)) &= 1 + \text{size}(x) \\
 \text{size}(\text{parevB}(x)) &= 1 + \text{size}(x) \\
 \text{size}(+\text{puncA}) &= 1 \\
 \text{size}(+\text{puncB}) &= 1 \\
 \text{size}(\text{'punc}) &= 1 \\
 \text{size}(\text{catevA}(\cdot)(x)) &= 1 + \text{size}(x)
 \end{aligned}$$

We lift this to lists of events in the natural way:

Definition B.111 (Event List Size).

$$\begin{aligned}
 \text{size}([]) &= 0 \\
 \text{size}(x : xs) &= \text{size}(x) + \text{size}(xs)
 \end{aligned}$$

To construct an *a priori* bound on the size of any event to appear in stream, we recurse on the type of the stream:

Definition B.112 (Event size bound).

$$\begin{aligned}
 \text{evSizeBound}(\varepsilon) &= 0 \\
 \text{evSizeBound}(1) &= 1 \\
 \text{evSizeBound}(s \parallel t) &= 1 + \max(\text{evSizeBound}(s), \text{evSizeBound}(t)) \\
 \text{evSizeBound}(s \cdot t) &= \max(1 + \text{evSizeBound}(s), \text{evSizeBound}(t)) \\
 \text{evSizeBound}(s + t) &= \max(1, \text{evSizeBound}(s), \text{evSizeBound}(t)) \\
 \text{evSizeBound}(s^*) &= \max(1, 1 + \text{evSizeBound}(s))
 \end{aligned}$$

We now re-state and prove Theorem 6.1:

THEOREM B.113 (BOUNDED EVENT SIZE). *For all s , there is some $N = \text{evSizeBound}(\cdot)(s)$ such that for any $xs : \text{events}(s)$ and any $x \in xs$, we have that $|x| \leq N$, where $|\cdot|$ denotes the size of the AST.*

PROOF. `events.v:sizeBound_correct` □

Serialization and Deserialization

We turn now to the final result of 6, that we can serialize a prefix p into a list of events xs , secure in the knowledge that when we deserialize xs we will obtain the same prefix p .

Towards this result, we introduce a series of lemmas that allow us to use the tag information encoded in each event to recover the prefix structure during deserialization. Observe that the shape of each lemma mirrors that of the corresponding serialization (\dagger) constructor.

LEMMA B.114 (ES TOP PAR RECOVERY). *If*

$$(1) \ zs \in \widehat{\text{parevA}}(xs) \parallel \widehat{\text{parevB}}(ys)$$

(2) $xs \hookrightarrow^s p$

(3) $ys \hookrightarrow^t p'$

then $zs \hookrightarrow^{s \parallel t} \text{parp}(p, p')$

PROOF. `events.v:evToPfx_list_par_structural` □

LEMMA B.115 (ESToP CAT RECOVERY). *If $xs \hookrightarrow^s (b)^\circ$ and $ys \hookrightarrow^t p$ then*

$$\widehat{\text{catevA}}(xs) ++ \text{punc} :: ys \hookrightarrow^{s \cdot t} \text{catpB}(b, p)$$

PROOF. `events.v:evToPfx_list_cat_structural` □

LEMMA B.116 (ESToP STAR RECOVERY). *If $xs \hookrightarrow^s (b)^\circ$ and $ys \hookrightarrow^{s^*} p$ then*

$$\text{puncB} :: \widehat{\text{catevA}}(xs) ++ \text{punc} :: ys \hookrightarrow^{s^*} \text{stpB}(b, p)$$

PROOF. `events.v:evToPfx_list_star_structural` □

THEOREM B.117 (SERIALIZATION/DESERIALIZATION ROUND TRIP). *If $p \dagger^s xs$, then $xs \hookrightarrow^s p$.*

PROOF. `events.v:roundTrip_dir1` □

C EXAMPLES

This appendix collects some additional examples of programming with our stream types.

