

# An Introduction to the StreamQRE Language

Rajeev ALUR and Konstantinos MAMOURAS  
*Department of Computer and Information Science*  
*University of Pennsylvania*  
*Philadelphia, PA 19104, USA*

**Abstract.** Real-time decision making in emerging IoT applications typically relies on computing quantitative summaries of large data streams in an efficient and incremental manner. We give here an introduction to the StreamQRE language, which has recently been proposed for the purpose of simplifying the task of programming the desired logic in such stream processing applications. StreamQRE provides natural and high-level constructs for processing streaming data, and it offers a novel integration of linguistic constructs from two distinct programming paradigms: streaming extensions of relational query languages and quantitative extensions of regular expressions. The former allows the programmer to employ relational constructs to partition the input data by keys and to integrate data streams from different sources, while the latter can be used to exploit the logical hierarchy in the input stream for modular specifications.

**Keywords.** data stream processing, Quantitative Regular Expressions

## Introduction

The last few years have witnessed an explosion of IoT systems in applications such as smart buildings, wearable devices, and healthcare [1]. A key component of an effective IoT system is the ability to make decisions in real-time in response to data it receives. For instance, a gateway router in a smart home should detect and respond in a timely manner to security threats based on monitored network traffic, and a healthcare system should issue alerts in real-time based on measurements collected from all the devices for all the monitored patients. While the exact logic for making decisions in different applications requires domain-specific insights, it typically relies on computing *quantitative* summaries of large data streams in an efficient and incremental manner. Programming the desired logic as a deployable implementation is challenging due to the enormous volume of data and hard constraints on available memory and response time.

The recently proposed language StreamQRE [2] (pronounced StreamQuery) is meant to assist IoT programmers: it makes the task of specifying the desired decision-making logic simpler by providing natural and high-level declarative constructs for processing streaming data, and the proposed compiler and runtime system facilitates deployment with guarantees on memory footprint and per-item processing time. The StreamQRE language extends *quantitative regular expressions*—an extension of clas-

sical regular expressions for associating numerical values with strings [3], with constructs typical in extensions of relational query languages for handling streaming data [4,5,6,7,8,9,10,11]. The novel integration of linguistic constructs allows the programmer to impart to the input data stream a logical hierarchical structure (for instance, view patient data as a sequence of episodes and view network traffic as a sequence of Voice-over-IP sessions) and also employ relational constructs to partition the input data by keys (e.g., patient identifiers and IP addresses).

The basic object in the language is a *streaming query*, which is modeled as a partial function from sequences of input data items to an output value (which can be a relation). We present the syntax and semantics of the StreamQRE language with type-theoretic foundations. In particular, each streaming query has an associated *rate* that captures its domain, that is, as it reads the input data stream, the prefixes that trigger the production of the output. In the StreamQRE calculus, the rates are required to be *regular*, captured by symbolic regular expressions, and the theoretical foundations of symbolic automata [12] lead to decision procedures for constructing well-typed expressions. Regular rates also generalize the concept of *punctuations* in streaming database literature [6].

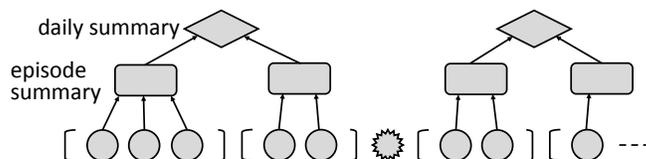
The StreamQRE language has a small set of core combinators with clear semantics. An atomic query processes individual items. The constructs `split` and `iter` are quantitative analogs of concatenation and Kleene-iteration, and integrate hierarchical pattern matching with familiar sequential iteration over a list of values. The global choice operator `or` allows selection between two expressions with disjoint rates. The combination operator `combine` allows combining output values produced by multiple expressions with equivalent rates processing input data stream in parallel. The key-based partitioning operator `map-collect` is a generalization of the widely used map-reduce construct that partitions the input data stream into a set of sub-streams, one per key, and returns a relation. Finally, the streaming composition operator  $\gg$  streams the sequence of outputs produced by one expression as an input stream to another, allowing construction of pipelines of operators.

The core StreamQRE constructs can be used to define a number of derived patterns that are useful in practice, such as *tumbling* and *sliding windows* [6], selection, and filtering. The language has been implemented as a Java library that supports the basic and derived constructs [13]. We show how to program in StreamQRE using an illustrative example regarding monitoring patient measurements, the recent Yahoo Streaming Benchmark for advertisement-related events [14], and the NEXMark Benchmark for auction bids [15]. These examples illustrate how hierarchically nested iterators and global case analysis facilitate modular stateful sequential programming, and key-based partitioning and relational operators facilitate traditional relational programming. The two styles offer alternatives for expressing the same query in some cases, while some queries are best expressed by intermingling the two views.

**Organization.** The remaining paper is organized as follows. Section 1 introduces the syntax and semantics of the StreamQRE language, and explains how each construct is used with an illustrative example regarding monitoring patient measurements. Section 2 shows how to use StreamQRE to program some common stream transformation. Section 3 presents some example queries for the Yahoo Streaming Benchmark, and Section 4 gives queries for the NEXMark Benchmark.

## 1. The StreamQRE Language

As a motivating example, suppose that a patient is being monitored for episodes of a physiological condition such as epilepsy [16], and the data stream consists of four types of events: (1) An event  $B$  marking the beginning of an episode, (2) a time-stamped measurement  $M(ts, val)$  by a sensor, (3) an event  $E$  marking the end of an episode, (4) and an event  $D$  marking the end of a day. Given such an input data stream, suppose we want to specify a policy  $f$  that outputs every day, the maximum over all episodes during that day, of the average of all measurements during an episode. A suitable abstraction is to impart a hierarchical structure to the stream:



The data stream is a sequence of days (illustrated as diamonds), where each day is a sequence of episodes (illustrated as rectangles), and each episode is a nonempty sequence of corresponding measurements (shown as circles) between a begin  $B$  marker (shown as an opening bracket) and an end  $E$  marker (shown as a closing bracket). The end-of-day marker is shown as a star. The regular expression  $((B \cdot M^+ \cdot E)^* \cdot D)^*$  over the event types  $B$ ,  $M$ ,  $E$  and  $D$  specifies naturally the desired hierarchical structure. For simplicity, we assume that episodes do not span day markers.

The policy  $f$  thus describes a hierarchical computation that follows the structure of this decomposition of the stream: the summary of each episode (pattern  $B \cdot M^+ \cdot E$ ) is an aggregation of the measurements (pattern  $M$ ) it contains, and similarly the summary of each day (pattern  $(B \cdot M^+ \cdot E)^* \cdot D$ ) is an aggregation of the summaries of the episodes it contains. In order for the policy to be fully specified, the hierarchical decomposition (parse tree) of the stream has to be unique. Otherwise, the summary would not be uniquely determined and the policy would be ambiguous. To guarantee uniqueness of parsing at compile time, each policy  $f$  describes a *symbolic unambiguous regular expression*, called its *rate*, which allows for at most one way of decomposing the input stream. The qualifier *symbolic* means that the alphabets (data types) can be of unbounded size, and that unary predicates are used to specify classes of letters (data items) [12]. The use of regular rates implies decidability of unambiguity. Even better, there are efficiently checkable typing rules that guarantee unambiguity for all policies [17,18,19].

To define quantitative queries, we first choose a typed signature which describes the basic data types and operations for manipulating them. We fix a collection of *basic types*, and we write  $A, B, \dots$  to range over them. This collection contains the type  $\text{Bool}$  of boolean values, and the unit type  $\text{Ut}$  whose unique inhabitant is denoted by  $\text{def}$ . It is also closed under the cartesian product operation  $\times$  for forming pairs of values. Typical examples of basic types are the natural numbers  $\text{Nat}$ , the integers  $\text{Int}$ , and the real numbers  $\mathbb{R}$ .

We also fix a collection of *basic operations* on the basic types, for example the  $k$ -ary operation  $op : A_1 \times \dots \times A_k \rightarrow B$ . The identity function on  $D$  is written as  $id_D : D \rightarrow D$ , and the operations  $\pi_1 : A \times B \rightarrow A$  and  $\pi_2 : A \times B \rightarrow B$  are the left and right projection respectively. We assume that the collection of operations contains all identities and pro-

jections, and is closed under pairing and function composition. To describe derived operations we use a variant of lambda notation that is similar to Java's lambda expressions [20]. That is, we write  $(A\ x) \rightarrow t(x)$  to mean  $\lambda x:A.t(x)$  and  $(A\ x,\ B\ y,\ C\ z) \rightarrow t(x,y,z)$  to mean  $\lambda x:A,y:B,z:C.t(x,y,z)$ . For example, the identity function on  $D$  is  $(D\ x) \rightarrow x$ , the left projection on  $A \times B$  is  $(A\ x,\ B\ y) \rightarrow x$ , the right projection on  $A \times B$  is  $(A\ x,\ B\ y) \rightarrow y$ , and  $(D\ x) \rightarrow \text{def}$  is the unique function from  $D$  to  $\text{Ut}$ . We will typically use lambda expressions in the context of queries from which the types of the input variables can be inferred, so we will omit them as in  $(x,y) \rightarrow x$ .

For every basic type  $D$ , assume that we have fixed a collection of *atomic predicates*, so that the satisfiability of their Boolean combinations (built up using the Boolean operations: and, or, not) is decidable. We write  $\varphi : D \rightarrow \text{Bool}$  to indicate that  $\varphi$  is a predicate on  $D$ , and we denote by  $\text{true}_D : D \rightarrow \text{Bool}$  the predicate that is always true. The predicate  $((\text{Int}\ x) \rightarrow x > 0) : \text{Int} \rightarrow \text{Bool}$  is true of the strictly positive integers.

**Example 1.** For the example patient-monitoring stream described previously, suppose that we now allow the stream to contain information for several patients. The data type  $D_P$  for this multiple-patient monitoring stream is the tagged (disjoint) union:

$$D_P = \{D\} \cup \{B(p), E(p) \mid p \in \text{PID}\} \cup \{M(p,t,v) \mid p \in \text{PID}, t \in T, \text{ and } v \in V\},$$

where  $\text{PID}$  is the set of patient identifiers,  $T$  is the set of timestamps, and  $V$  is the set of scalars for the measurements. The projection functions  $\text{typ} : D_P \rightarrow \{D,B,E,M\}$ ,  $\text{pId} : D_P \rightarrow \text{PID}$ ,  $\text{ts} : D_P \rightarrow T$  and  $\text{val} : D_P \rightarrow V$  get the type, patient identifier, timestamp, and value of a data item respectively (when undefined, the functions simply return some default value). For a data item  $x \in D_P$ , we write  $x.\text{typ}$ ,  $x.\text{pId}$ ,  $x.\text{ts}$  and  $x.\text{val}$  to denote the application of these functions.

**Symbolic regular expressions.** For a type  $D$ , we define the set of *symbolic regular expressions over  $D$*  [21], denoted  $\text{RE}(D)$ , with the following grammar:

$r ::= \varphi$		[predicate on $D$ ]
$\varepsilon$		[empty sequence]
$r \sqcup r$		[nondeterministic choice]
$r \cdot r$		[concatenation]
$r^*$	.	[iteration]

The concatenation symbol  $\cdot$  is sometimes omitted, that is, we write  $rs$  instead of  $r \cdot s$ . The expression  $r^+$  (iteration at least once) abbreviates  $r \cdot r^*$ . We write  $r : \text{RE}(D)$  to indicate the  $r$  is a regular expression over  $D$ . Every expression  $r : \text{RE}(D)$  is interpreted as a set  $\llbracket r \rrbracket \subseteq D^*$  of finite sequences over  $D$ .

$$\begin{aligned} \llbracket \varphi \rrbracket &\triangleq \{d \in D \mid \varphi(d) = \text{true}\} & \llbracket r \sqcup s \rrbracket &\triangleq \llbracket r \rrbracket \cup \llbracket s \rrbracket & \llbracket r^* \rrbracket &\triangleq \bigcup_{n \geq 0} \llbracket r \rrbracket^n \\ \llbracket \varepsilon \rrbracket &\triangleq \{\varepsilon\} & \llbracket r \cdot s \rrbracket &\triangleq \llbracket r \rrbracket \cdot \llbracket s \rrbracket \end{aligned}$$

For subsets  $X, Y \subseteq D^*$ , we define  $X \cdot Y = \{uv \mid u \in X \text{ and } v \in Y\}$  and  $X^n$  is given by induction on  $n$  as follows:  $X^0 = \{\varepsilon\}$  and  $X^{n+1} = X^n \cdot X$ . Two expressions are said to be *equivalent* if they denote the same language.

	data item	stream seen so far	current output (if any)
		$\varepsilon$	$f(\varepsilon)$
	$d_1$	$d_1$	$f(d_1)$
	$d_2$	$d_1 d_2$	$f(d_1 d_2)$
	$d_3$	$d_1 d_2 d_3$	$f(d_1 d_2 d_3)$
	$d_4$	$d_1 d_2 d_3 d_4$	$f(d_1 d_2 d_3 d_4)$
	$d_5$	$d_1 d_2 d_3 d_4 d_5$	$f(d_1 d_2 d_3 d_4 d_5)$
	...		

**Figure 1.** A streaming transformation  $f$  specifies the output for every prefix of the stream.

**Example 2.** The symbolic regular expression  $((\text{Nat } x) \rightarrow \text{true})^* \cdot ((\text{Nat } x) \rightarrow x > 0)$  over the type  $\text{Nat}$  of natural numbers denotes all sequences that end with a strictly positive number.

**Unambiguity.** The notion of unambiguity for regular expressions [17] is a way of formalizing the requirement of uniqueness of parsing. The languages  $L_1, L_2$  are said to be *unambiguously concatenable* if for every word  $w \in L_1 \cdot L_2$  there are unique  $w_1 \in L_1$  and  $w_2 \in L_2$  with  $w = w_1 w_2$ . The language  $L$  is said to be *unambiguously iterable* if for every word  $w \in L^*$  there is a unique integer  $n \geq 0$  and unique  $w_i \in L$  with  $w = w_1 \cdots w_n$ . The definitions of unambiguous concatenability and unambiguous iterability extend to regular expressions in the obvious way. Now, a regular expression is said to be *unambiguous* if it satisfies the following:

1. For every subexpression  $e_1 \sqcup e_2$ ,  $e_1$  and  $e_2$  are *disjoint*.
2. For every subexpression  $e_1 \cdot e_2$ ,  $e_1$  and  $e_2$  are *unambiguously concatenable*.
3. For every subexpression  $e^*$ ,  $e$  is *unambiguously iterable*.

Checking whether a regular expression is unambiguous can be done in polynomial time. For the case of symbolic regular expressions this results still holds, under the assumption that satisfiability of the predicates can be decided in unit time [3].

**Example 3.** Consider the finite alphabet  $\Sigma = \{a, b\}$ . The regular expression  $r = (a + b)^* b (a + b)^*$  denotes the set of sequences with at least one occurrence of  $b$ . It is ambiguous, because the subexpressions  $(a + b)^* b$  and  $(a + b)^*$  are not unambiguously concatenable: the word  $w = ababa$  matches  $r$ , but there are two different splits  $w = ab \cdot aba$  and  $w = abab \cdot a$  that witness the ambiguity of parsing. The regular expressions  $a^* b (a + b)^*$  and  $(a + b)^* b a^*$  are both equivalent to  $r$ , and they are unambiguous.

**Streaming transformations.** The basic object in the StreamQRE language is the *query*, which describes the transformation of an input stream into an output stream. At any given moment in time, only a finite number of data items have arrived, therefore a stream transformation can be modeled as a function from  $D^*$ , where  $D$  is the type of input data items, to  $C$ , where  $C$  is the type of the outputs. In other words, the transformation describes how to aggregate the entire stream seen so far into an output value. As the input stream gets extended with more and more items, the emitted outputs form a stream of elements of  $C$ . We want to allow for the possibility of not having an output with every new element arrival, therefore a streaming transformation is modeled as a partial function  $D^* \dashrightarrow C$ . See Figure 1 for an illustration.

**Example 4.** For example, suppose we want to describe a *filtering* transformation on a stream of integers, where only the nonnegative numbers are retained and the negative numbers are filtered out. The function  $f : \text{Int}^* \rightarrow \text{Int}$  that describes this transformation is defined on the nonempty sequences  $v_1 v_2 \dots v_n$  with  $v_n \geq 0$ , and the value is the last number of the sequence, i.e. the current item.

	data item	stream seen so far	current output (if any)
		$\epsilon$	
	3	3	3
	-2	3 -2	
	4	3 -2 4	4
	8	3 -2 4 8	8
	-1	3 -2 4 8 -1	
	5	3 -2 4 8 -1 5	5
	...		

time ↓

The set of stream prefixes for which this function is defined is denoted by the symbolic regular expression  $((\text{Int } x) \rightarrow \text{true})^* \cdot ((\text{Int } x) \rightarrow x > 0)$ . After having consumed the input stream 3 -2 4 8 -1 5, the overall output stream is 3 4 8 5.

The rate of a transformation describes which stream prefixes trigger the production of output. In some other formalisms, such as transducers [22,23] and synchronous languages [24], output is typically produced at every new data item arrival. In StreamQRE, on the other hand, output does not have to be produced with every new item. For example, after processing  $n$  data items, the output stream generally consists of  $k \leq n$  items.

**Streaming Queries.** We now introduce formally the language of *Streaming Quantitative Regular Expressions (QREs)* for representing stream transformations. For brevity, we also call these expressions *queries*. A query represents a streaming transformation whose domain is a regular set over the input data type. The *rate* of a query  $f$ , written  $R(f)$ , is a symbolic regular expression that denotes the domain of the transformation that  $f$  represents. The definition of the query language has to be given simultaneously with the definition of rates (by mutual induction), since the query constructs have typing restrictions that involve the rates. We annotate a query  $f$  with a type  $\text{QRE}\langle D, C \rangle$  to denote that the input stream has elements of type  $D$  and the outputs are of type  $C$ . Figure 2 shows the full formal syntax of streaming queries. The decidability of type checking follows from results in [17,18,19] and it is also discussed in [3].

**Example 5 (Rate of a query).** In the patient monitoring example described in the beginning of this section, the statistical summary of a patient’s measurements should be output at the end of each day, and thus, depends only on the types of events in a regular manner. The rate in this case is the regular expression  $((B \cdot M^+ \cdot E)^* \cdot D)^*$ . The tag  $B$  in this expression is abbreviation for the predicate  $(x \rightarrow x.\text{typ} = B)$ . We also write  $\neg B$  to stand for the predicate  $(x \rightarrow x.\text{typ} \neq B)$ . Similar abbreviations are considered for the tags  $M$ ,  $E$  and  $D$ . We will be using these abbreviations freely from now on, since their meaning is obvious from the context.