C.1 Examples Illustrating Kernel λ^{ST}

The parallel swap term from Section 3 is:

$$z : s \parallel t \vdash \text{let}_- (x; y) = z \text{ in } ((y : t @ (x : s, -)), (x : s @ -, y : t)) : t \parallel s$$

And the “prepend” term is:

$$x : s \vdash \text{subctx} (\text{SubtyUnitSemicL}, (()); (x : s @ -)) : 1 \cdot s$$

C.2 Map

Below is the sugared term for map, as presented in Section 5.

```
map (e : s → t) (xs : s*) : t* =
  fix. case xs of
    | nil => nil
    | y :: ys => e(y) :: rec(ys)
```

The Haskell encoding of the core term can be found in `Events.hs:maptm`.

C.3 Filter

Here is the code for the `filter` transformer described in Section 5.

```
filter (e : s → 2) (xs : s*) : s* =
  fix. case xs of
    | nil => nil
    | y :: ys => let b = e(y) in
      case b of
        | inl() => y :: rec(ys)
        | inr() => rec(ys)
```

The core term corresponding to this one can be found at `Examples.hs:filtertm`.

C.4 Running Fold

We can also define a *running* fold operation on star streams, which outputs its partial results as it goes.

```
runningfold (e : ⟨t⟩ | s → t) (xs : s*) : t* =
  fix with (acc = init).
  case xs of
    | nil => nil
    | y :: ys =>
      let acc' = e(y)[acc] in
      wait acc' then
        ⟨ acc' ⟩ :: rec(ys)[acc']
      end
```

This core term can be found at `Examples.hs:foldtm`.

C.5 Thresholding

The `parseLL` transformer parses a homogeneous stream of ints — as it would be sent from an actual IOT device — into the parsed representation that we can use to define the per-run averaging operation from Section 5. This transformer consumes the input stream, grouping together runs of

elements larger than the threshold t . The program operates the same as the functional program of type $[Int] \rightarrow [NonEmpty Int]$ which computes these runs by way of a fold. The function `parseLL` consumes the input stream until a run starts (an element is larger than the threshold t), and then starts running an auxiliary function `parseLLAux` which actually builds the runs.

```

parseLL (t : int) (xs : int*) : (int · int*)* =
  fix. case xs of
    | nil => nil
    | y :: ys => wait y then
      case ⟨ y > t ⟩ of
        | true => let (zs,zss) = parseLLAux t xs in (y;zs);zss
        | false => rec(ys)
      end

```

`parseLLAux` takes in a stream xs of ints, and returns (1) the maximal initial prefix of the stream which is all larger than the threshold, followed by (2) the stream of runs after the first. This strengthened specification is required to define the operation as a fold-left. `parseLLAux` maintains a Boolean b which indicates if the stream is currently in the middle of a run.

```

parseLLAux (t : int) (xs : int*) : int* · (int · int*)* =
  fix with (b = true).
    case xs of
      | nil => (nil ; nil)
      | y :: ys => wait y then
        case ⟨ y > t ⟩ of
          | true => let (zs,zss) = rec(ys)[true] in
              case b of
                | true => (y;zs;zss)
                | false => (nil;(y;zs)::zss)
          | false => rec(ys)[false]
        end
    end

```

The program begins by casing on the input, and buffering in the head. If the head is less than the threshold, we simply continue. If the head is above the threshold, then we make a recursive call to compute the runs of the rest. If we are already on a run (b is true), we cons y onto the current head zs , and return $(y;zs;zss)$. If we were not on a run, then the initial prefix is empty, and we return $(nil;(y;zs)::zss)$. The core terms for this program can be found at `Examples.hs:parseLL` and `Examples.hs:parseLLAux`.