$\frac{\text{satisfiable } \varphi : D \rightarrow \text{Bool} \quad op : D \rightarrow C}{\text{atom}(\varphi, op) : \text{QRE}\langle D, C \rangle} \text{ (atomic)}$ $R(\text{atom}(\varphi, op)) = \varphi$	$\frac{c \in C}{\text{eps}(c) : \text{QRE}\langle D, C \rangle} \text{ (empty)}$ $R(\text{eps}(c)) = \varepsilon$
$\frac{\begin{array}{l} f : \text{QRE}\langle D, A \rangle \quad g : \text{QRE}\langle D, B \rangle \quad op : A \times B \rightarrow C \\ R(f) \text{ and } R(g) \text{ are unambiguously concatenable} \end{array}}{\text{split}(f, g, op) : \text{QRE}\langle D, C \rangle} \text{ (concatenation)}$ $R(\text{split}(f, g, op)) = R(f) \cdot R(g)$	
$\frac{\begin{array}{l} \text{init} : \text{QRE}\langle D, B \rangle \quad \text{body} : \text{QRE}\langle D, A \rangle \quad op : B \times A \rightarrow B \\ R(\text{body}) \text{ is unambiguously iterable} \\ R(\text{init}) \text{ and } R(\text{body})^* \text{ are unambiguously concatenable} \end{array}}{\text{iter}(\text{init}, \text{body}, op) : \text{QRE}\langle D, B \rangle} \text{ (iteration)}$ $R(\text{iter}(f, g, op)) = R(f) \cdot R(g)^*$	
$\frac{\begin{array}{l} f : \text{QRE}\langle D, C \rangle \quad g : \text{QRE}\langle D, C \rangle \quad R(f) \text{ and } R(g) \text{ are disjoint} \end{array}}{\text{or}(f, g) : \text{QRE}\langle D, C \rangle} \text{ (choice)}$ $R(\text{or}(f, g)) = R(f) \sqcup R(g)$	
$\frac{f : \text{QRE}\langle D, A \rangle \quad op : A \rightarrow B}{\text{apply}(f, op) : \text{QRE}\langle D, B \rangle} \text{ (application)}$ $R(\text{apply}(f)) = R(f)$	
$\frac{\begin{array}{l} f : \text{QRE}\langle D, A \rangle \quad g : \text{QRE}\langle D, B \rangle \quad op : A \times B \rightarrow C \quad R(f), R(g) \text{ equivalent} \end{array}}{\text{combine}(f, g, op) : \text{QRE}\langle D, C \rangle} \text{ (combination)}$ $R(\text{combine}(f, g, op)) = R(f)$	
$\frac{f : \text{QRE}\langle D, C \rangle \quad g : \text{QRE}\langle C, E \rangle}{f \gg g : \text{QRE}\langle D, E \rangle} \text{ (streaming composition)}$	
$\frac{\begin{array}{l} \varphi_S : D \rightarrow \text{Bool} \quad m : D \rightarrow K \quad f : \text{QRE}\langle D, C \rangle \quad r : \text{RE}\langle D \rangle \quad R(f) \subseteq r \setminus \varphi_S^* \\ r = s[(-\varphi_S)^* \varphi_S / \varphi_S], \text{ where the only predicate } s : \text{RE}\langle D \rangle \text{ can contain is } \varphi_S \end{array}}{\text{map-collect}(\varphi_S, m, f, r) : \text{QRE}\langle D, \text{Map}\langle K, C \rangle \rangle} \text{ (key-based partitioning)}$ $R(\text{map-collect}(\varphi, m, f, r)) = r$	

**Figure 2.** The syntax of Streaming Quantitative Regular Expressions.

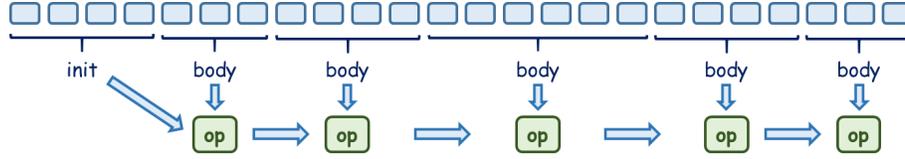
**Atomic queries.** The basic building blocks of queries are expressions that describe the processing of a single data item. Suppose  $\varphi : D \rightarrow \text{Bool}$  is a predicate over the data item type  $D$  and  $op : D \rightarrow C$  is an operation from  $D$  to the output type  $C$ . Then, the *atomic query*  $\text{atom}(\varphi, op) : \text{QRE}\langle D, C \rangle$ , with rate  $\varphi$ , is defined on single-item streams that satisfy the predicate  $\varphi$ . The output is the value of  $op$  on the input element.

*Notation:* It is very common for  $op$  to be the identity function, and  $\varphi$  to be the always-true predicate. So, we abbreviate the query  $\text{atom}(\varphi, (D\ x) \rightarrow x)$  by  $\text{atom}(\varphi)$ , and the query  $\text{atom}((D\ x) \rightarrow \text{true})$  by  $\text{atom}()$ .

**Example 6.** For the stream of monitored patients, the query that matches a single measurement item and returns its value is  $f = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow x.\text{val})$ . The type of  $f$  is  $\text{QRE}\langle D_p, V \rangle$  and its rate is  $M$ .

**Empty sequence.** The query  $\text{eps}(c) : \text{QRE}\langle D, C \rangle$ , where  $c$  is a value of type  $C$ , is only defined on the empty sequence  $\varepsilon$  and it returns the output  $c$ .

**Iteration.** Suppose that  $\text{init} : \text{QRE}\langle D, B \rangle$  describes the computation for initializing an aggregate value of type  $B$ , and  $\text{body} : \text{QRE}\langle D, A \rangle$  describes a computation that we want to iterate over consecutive subsequences of the input stream, in order to aggregate the values (of type  $A$ ) sequentially using an aggregator  $op : B \times A \rightarrow B$ .



More specifically, we split the input stream  $w$  into subsequences  $w = u\ w_1\ w_2 \dots w_n$ , where  $u$  matches  $\text{init}$  and each  $w_i$  matches  $\text{body}$ . We apply  $\text{init}$  to  $u$  and  $\text{body}$  to each of the  $w_i$ , thus producing the output values  $b_0$  and  $a_1\ a_2 \dots a_n$  with  $b_0 = \text{init}(u)$  and  $a_i = \text{body}(w_i)$ . Finally, we combine these results using the list iterator *left fold* with start value  $b_0$  and aggregation operation  $op : B \times A \rightarrow B$  by folding the list of values  $a_1\ a_2 \dots a_n$ . This can be formalized with the combinator  $\text{fold} : B \times (B \times A \rightarrow B) \times A^* \rightarrow B$ , which takes an initial value  $b \in B$  and a stepping map  $op : B \times A \rightarrow B$ , and iterates through a sequence of values of type  $A$ :

$$\text{fold}(b, op, \varepsilon) = b \qquad \text{fold}(b, op, \gamma a) = op(\text{fold}(b, op, \gamma), a)$$

for all sequences  $\gamma \in A^*$  and all values  $a \in A$ . E.g.,  $\text{fold}(b, op, a_1 a_2) = op(op(b, a_1), a_2)$ .

The query  $g = \text{iter}(\text{init}, \text{body}, op) : \text{QRE}\langle D, B \rangle$  describes the computation of the previous paragraph. In order for  $g$  to be well-defined as a function, every input stream  $w$  that matches  $g$  must be uniquely decomposable into  $w = u w_1 w_2 \dots w_n$  with  $u$  matching  $\text{init}$  and each  $w_i$  matching  $\text{body}$ . This requirement can be expressed equivalently as follows: the rate  $R(\text{body})$  is unambiguously iterable, and the rates  $R(\text{init}), R(\text{body})^*$  are unambiguously concatenable.

These sequential iterators can be nested imparting a hierarchical structure to the input data stream facilitating modular programming. In the single-patient monitoring stream, for example, we can associate an iterator with the episode nodes to summarize the sequence of measurements in an episode, and another iterator with the day nodes to summarize the sequence of episodes during a day.

**Example 7.** For the stream of monitored patients, the query  $g$  below matches a sequence of measurements and returns the sum of their values.

$f : \text{QRE}\langle D_P, V \rangle = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow x.\text{val}) \quad // \text{ rate } M$   
 $g : \text{QRE}\langle D_P, V \rangle = \text{iter}(\text{eps}(0), f, (x, y) \rightarrow x + y) \quad // \text{ rate } M^*$

**Combination and application.** Assume the queries  $f$  and  $g$  describe stream transformations with outputs of type  $A$  and  $B$  respectively that process the same set of input sequences, and  $op$  is an operation of type  $A \times B \rightarrow C$ . The query  $\text{combine}(f, g, op)$  describes the computation where the input is processed according to both  $f$  and  $g$  in parallel and their results are combined using  $op$ . Of course, this computation is meaningful only when both  $f$  and  $g$  are defined on the input sequence. So, we demand w.l.o.g. that the rates of  $f$  and  $g$  are equivalent.

This binary combination construct generalizes to an arbitrary number of queries. For example, we write  $\text{combine}(f, g, h, (x, y, z) \rightarrow op(x, y, z))$  for the ternary variant. In particular, we write  $\text{apply}(f, op)$  for the case of one argument.

**Example 8 (Average).** For the stream of monitored patients, the query  $h$  below matches a nonempty sequence of measurements and returns the average of their values.

$f_1 : \text{QRE}\langle D_P, V \rangle = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow x.\text{val}) \quad // \text{ rate } M$   
 $g_1 : \text{QRE}\langle D_P, V \rangle = \text{iter}(f_1, f_1, (x, y) \rightarrow x + y) \quad // \text{ rate } M^+$   
 $f_2 : \text{QRE}\langle D_P, V \rangle = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow 1) \quad // \text{ rate } M$   
 $g_2 : \text{QRE}\langle D_P, V \rangle = \text{iter}(f_2, f_2, (x, y) \rightarrow x + y) \quad // \text{ rate } M^+$   
 $h : \text{QRE}\langle D_P, V \rangle = \text{combine}(g_1, g_2, (x, y) \rightarrow x/y) \quad // \text{ rate } M^+$

The query  $g_1$  computes the sum of a nonempty sequence of measurements, and the query  $g_2$  computes the length of a nonempty sequence of measurements. An alternative implementation uses a single iteration construct and an accumulator that is the pair of the running sum and the running count.

$f' : \text{QRE}\langle D_P, V \times V \rangle = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow (x.\text{val}, 1)) \quad // \text{ rate } M$   
 $g' : \text{QRE}\langle D_P, V \times V \rangle = \text{iter}(f', f', (x, y) \rightarrow x + y) \quad // \text{ rate } M^+$   
 $h' : \text{QRE}\langle D_P, V \rangle = \text{apply}(g', x \rightarrow \pi_1(x)/\pi_2(x)) \quad // \text{ rate } M^+$

The  $+$  operation in  $g'$  is componentwise addition of pairs of values. The query  $h'$  computes the average by dividing the running sum by the running count.

**Example 9 (Standard Deviation).** For a sequence  $x_1, x_2, \dots, x_n$  of numbers, the *mean* is the simple average  $\mu = (\sum_i x_i)/n$ , and the *standard deviation* is  $\sigma = (\sum_i (x_i - \mu)^2)/n$ . Equivalently, we can calculate the quantity  $\sigma \cdot n$  as follows:

$$\sigma \cdot n = \sum_i (x_i^2 + \mu^2 - 2\mu x_i) = (\sum_i x_i^2) + n\mu^2 - 2\mu(\sum_i x_i) = (\sum_i x_i^2) - (\sum_i x_i)^2/n.$$

So, both the mean and the standard deviation can be calculated from the quantities  $\sum_i x_i$  and  $\sum_i x_i^2$  and the count  $n$ . This suggests the following query for the streaming computation of the standard deviation for a sequence of patient measurements:

```

f1 : QRE⟨DP, V⟩ = atom(x → x.typ = M, x → 1)           // rate M
g1 : QRE⟨DP, V⟩ = iter(f1, f1, (x, y) → x + y)         // rate M+
f2 : QRE⟨DP, V⟩ = atom(x → x.typ = M, x → x.val)       // rate M
g2 : QRE⟨DP, V⟩ = iter(f2, f2, (x, y) → x + y)         // rate M+
f3 : QRE⟨DP, V⟩ = atom(x → x.typ = M, x → x.val · x.val) // rate M
g3 : QRE⟨DP, V⟩ = iter(f3, f3, (x, y) → x + y)         // rate M+
h : QRE⟨DP, V⟩ = combine(g1, g2, g3, (x, y, z) → z - y2/x) // rate M+

```

The query  $g_1$  computes the running count, the query  $g_2$  computes the running sum  $\sum_i x_i$ , and the query  $g_3$  computes the running sum of squares  $\sum_i x_i^2$ .

**Example 10 (Integration).** For the single-patient monitoring stream, assume that the diagnosis depends on the average value of the *piecewise-linear interpolant* of the sampled measurements. Computing this quantity corresponds to integrating the interpolant over the interval of the measurements. That is, the quantitative summary of a given sequence  $(t_1, v_1), (t_2, v_2), \dots, (t_n, v_n)$  of timestamped values with  $t_1 < t_2 < \dots < t_n$  is

$$A_n = \frac{1}{(t_n - t_1)} \cdot \sum_{i=1}^{n-1} \frac{(v_i + v_{i+1})(t_{i+1} - t_i)}{2}.$$

To compute the quantity  $S_n = \sum_{i=1}^{n-1} (v_i + v_{i+1})(t_{i+1} - t_i)$  incrementally we must maintain the vector  $X_n = (t_n, v_n, S_n)$ , where  $t_n$  is the last timestamp,  $v_n$  is the last value, and  $S_n$  is the running sum. We then put  $X_1 = (t_1, v_1, 0)$  and

$$X_{n+1} = (t_{n+1}, v_{n+1}, S_n + (v_n + v_{n+1})(t_{n+1} - t_n)).$$

From  $X_n = (t_n, v_n, S_n)$  and the first timestamp  $t_1$  we can then compute  $A_n = S_n / 2(t_n - t_1)$ .

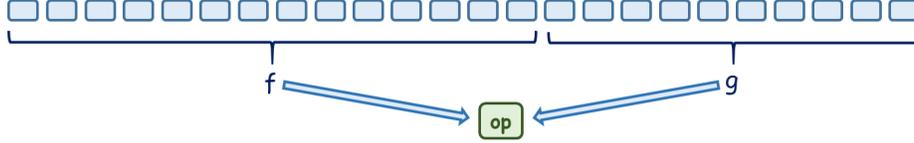
```

f : QRE⟨DP, DP⟩ = atom(x → x.typ = M)           // rate M
g1 : QRE⟨DP, T⟩ = iter(apply(f, x → x.ts), f, (x, y) → x) // rate M+
op = (T × V × V x, DP y) →
      (y.ts, y.val, π3(x) + (π2(x) + y.val)(y.ts - π1(x)))
g2 : QRE⟨DP, T × V × V⟩ = iter(apply(f, x → (x.ts, x.val, 0)), f, op) // rate M+
op' = (T x, T × V × V y) →
      if (π1(y) - x > 0) then π3(y) / 2(π1(y) - x) else π2(y)
h = combine(g1, g2, op') // rate M+

```

The query  $g_1$  passes along the first timestamp, and the query  $g_2$  calculates the vector  $X_n$ . The top-level query  $h$  calculates the average of the piecewise-linear interpolant when  $n \geq 2$ , and returns the value  $v_1$  when  $n = 1$ .

**Quantitative concatenation.** Suppose that we want to perform two streaming computations in sequence: first execute the query  $f : \text{QRE}\langle D, A \rangle$ , then the query  $g : \text{QRE}\langle D, B \rangle$ , and finally combine the two results using the operation  $op : A \times B \rightarrow C$ .



More specifically, we split the input stream into two parts  $w = w_1w_2$ , process the first part  $w_1$  according to  $f$  with output  $f(w_1)$ , process the second part  $w_2$  according to  $g$  with output  $g(w_2)$ , and produce the final result  $op(f(w_1), g(w_2))$  by applying  $op$  to the intermediate results.

The query  $\text{split}(f, g, op) : \text{QRE}\langle D, C \rangle$  describes this computation. In order for this construction to be well-defined as a function, every input  $w$  that matches  $\text{split}(f, g, op)$  must be uniquely decomposable into  $w = w_1w_2$  with  $w_1$  matching  $f$  and  $w_2$  matching  $g$ . In other words, the rates of  $f$  and  $g$  must be unambiguously concatenable.

The binary  $\text{split}$  construct extends naturally to more than two arguments. For example, the ternary version would be  $\text{split}(f, g, h, (x, y, z) \rightarrow op(x, y, z))$ .

**Example 11.** For the stream of monitored patients, we say that a measurement is *high-risk* if its value exceeds 50. The query  $h : \text{QRE}\langle D_P, V \rangle$  below matches a sequence of measurements containing at least one high-risk measurement, and returns the maximum value after the last occurrence of a high-risk measurement.

```

f1 : QRE⟨DP, Ut⟩ = atom(x → x.typ = M, x → def)           // rate M
g1 : QRE⟨DP, Ut⟩ = iter(eps(def), f1, (x, y) → def)       // rate M*
f2 : QRE⟨DP, Ut⟩ = atom(x → x.typ = M and x.val > 50, x → def) // M(v > 50)
f3 : QRE⟨DP, V⟩ = atom(x → x.typ = M and x.val ≤ 50, x → x.val) // M(v ≤ 50)
g3 : QRE⟨DP, V⟩ = iter(eps(-∞), f3, (x, y) → max(x, y)) // M(v ≤ 50)*
h : QRE⟨DP, V⟩ = split(g1, f2, g3, (x, y, z) → z) // M* · M(v > 50) · M(v ≤ 50)*

```

The query  $g_1$  processes a sequence of measurements and returns nothing, the query  $f_2$  processes a single high-risk measurement and returns nothing, and the query  $g_3$  processes a sequence of non-high-risk measurements and returns the maximum value. The top-level query  $h$  executes  $g_1$ ,  $f_2$  and  $g_3$  in sequence and returns the output of  $g_3$ .

**Global choice.** Given queries  $f$  and  $g$  of the same type with disjoint rates  $r$  and  $s$ , the query  $\text{or}(f, g)$  applies either  $f$  or  $g$  to the input stream depending on which one is defined. The rate of  $\text{or}(f, g)$  is the union  $r \sqcup s$ . This *choice* construction allows a case analysis based on a global regular property of the input stream. In our patient example, suppose we want to compute a statistic across days, where the contribution of each day is computed differently depending on whether or not a specific physiological event occurs sometime during the day. Then, we can write a query summarizing the daily activity with a rate capturing good days (the ones without any significant event) and a different query with a rate capturing bad days, and iterate over their disjoint union.