C.6 Partitioning

A crucial streaming idiom is partitioning, where a homogeneous stream of data is split into two or more parallel streams, which are then routed to different downstream nodes in the dataflow graph. The purpose of partitioning is to expose parallelism: the different downstream operators can be run separately, potentially on different physical machines. Depending on the situation, a programmer may choose to use different partitioning strategies. In λ^{ST} , some common partitioning strategies are implementable as transformers.

Round Robin Partitioning. A round-robin partitioner fairly distributes an incoming stream of type s^* into a parallel pair of streams $s^* || s^*$. It does this by sending the first element to the left branch, the second to the right, the third to the left, and so on. In λ^{ST} , we write this by maintaining a Boolean accumulator, and negating after each item. If the Boolean is true, we send the element left, if it's false, we send it right.

```

roundRobin (xs : s*) : s* || s* =
fix with (b = false).
  case xs of
  | nil => (nil, nil)
  | y :: ys => let (zs,ws) = rec(ys)[not b] in
                if b then (y :: zs,ws) else (zs, y :: ws)

```

The core term for this program can be found at `Examples.hs:roundRobin`

Hash-Based Partitioning. A hash-based partitioner routes stream elements based on a hash of their data, mod the number of routes k . We fix $k = 2$ for simplicity, and abstract this hashing function as a transformer `hm2` which accepts a stream of type s and produces a stream of type $s + s$. Conceptually, we assume that this program computes a hash mod 2, and returns the input in the left or right position accordingly. Then we can write the hash-based partitioner like this:

```

sumRouter (xs : s*) : s* || s* =
  fix. case xs of
  | nil => (nil, nil)
  | y :: ys => let (zs,ws) = rec(ys) in
                case hm2(y) of
                | inl(z) => (z :: zs, ws)
                | inr(w) => (zs, w :: ws)

```

The core term for this program can be found at `Examples.hs:sumRouter`

C.7 Windowing and Punctuation

Many kinds of windows have been considered in the literature. The most common windows are event-based – windows defined by the number of elements they’ll contain – and time-based – windows which contain all the events from a fixed length of time. Windows can also be tumbling – the next window starts after the previous ends – or sliding – every event could begin a new window.

In λ^{ST} , windowed operators are just maps over a stream whose elements are windows. Given a per-window stream transformer f which takes windows s^* to a result type t , and a “windowing strategy” win which takes a stream r^* and turns it into a stream of windows $(s^*)^*$, we can write a windowed operation of type $r^* \rightarrow t^*$ as follows: $\text{xs} : r^* \vdash \text{map}(f)(\text{win}(\text{xs})) : t^*$.

For example, if we wanted to compute a size-3 sliding sum of a stream of `Ints`, we would use a windower win which takes Int^* to $(\text{Int}^*)^*$ where the inner streams are the windows, and f from Int^* to `Int` is the sum operation.

Every per-window function commonly used in stream processing practice operates on entire windows at once, which is accomplished in λ^{ST} by `wait-ing` on the whole window, and then aggregating it with an embedded historical program. For this reason, we focus primarily on the window construction aspect.

Fixed-Size Tumbling Windows. The k -size tumbling windower creates windows of size k , where each new window starts immediately after the last window ended. For instance when $k = 2$, a stream `1, 2, 4, 7, 3, 8, ...` turns into a stream $\langle 1, 2 \rangle, \langle 4, 7 \rangle, \langle 3, 8 \rangle, \dots$. The code for a fixed-size tumbling window is exactly the functional code for computing k -strides of a list using a fold. We maintain a counter of the current window size, and compute a `concat-pair` of the first (partial) window, and the windows of the remaining stream. The actual size- k tumbling windower is computed by post-composing this transformer with a `cons` operation.

```

tumblingWindow (k : int) (xs : s*) : s* . s** =

```

```

fix with (n = 0).
  case xs of
  | nil => (nil;nil)
  | y :: ys => if n + 1 >= k then let (cur;rest) = rec(ys)[0] in (y :: nil; cur :: rest)
                else let (cur;rest) = rec(ys)[n+1] in (y :: cur; rest)