**Example 12.** For the stream of monitored patients, we describe a query that processes a nonempty sequence of episodes (hence, with rate  $(B \cdot M^+ \cdot E)^+$ ) and outputs at the end of each episode its summary. The summary of the episode is the average of all measurements if there are no high-risk measurements (high-risk: value exceeds 50), otherwise it is the average of only the high-risk measurements. We start with query  $g$ , which processes a nonempty sequence of non-high-risk measurements and returns the sum/count:

$$\begin{aligned} \varphi : D_p \rightarrow \text{Bool} &= x \rightarrow x.\text{typ} = M \text{ and } x.\text{val} \leq 50 && // \text{ predicate on } D_p \\ f : \text{QRE}\langle D_p, V \times V \rangle &= \text{atom}(\varphi, x \rightarrow (x.\text{val}, 1)) && // \text{ rate } M(v \leq 50) \\ g : \text{QRE}\langle D_p, V \times V \rangle &= \text{iter}(f, f, (x, y) \rightarrow x + y) && // \text{ rate } M(v \leq 50)^+ \end{aligned}$$

Similarly, the query  $h$  below processes a sequence with at least one high-risk measurement and returns the sum and count of the high-risk measurements. The idea for describing this computation comes from the following observation: the language over the alphabet  $\Sigma = \{a, b\}$  that contains at least one occurrence of  $b$  is denoted by the ambiguous expression  $(a + b)^* b (a + b)^*$ , which is equivalent to the unambiguous  $a^* (ba^*)^+$ .

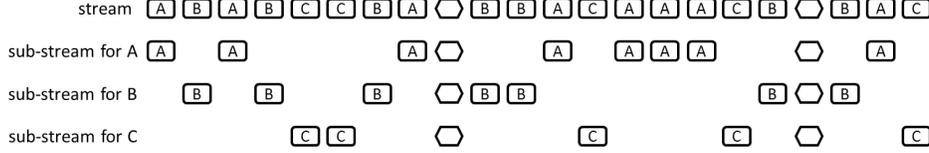
$$\begin{aligned} f_1^* : \text{QRE}\langle D_p, \text{Ut} \rangle &= \text{iter}(\text{eps}(\text{def}), \text{atom}(\varphi), (x, y) \rightarrow \text{def}) && // \text{ rate } M(v \leq 50)^* \\ \psi : D_p \rightarrow \text{Bool} &= x \rightarrow x.\text{typ} = M \text{ and } x.\text{val} > 50 && // \text{ predicate on } D_p \\ f_2 : \text{QRE}\langle D_p, V \times V \rangle &= \text{atom}(\psi, x \rightarrow (x.\text{val}, 1)) && // M(v > 50) \\ g' : \text{QRE}\langle D_p, V \times V \rangle &= \text{split}(f_2, f_1^*, (x, y) \rightarrow x) && // r \triangleq M(v > 50) \cdot M(v \leq 50)^* \\ g'' : \text{QRE}\langle D_p, V \times V \rangle &= \text{iter}(g', g', (x, y) \rightarrow x + y) && // \text{ rate } r^+ \\ h : \text{QRE}\langle D_p, V \times V \rangle &= \text{split}(f_1^*, g'', (x, y) \rightarrow y) && // \text{ rate } M(v \leq 50)^* \cdot r^+ \end{aligned}$$

We have written the queries  $g$  and  $h$  which process sequences of measurements differently based on the occurrence of high-risk measurements. The top-level query  $m : \text{QRE}\langle D_p, V \rangle$  is then given below:

$$\begin{aligned} k : \text{QRE}\langle D_p, V \times V \rangle &= \text{or}(g, h) && // \text{ rate } M^+ \\ f_B : \text{QRE}\langle D_p, D_p \rangle &= \text{atom}(x \rightarrow x.\text{typ} = B) && // \text{ rate } B \\ f_E : \text{QRE}\langle D_p, D_p \rangle &= \text{atom}(x \rightarrow x.\text{typ} = E) && // \text{ rate } E \\ \text{ep} : \text{QRE}\langle D_p, V \rangle &= \text{split}(f_B, k, f_E, (x, y, z) \rightarrow \pi_1(y) / \pi_2(y)) && // \text{ rate } B \cdot M^+ \cdot E \\ m : \text{QRE}\langle D_p, V \rangle &= \text{iter}(\text{ep}, \text{ep}, (x, y) \rightarrow y) && // \text{ rate } (B \cdot M^+ \cdot E)^+ \end{aligned}$$

To see why the rate of  $\text{or}(g, h)$  is  $M^+$ , it suffices to notice that the regular expressions  $(a \sqcup b)^+$  and  $a^+ \sqcup a^* (ba^*)^+$  are equivalent.

**Key-based partitioning.** The input data stream for our running example contains measurements from different patients, and suppose we have written a query  $f$  that computes a summary of data items corresponding to a single patient. Then, to compute an aggregate across patients, the most natural way is to partition the input stream by a key, the patient



**Figure 3.** Partitioning a stream into several parallel sub-streams according to a key (letter in box).

identifier in this case, supply the corresponding projected sub-stream to a copy of  $f$ , one per key, and collect the set of resulting values.

In order to synchronize the per-key computations, we specify a predicate  $\varphi_S : D \rightarrow \text{Bool}$  which defines the *synchronization elements*. The rest of the elements, which satisfy the negation  $\neg\varphi_S$ , are the *keyed elements*. We typically write  $K$  for the set of keys, and  $map : D \rightarrow K$  for the function that projects the key from an item (the value of  $map$  on synchronization items is irrelevant). For the patient input data type of Example 1 we choose:  $\varphi_S = (x \rightarrow x.\text{typ} = D)$  and  $K = \text{PID}$ . The partitioning ensures that the synchronization elements are preserved so that the outputs of different copies of  $f$  are synchronized correctly (for example, if each  $f$  outputs a patient summary at the end of the day, then each sub-stream needs to contain all the end-of-day markers). Note that the output of such a composite streaming function is a mapping  $T : \text{Map}\langle K, C \rangle$  from keys to values, where  $C$  is the output type of  $f$  and  $T(k)$  is the output of the computation of  $f$  for key  $k$ . This key-based partitioning operation is our analog of the map-reduce operation [25] and lends naturally to distributed processing.

We describe the partitioning of the input stream using terminology from concurrent programming. For every key  $k$ , imagine that there is a thread that receives and processes the sub-stream with the data items that concern  $k$ . This includes *all* synchronization items, and those keyed data items  $x$  for which  $map(x) = k$ . So, an item of  $D$  is sent to only one thread (as prescribed by the key), but an item satisfying  $\varphi_S$  is sent to all threads. See Figure 3 for an illustration of the partitioning into sub-streams. Each thread computes independently, and the synchronization elements are used for collecting the results of the threads. We specify a symbolic regular expression  $r$  over  $D$ , which *enforces* a rate of output for the overall computation. For example, if  $r = (((\neg D)^* \cdot D)^2)^*$  then we intend to have output every other day. The rate should only specify sequences that end in a synchronization item.

Suppose  $f : \text{QRE}\langle D, C \rangle$  is a query that describes the per-key (i.e., per-thread) computation, and  $r$  is the overall output rate that we want to enforce. Then, the query

$$\text{map-collect}(\varphi_S, \text{map}, f, r) : \text{QRE}\langle D, \text{Map}\langle K, C \rangle \rangle$$

describes the simultaneous computation for all keys, where the overall output is given whenever the stream matches  $r$ . The overall output is the map obtained by collecting the outputs of all threads that match. W.l.o.g. we assume that the rate of  $f$  is contained in  $r$ , and that it only contains streams with at least one occurrence of a keyed data item. The rate  $r$  should only depend on the occurrence of synchronization elements, so we demand that  $r = s[(\neg\varphi_S)^* \varphi_S / \varphi_S]$  where the only predicate that  $s : \text{RE}\langle D \rangle$  is allowed to contain is  $\varphi_S$ . We write  $s[\psi / \varphi]$  to denote the result of replacing every occurrence of  $\varphi$  in  $s$  with  $\psi$ . For example, if  $s = D^*$  (indicating output at every day marker) and  $\varphi_S = D$ ,

then  $r = s[(-\varphi_S)^* \varphi_S / \varphi_S] = ((-D)^* D)^*$ . These restrictions do not affect expressiveness, but are useful for efficient evaluation.

**Example 13.** Suppose we want to output at the end of each day a table with summaries for the patients that had at least one episode within the day. Assuming that the data stream consists of items for a single patient, we first write the query that produces the episode summary for a single patient:

$$\begin{aligned}
f_M : \text{QRE}\langle D_P, V \times V \rangle &= \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow (x.\text{val}, 1)) && // \text{rate } M \\
g : \text{QRE}\langle D_P, V \times V \rangle &= \text{iter}(f, f, (x, y) \rightarrow x + y) && // \text{rate } M^+ \\
f_B : \text{QRE}\langle D_P, D_P \rangle &= \text{atom}(x \rightarrow x.\text{typ} = B) && // \text{rate } B \\
f_E : \text{QRE}\langle D_P, D_P \rangle &= \text{atom}(x \rightarrow x.\text{typ} = E) && // \text{rate } E \\
h : \text{QRE}\langle D_P, V \rangle &= \text{split}(f_B, g, f_E, (x, y, z) \rightarrow \pi_1(y) / \pi_2(y)) && // \text{rate } B \cdot M^+ \cdot E
\end{aligned}$$

The query  $h$ , matches a full episode and returns the average of the measurements of the episode. The daily summary is then given by:

$$\begin{aligned}
k : \text{QRE}\langle D_P, V \rangle &= \text{iter}(h, h, (x, y) \rightarrow \max(x, y)) && // \text{rate } (B \cdot M^+ \cdot E)^+ \\
l : \text{QRE}\langle D_P, V \rangle &= \text{split}(k, \text{atom}(x \rightarrow x.\text{typ} = D), (x, y) \rightarrow x) && // \text{rate } (B \cdot M^+ \cdot E)^+ \cdot D
\end{aligned}$$

The query  $l$  matches days with at least one episode and outputs the maximum episode summary. The top-level query for the stream that concerns all patients is then:

$$\begin{aligned}
m &= \text{map-collect}(x \rightarrow x.\text{typ} = D, x \rightarrow x.\text{pId}, l, (-D)^* \cdot D) && // \text{rate } (-D)^* \cdot D \\
n &= \text{iter}(m, m, (x, y) \rightarrow y) && // \text{rate } ((-D)^* \cdot D)^+
\end{aligned}$$

The synchronization items for the query  $m : \text{QRE}\langle D_P, \text{Map}\langle \text{PID}, V \rangle \rangle$  are the day markers, and the keys are the patient identifiers. The rate of  $l$  satisfies the typing restrictions of the definition, because every sequence that matches  $l$  contains at least one keyed item and also matches the overall rate  $(-D)^* \cdot D$ . So, the query  $m$  processes a single day and outputs the table of daily summaries for all patients that have had an episode in the day. Finally, the query  $n$  iterates  $m$  for every consecutive day.

**Streaming composition.** A natural operation for query languages over streaming data is streaming composition: given two streaming queries  $f$  and  $g$ ,  $f \gg g$  represents the computation in which the stream of outputs produced by  $f$  is supplied as the input stream to  $g$ . Such a composition is useful in setting up the query as a pipeline of several stages. We allow the operation  $\gg$  to appear *only at the top-level* of a query. So, a general query is a pipeline of  $\gg$ -free queries. At the top level, no type checking needs to be done for the rates, so we do not define the function  $R$  for queries  $f \gg g$ .

**Example 14.** Suppose the input stream concerns a single patient, and we want to compute at the end of each day the minimum and maximum measurement within the day. The first stage of the computation filters out the irrelevant  $B$  and  $E$  markers:

$f : \text{QRE}\langle D_P, D_P \rangle = \text{atom}(x \rightarrow x.\text{typ} = M \text{ or } x.\text{typ} = D) \quad // \text{ rate } M \sqcup D$   
 $g : \text{QRE}\langle D_P, \text{Ut} \rangle = \text{iter}(\text{eps}(\text{def}), \text{atom}(), (x, y) \rightarrow \text{def}) \quad // \text{ rate } (M \sqcup D \sqcup B \sqcup E)^*$   
 $h : \text{QRE}\langle D_P, D_P \rangle = \text{split}(g, f, (x, y) \rightarrow y) \quad // \text{ rate } (M \sqcup D \sqcup B \sqcup E)^* \cdot (M \sqcup D)$

The query  $h$  matches any sequence that ends in a measurement or day marker and returns the last item. Thus,  $h$  filters out  $B$  and  $E$  items. For the second stage of the computation, we assume that the stream consists of only  $M$  and  $D$  items.

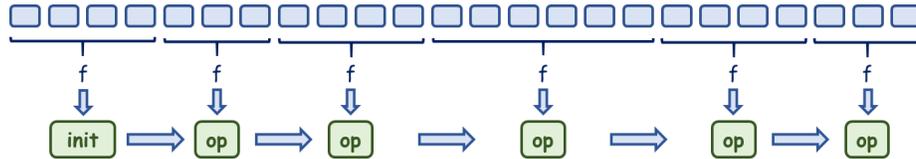
$k : \text{QRE}\langle D_P, V \rangle = \text{atom}(x \rightarrow x.\text{typ} = M, x \rightarrow x.\text{val}) \quad // \text{ rate } M$   
 $l : \text{QRE}\langle D_P, V \rangle = \text{iter}(\text{eps}(-\infty), k, (x, y) \rightarrow \max(x, y)) \quad // \text{ rate } M^*$   
 $m : \text{QRE}\langle D_P, V \rangle = \text{split}(l, \text{atom}(x \rightarrow x.\text{typ} = D), (x, y) \rightarrow x) \quad // \text{ rate } M^* \cdot D$   
 $n : \text{QRE}\langle D_P, V \rangle = \text{iter}(m, m, (x, y) \rightarrow y) \quad // \text{ rate } (M^* \cdot D)^+$

The top-level query is then the pipeline  $h \gg n$  and its domain is  $((M \sqcup B \sqcup E)^* \cdot D)^+$ .

## 2. Common Patterns

The core language of Figure 2 is expressive enough to describe many common stream transformations. We present below several derived patterns, including stream filtering, stream mapping, and aggregation over windows.

**Iteration at least once.** Let  $f : \text{QRE}\langle D, A \rangle$  be a query with output type  $A$ ,  $\text{init} : A \rightarrow B$  be the initialization function, and  $\text{op} : B \times A \rightarrow B$  be the aggregation function. The query  $\text{iter}_1(f, \text{init}, \text{op})$ , with output type  $B$ , splits the input stream  $w$  unambiguously into consecutive parts  $w_1 w_2 \dots w_n$  each of which matches  $f$ , applies  $f$  to each  $w_i$  producing a sequence of output values  $a_0 a_1 a_2 \dots a_n$ , i.e.  $a_i = f(w_i)$ , and combines the results  $a_1 a_2 \dots a_n$  using the list iterator *left fold* with start value  $\text{init}(a_0) \in B$  and accumulation operation  $\text{op} : B \times A \rightarrow B$ .



The construct  $\text{iter}_1$  can be encoded using  $\text{iter}$  as follows:

$$\text{iter}_1(f, \text{init}, \text{op}) \triangleq \text{iter}(\text{apply}(f, \text{init}), f, \text{op}).$$

The type of  $\text{iter}_1(f, \text{init}, \text{op})$  is  $\text{QRE}\langle D, B \rangle$  and its rate is  $R(f)^+$ . We use the abbreviation  $\text{iter}_1(f, \text{op})$  for the common case where  $A = B$  and  $\text{init}$  is the identity function.

**Matching without output.** Suppose  $r$  is an unambiguous symbolic regex over the data item type  $D$ . The query  $\text{match}(r)$ , whose rate is equal to  $r$ , does not produce any output when it matches. This is essentially the same as producing `def` as output for a match. The `match` construct can be encoded as follows:

$$\begin{aligned}\text{match}(\varphi) &\triangleq \text{atom}(\varphi, x \rightarrow \text{def}) \\ \text{match}(r_1 \sqcup r_2) &\triangleq \text{or}(\text{match}(r_1), \text{match}(r_2)) \\ \text{match}(r_1 \cdot r_2) &\triangleq \text{split}(\text{match}(r_1), \text{match}(r_2), (x, y) \rightarrow \text{def}) \\ \text{match}(r^*) &\triangleq \text{iter}(\text{eps}(\text{def}), \text{match}(r), (x, y) \rightarrow \text{def})\end{aligned}$$

An easy induction establishes that  $R(\text{match}(r)) = r$ .

**Stream filtering.** Let  $\varphi$  be a predicate over the type of input data items  $D$ . We want to describe the streaming transformation that filters out all items that do not satisfy  $\varphi$ . We implement this with the query  $\text{filter}(\varphi)$ , which matches all stream prefixes that end with an item satisfying  $\varphi$ .

$$\text{filter}(\varphi) \triangleq \text{split}(\text{match}(\text{true}_D^*), \text{atom}(\varphi), (x, y) \rightarrow y)$$

The type of  $\text{filter}(\varphi)$  is  $\text{QRE}\langle D, D \rangle$  and its rate is  $\text{true}_D^* \cdot \varphi$ .

**Stream mapping.** The mapping of an input stream of type  $D$  to an output stream of type  $C$  according to the operation  $op : D \rightarrow C$  is given by the following query:

$$\text{map}(op) \triangleq \text{split}(\text{match}(\text{true}_D^*), \text{atom}(\text{true}_D, op), (x, y) \rightarrow y).$$

Its type is  $\text{QRE}\langle D, C \rangle$  and its rate is  $\text{true}_D^* \cdot \text{true}_D = \text{true}_D^+$ .

**Example 15.** Using filtering, mapping and streaming composition we can implement the average of a sequence of scalars with a very common idiom:

$$\begin{aligned}\mathbf{f} : \text{QRE}\langle V, V \times V \rangle &= \text{map}(x \rightarrow (x, 1)) && // \text{rate } V^+ \\ \mathbf{g} : \text{QRE}\langle V \times V, V \times V \rangle &= \text{iter}_1(\text{atom}(), (x, y) \rightarrow x + y) && // \text{rate } (V \times V)^+ \\ \mathbf{h} : \text{QRE}\langle V \times V, V \rangle &= \text{map}(x \rightarrow \pi_1(x) / \pi_2(x)) && // \text{rate } (V \times V)^+\end{aligned}$$

and the top-level query is the pipeline  $\mathbf{f} \gg \mathbf{g} \gg \mathbf{h}$ , whose type is  $\text{QRE}\langle V, V \rangle$ .

**Iteration exactly  $n$  times.** Let  $n \geq 1$  and  $\mathbf{f} : \text{QRE}\langle D, A \rangle$  be a query to iterate exactly  $n$  times. The aggregation is specified by the initialization function  $\text{init} : A \rightarrow B$  (for the first value) and the aggregation function  $op : B \times A \rightarrow B$ . The construct  $\text{iter}^n$  describes iteration (and aggregation) exactly  $n$  times, and can be encoded as follows:

$$\begin{aligned}\text{iter}^1(\mathbf{f}, \text{init}, op) &\triangleq \text{apply}(\mathbf{f}, \text{init}) \\ \text{iter}^{n+1}(\mathbf{f}, \text{init}, op) &\triangleq \text{split}(\text{iter}^n(\mathbf{f}, \text{init}, op), \mathbf{f}, op)\end{aligned}$$

The type of  $\text{iter}^n(\mathbf{f}, \text{init}, op)$  is  $\text{QRE}\langle D, B \rangle$  and its rate is  $R(\mathbf{f})^n$  ( $n$ -fold concatenation).