```

The core term for this program can be found at `Examples.hs:tumblingWindow`

k -size window transformers can actually have the even stronger output type $(s^k)^*$, where s^k is the k -fold concatenation of s . If the window function being used requires that the windows all have exactly size k (like taking pairwise differences for $k = 2$), this type can be used instead. The following program implements size-2 windows with this stronger type by casing two-deep into the stream at a time, and pairing up elements into concatenation pairs.

```

wind2 (xs : s*) : (s · s)*
  fix. case xs of
  | nil => nil
  | y :: ys => case ys of
    | nil => nil
    | z :: zs => (y;z) :: rec(zs)

```

The core term for this program can be found at `Examples.hs:sumRouter`

Fixed-Size Sliding Windows. A k -sized sliding windower produces a new window for each new element, including both the new element and the $k - 1$ previous ones. The code for this windower keeps the current window under construction in memory. When each new stream element arrives, we emit the current window. For the first k elements, we only add to the window. After k , we start evicting from the window.

```

slidingWindower (xs : s*) : s** =
  fix with (acc = []).
  case xs of
  | nil => acc :: nil
  | y::ys => wait y then
    let next = (if |acc| < k then y :: acc else y :: (init acc)) in
    (next) :: rec(ys)[next]
  end

```

The core term for this program can be found at `Examples.hs:slidingWindower`

Punctuation-Based Windows. Time-based windows are commonly implemented by way of *punctuation*: unit elements inserted into a stream to authoritatively mark that a period of time has ended. This is required because in the presence of network delays, it's impossible to know if a time period is over (and so a window can be emitted) or if there are more elements in the period to arrive. A punctuated stream has type $(1 + s)^*$, where the punctuation events mark the end of each time period.

The following code computes a windowed stream $(s^*)^*$ from a punctuated stream $(1 + s)^*$ by emitting windows which are the (potentially empty) runs of ss between punctuation marks. Like fixed-size tumbling windows, we generalize the return type to compute with a fold, and then the actual operator is defined by post-composing with a `cons`.

```

puncWindow (xs : (1+s)*) : s* · s** =
  fix.
  case xs of
  | nil => (nil ; nil)

```

```

| y :: ys => let (cur;rest) = rec(ys) in
    case y of
    | inl() => (nil ; cur :: rest)
    | inr(z) => (z :: cur ; rest)

```

The core term for this program can be found at `Examples.hs:puncWindow`

Merging Streams and Synchronizing Punctuation. Parallel streams of star type can be *synchronized*, pairing off one element from one stream with one element of another. Given a stream of type $s^* || t^*$, we can produce a stream of type $(s || t)^*$. This type's similarity to the standard functional program `zip` is more than just surface level: the program below has essentially the same code.

```

sync (xs : s*, ys : t*) : (s || t)* =
  fix.
  case xs of
  | nil => nil
  | x' :: xs' =>
    case ys of
    | nil => nil
    | y' :: ys' =>
      wait x',y' then
        <<x',y'>> :: rec(xs,ys)
  end

```

The core term for this program can be found at `Examples.hs:sync`

Semantically, this program waits until a full element from each of the parallel input streams has arrived, sends them both out, and then continues with zipping the two tails. This is necessarily blocking: the output type guarantees that exactly one s and t will be produced before the next pair begins, and so we must wait for both to arrive before sending the other out. The upshot is that because this program is well typed in λ^{ST} , it is necessarily deterministic. This gives us the deterministic merge operation that was needed to prevent the bug when averaging data from a pair of sensors in Section 2.

Moreover, for parallel streams of windows, synchronization enables databases-style *streaming joins*. Given parallel streams $(s^*)^*$ and $(t^*)^*$, we can synchronize to get $(s^* || t^*)^*$, and then apply a join operation to each parallel pair of windows.