**Pattern-based tumbling windows.** The term tumbling windows is used to describe the splitting of the stream into contiguous non-overlapping regions [6]. Suppose we want to describe the streaming function that iterates  $f$  at least once and reports the result given by  $f$  at every match. The following query expresses this behavior:

$$\text{iter-last}(f) \triangleq \text{iter}(f, f, (x,y) \rightarrow y).$$

The rate of  $\text{iter-last}(f)$  is equal to  $R(f)^+$ .

**Example 16.** Suppose that the query  $f : \text{QRE}(D_p, V)$  has rate  $r = (B \cdot M^+ \cdot E)^* \cdot D$  and computes the daily summary for a single patient. Then, the query  $\text{iter-last}(f)$  has rate  $r^+$  and computes the daily summary at the end of every day. Finally, the query

$$\text{iter-last}(f) \gg \text{iter}_1(\text{atom}(), (x,y) \rightarrow \max(x,y))$$

computes at the end of every day the maximum daily summary so far. Notice that this query can be equivalently expressed without  $\gg$  as  $\text{iter}_1(f, (x,y) \rightarrow \max(x,y))$ .

**Sliding windows (slide by pattern).** To express a policy such as “output the statistical summary of events in the past ten hours every five minutes” existing relational query languages provide an explicit sliding window primitive [6]. We can support this primitive, which can be compiled into the base language by massaging the input data stream with the introduction of suitable tags (marking five-minute time intervals in this example). The insertion of the tags then allows to express both the window and the sliding using very general regular patterns. Let  $n \geq 1$  be the size of the window, and  $f : \text{QRE}(D, A)$  be the query that processes a unit pattern. The aggregation over the window is specified by the function  $\text{init} : A \rightarrow B$  for initialization and the aggregation function  $\text{op} : B \times A \rightarrow B$ . We give a query that computes the aggregation over the last  $n$  units of the stream (or over all units if the stream has less than  $n$  units):

$$\begin{aligned} g &= \text{or}(\text{iter}^1(f, \text{init}, \text{op}), \dots, \text{iter}^{n-1}(f, \text{init}, \text{op})) \\ h &= \text{split}(\text{match}(R(f)^*), \text{iter}^n(f, \text{init}, \text{op}), (x,y) \rightarrow y) \end{aligned}$$

and  $\text{wnd}(f, n, \text{init}, \text{op}) = \text{or}(g, h)$  with rate  $R(f)^+$ .

### 3. The Yahoo Streaming Benchmark

The Yahoo Benchmark [14] specifies a stream of advertisement-related events for an analytics pipeline. It specifies a set of campaigns and a set of advertisements, where each ad belongs to exactly one campaign. The static map from ads to campaigns is computed ahead-of-time and stored in memory. Each element of the data stream is of the form

$$(\text{userId}, \text{pageId}, \text{adId}, \text{eventType}, \text{eventTime}),$$

indicating the interaction of a user with an advertisement, where  $\text{eventType}$  is one of  $\{\text{view}, \text{click}, \text{purchase}\}$ . The component  $\text{eventTime}$  is the timestamp of the event.

### 3.1. First Yahoo query

The basic benchmark query (described in [14]) computes, at the end of every second, a map from each campaign to the number of views associated with that campaign within the last second. For each event tuple, this involves a lookup to determine the campaign associated with the advertisement viewed. The reference implementation published with the Yahoo benchmark involves a multi-stage pipeline:

- (a) *stage 1*: filter view events,
- (b) *stage 2*: project the ad id from each view tuple,
- (c) *stage 3*: lookup the campaign id of each ad,
- (d) *stage 4*: compute for every one-second window the number of events (views) associated with each campaign.

The query involves key-based partitioning on only one property, namely the derived campaign id of the event. We present three ways of expressing this query in StreamQRE. We assume w.l.o.g. that the stream also contains events  $S$  that serve as end-of-second markers.

**Implementation (I)** We reproduce faithfully the reference implementation of the Yahoo benchmark [14] by constructing the following multi-stage pipeline:

- (1) query  $g_1$ : filter view and end-of-second events,
- (2) query  $g_2$ : project the ad id from each view tuple,
- (3) query  $g_3$ : lookup the campaign id of each advertisement,
- (4) query  $g_7$ : compute for every one-second window the number of events (views) associated with each campaign.

We write  $ADID$  for the type of ad identifiers,  $CID$  for the type of campaign identifiers, and  $D_Y$  for the data type of the input events.

$$\begin{aligned}
 g_1 &: \text{QRE}\langle D_Y, D_Y \rangle = \text{filter}(x \rightarrow x = S \text{ or } x.\text{isView}) \\
 g_2 &: \text{QRE}\langle D_Y, ADID \cup \{S\} \rangle = \text{map}(x \rightarrow \text{if } (x = S) \text{ then } S \text{ else } x.\text{adId}) \\
 g_3 &: \text{QRE}\langle ADID \cup \{S\}, CID \cup \{S\} \rangle = \text{map}(x \rightarrow \text{if } (x = S) \text{ then } S \text{ else } \text{lookup}(x.\text{adId})) \\
 g_4 &: \text{QRE}\langle CID \cup \{S\}, \text{Nat} \rangle = \text{iter}(\text{eps}(0), \text{atom}(x \rightarrow x \neq S), (x, y) \rightarrow x + 1) \\
 g_5 &: \text{QRE}\langle CID \cup \{S\}, \text{Nat} \rangle = \text{split}(g_4, \text{atom}(x \rightarrow x = S), (x, y) \rightarrow x) \\
 g_6 &: \text{QRE}\langle CID \cup \{S\}, \text{Nat} \rangle = \text{map-collect}(x \rightarrow x = S, x \rightarrow x.\text{cId}, g_5, (\neg S)^* \cdot S) \\
 g_7 &: \text{QRE}\langle CID \cup \{S\}, \text{Nat} \rangle = \text{iter-last}(g_6)
 \end{aligned}$$

The function  $\text{lookup} : ADID \rightarrow CID$  models the mapping of an ad to the campaign it belongs to, and  $x.\text{isView}$  abbreviates the boolean expression  $(x.\text{eventType} = \text{view})$ . The auxiliary query  $g_5$ , with rate  $(\neg S)^* \cdot S$ , calculates the length of event sequences of a single campaign for one second. With  $g_7$  we perform (every second) key-based partitioning based on the campaign id  $\text{cId}$ . The top-level query to compute the number of views for each campaign per second is the pipeline

$$g_8 = g_1 \gg g_2 \gg g_3 \gg g_7.$$

Since the domain of  $g_1 \gg g_2 \gg g_3$  is  $D_Y^* \cdot (x \rightarrow x.\text{isView} \text{ or } x = S)$ , and  $g_7$  reports at every  $S$  marker, the domain of  $g_8$  is  $D_Y^* \cdot S$  (output at the end of every second).

**Implementation (II)** The stages  $g_1$ ,  $g_2$  and  $g_3$  of the pipeline of implementation (I) can be collapsed into a single query  $h_3$  as follows:

$$\begin{aligned} h_1 &: \text{QRE}\langle D_Y, \text{CID} \cup \{S\} \rangle = \text{atom}(x \rightarrow x.\text{isView}, x \rightarrow \text{lookup}(x.\text{adId})) \\ h_2 &: \text{QRE}\langle D_Y, \text{CID} \cup \{S\} \rangle = \text{or}(h_1, \text{atom}(x \rightarrow x = S, x \rightarrow S)) \\ h_3 &: \text{QRE}\langle D_Y, \text{CID} \cup \{S\} \rangle = \text{split}(\text{match}(((D_Y x) \rightarrow \text{true})^*), h_2, (x, y) \rightarrow y) \end{aligned}$$

The benchmark query can now be written as a two-stage pipeline  $h_4 = h_3 \gg g_7$ .

**Implementation (III)** The previous implementation uses the streaming composition operator to simplify the map-collect part of the query. The channel handling the events of each campaign assumes that all incoming events correspond to views, and therefore simply counts the number of tuples flowing in. We can eliminate the streaming composition by inspecting the event type of each incoming tuple during per-campaign processing:

$$\begin{aligned} k_1 &: \text{QRE}\langle D_Y, \text{Nat} \rangle = \text{atom}(x \rightarrow x.\text{isView}, x \rightarrow 1) && // \text{rate } V \text{ (view)} \\ k_2 &: \text{QRE}\langle D_Y, \text{Nat} \rangle = \text{atom}(x \rightarrow \text{not}(x.\text{isView} \text{ or } x = S), x \rightarrow 0) && // \text{rate } \neg(V \sqcup S) \\ k_3 &: \text{QRE}\langle D_Y, \text{Nat} \rangle = \text{or}(k_1, k_2) && // \text{rate } \neg S \\ k_4 &: \text{QRE}\langle D_Y, \text{Nat} \rangle = \text{iter}(\text{eps}(0), k_3, (x, y) \rightarrow x + y) && // \text{rate } (\neg S)^* \\ k_5 &: \text{QRE}\langle D_Y, \text{Nat} \rangle = \text{split}(k_4, \text{atom}(x \rightarrow x = S), (x, y) \rightarrow x) && // \text{rate } (\neg S)^* \cdot S \end{aligned}$$

So,  $k_5$  counts the number of views in a stream of the form  $(\neg S)^* \cdot S$ . We now have a third way of representing the benchmark query:

$$\begin{aligned} k_6 &: \text{QRE}\langle D_Y, \text{Map}\langle \text{CID}, \text{Nat} \rangle \rangle = \text{map-collect}(x \rightarrow x = S, \\ &\quad x \rightarrow \text{lookup}(x.\text{adId}), k_5, (\neg S)^* \cdot S) \\ k_7 &: \text{QRE}\langle D_Y, \text{Map}\langle \text{CID}, \text{Nat} \rangle \rangle = \text{iter-last}(k_6). \end{aligned}$$

In implementations (I) and (II) the queries only looked up the campaign ids for view events, while query  $k_7$  computes the campaign id for *each* incoming tuple. It therefore makes more campaign id lookups.

### 3.2. Second Yahoo query

We extend the Yahoo benchmark with a more complex query. An important part of organizing a marketing campaign is quantifying how successful ads are. We define *success* as the number of users who purchase the product after viewing an ad for it. Our query outputs, at the end of every second, a map from campaigns to the most successful ad of the campaign so far, together with its success score.

Assume that we have fixed a specific ad and a specific user, and the stream consists only of events for these. The pattern for *success*, given by the regular expression

$$r = (\neg V)^* \cdot V \cdot (\neg P)^* \cdot P \cdot D_Y^* \cdot S,$$

indicates that the user purchases the product after seeing the ad for it. For simplicity, we have written  $V$  to denote the occurrence of view events, and  $P$  for purchase events. The trailing pattern  $D_Y^* \cdot S$  is used for matching until the end of an end-of-second marker.

Now, suppose that we have fixed a specific ad, and we want to compute its *score*: the number of users that have purchased the product after viewing the ad.

$$\begin{aligned} g_1 : \text{QRE}\langle D_Y, \text{Map}\langle \text{UID}, \text{Ut} \rangle \rangle &= \text{map-collect}(x \rightarrow x = S, \\ &\quad x \rightarrow x.\text{userId}, \text{match}(r), D_Y^* \cdot S) \\ g_2 : \text{QRE}\langle D_Y, \text{Nat} \rangle &= \text{apply}(g_1, x \rightarrow x.\text{size}) \end{aligned}$$

where  $\text{UID}$  is the set of user identifiers, and  $\text{size} : \text{Map}\langle K, C \rangle \rightarrow \text{Nat}$  returns the size of a map data structure (number of keys that are mapped to some value). The rate of  $g_2$  is  $D_Y^* \cdot S$ , that is,  $g_2$  produces at the end of every second the success score of the ad. Moreover, we observe that the specified rate  $D_Y^* \cdot S$  is equivalent to  $((\neg S)^* S)^+$ .

Given the event stream of a single campaign, we can divide it into sub-streams of individual ads, compute the success score for each one of them, and thus determine the most successful ad of the campaign:

$$\begin{aligned} g_3 : \text{QRE}\langle D_Y, \text{Map}\langle \text{ADID}, \text{Nat} \rangle \rangle &= \text{map-collect}(x \rightarrow x = S, x \rightarrow x.\text{adId}, g_2, D_Y^* \cdot S) \\ g_4 : \text{QRE}\langle D_Y, \text{ADID} \times \text{Nat} \rangle &= \text{apply}(g_3, x \rightarrow x.\text{argmax}), \end{aligned}$$

where  $\text{argmax} : \text{Map}\langle K, C \rangle \rightarrow K \times C$  (for a type  $C$  that is linearly ordered) calculates the key-value pair  $(k, c)$  that has the maximum value. In query  $g_4$ , it calculates the pair  $(\text{adId}, \text{score})$  for the ad with the maximum score. Finally, we use the `map-collect` construct to map each campaign id to the most successful ad:

$$g_5 : \text{QRE}\langle D_Y, \text{Map}\langle \text{CID}, \text{ADID} \times \text{Nat} \rangle \rangle = \text{map-collect}(x \rightarrow x = S, x \rightarrow x.\text{cId}, g_4, D_Y^* \cdot S)$$

The rates of  $g_1$ ,  $g_2$ ,  $g_3$ ,  $g_4$  and  $g_5$  are all equal to  $D_Y^* \cdot S$ . They produce output at every  $S$  marker occurrence in the stream.

#### 4. The NEXMark Streaming Benchmark

The Niagara Extension to XMark benchmark (NEXMark) [15] concerns the monitoring of an on-line auction system, such as eBay. Four kinds of events are recorded in the event stream: (a) **Person** events, which describe the registering of a new person to the auction system, (b) **Item** events, which mark the start of an auction for a specified item, (c) **Close** events, which mark the end of an auction for a specified item, and (d) **Bid** events, which record the bids made for items that are being auctioned.

```

Person(personId,name,ts)
Item(itemId,sellerId,initPrice,ts,dur,category)
Close(itemId,ts)
Bid(itemId,bidderId,bidIncrement,ts)

```

Every event contains the field `ts`, which is the timestamp of when the event occurred. Every new auction event (of type `Item`) specifies an initial price `initPrice` for the item, the duration `dur` of the auction, and the category to which the item belongs. Every bid event contains the bid increment, that is, the increment by which the previous bid is raised. So, to find the current bid for an item we need to add the initial price of the item together with all the bid increments for the item so far. We will describe two out of the five queries that are considered in [2], and which are minor variants of some of the queries of the NEXMark benchmark:

**Query.** Calculate the number of currently open auctions. The output should be updated at every auction start and close.

```

f : QRE⟨DN,DN⟩ = filter(x->x.isItem or x.isClose)
g : QRE⟨DN,Nat⟩ = map(x->if (x.isItem) then +1 else -1)
h : QRE⟨Nat,Nat⟩ = iter(eps(0), atom(), (x,y)->x+y)
k : QRE⟨DN,Nat⟩ = f >> g >> h

```

We write  $D_N$  for the input data type of the NEXMark benchmark.

**Query.** Find the item with the most bids in the last 24 hours. The output should be updated every minute. We assume that the stream has end-of-minute markers  $M$ .

First, suppose that we process a stream consisting only of bid events that ends with an end-of-minute marker. We compute a bid count for every item that appears:

```

f : QRE⟨DN,Nat⟩ = iter(eps(0), atom(x->x.isBid), (x,y)->x+1) // B*
g : QRE⟨DN,Nat⟩ = split(f, atom(x->x=M), (x,y)->x) // rate B* · M
h : QRE⟨DN,Map⟨IID,Nat⟩⟩ = map-collect(x->x=M,
                                     x->x.itemId, g, (¬M)* · M) // rate (¬M)* · M
k : QRE⟨DN,MSet⟨IID⟩⟩ = apply(h, x->x.toMSet()) // rate (¬M)* · M

```

The type of item identifiers is `IID`, and `toMSet : Map⟨K,Nat⟩ → MSet⟨K⟩` is an operation that turns a map object into a multiset. Now, we write the top-level query:

```

l : QRE⟨DN,MSet⟨IID⟩⟩ = wnd(k, 24 · 60, x->x, (x,y)->x ⊔ y) // rate ((¬M)* · M)+
m : QRE⟨DN,DN⟩ = filter(x->x.isBid or (x=M)) // rate DN* · (B ⊔ M)
n : QRE⟨DN,IID × Nat⟩ = m >> apply(l, x->x.argmax)

```

We write  $\uplus$  to denote multiset union. The function `argmax : MSet⟨C⟩ → C × Nat` returns the member of the multiset with the highest count.

```

// Process a single measurement: rate M
QRe<DPatients, DPatients> meas =
    Q.atomic(x -> x.isMeasurement(), x -> x);

// Sum of sequence of measurements: rate M*
QRe<DPatients, Double> sum =
    Q.iter(Q.eps(0.0), meas, (x,y) -> x+y.getValue());

// Length of sequence of measurements: rate M*
QRe<DPatients, Integer> count =
    Q.iter(Q.eps(0), meas, (x,y) -> x+1);

// Average of sequence of measurements: rate M*
QRe<DPatients, Double> measAvg =
    Q.combine(sum, count, (x,y) -> x/y);

Iterator<DPatients> stream = ... // input stream

// evaluator for the query
Eval<DPatients, Double> e = measAvg.getEval();

// execution loop
Double output = e.start(); // returns null, if undefined
while (stream.hasNext()) {
    DPatients d = stream.next();
    output = e.next(d); // returns null, if undefined
}

```

**Figure 4.** Computing the average of a nonempty sequence of measurements.

## 5. The StreamQRE Library in Java

StreamQRE has been implemented as a Java library [13] in order to facilitate the easy integration with user-defined types and operations. The implementation covers all the core constructs of Figure 2, and also provides optimizations for the derived constructs of Section 2 (stream filtering, stream mapping, sliding windows, etc.).

Figure 4 gives a simple example that illustrates how to program with the StreamQRE Java library. The query `measAvg` describes the computation of the average of a sequence of measurements (for the patient data stream). The method `getEval`, which stands for “get evaluator”, is used to obtain an object that encapsulates the evaluation algorithm for the query. On this evaluator object, the methods `start` and `next` are used to initialize the algorithm and consume data items respectively.

## 6. Conclusion

We have given an introduction to the StreamQRE language [2], a high-level formalism for processing streaming data. The query language integrates two paradigms for pro-

programming with streams: streaming relational languages with windowing constructs, and state-machine-based models for pattern-matching and performing sequence-aware computations. The language consists of a small but powerful core language, which has a formal denotational semantics and a decidable type system. The expressiveness of the language has been illustrated by encoding common patterns and programming significant examples.

A query of the StreamQRE language can be compiled into a streaming algorithm with strong efficiency guarantees [2], both for space usage and processing time per element. An experimental evaluation of StreamQRE is reported in [2], which shows that the StreamQRE implementation is competitive with popular streaming engines such as RxJava [26], Esper [27], and Flink [28].

## Acknowledgements

We thank our collaborators Zachary Ives, Sanjeev Khanna and Mukund Raghothaman. This research was supported by NSF Expeditions award CCF 1138996.

## References

- [1] R. Alur, E. Berger, A. Drobnis, L. Fix, K. Fu, G. Hager, D. Lopresti, K. Nahrstedt, E. Mynatt, S. Patel, J. Rexford, J. Stankovic, and B. Zorn. Systems computing challenges in the Internet of Things. In *Computing Community Consortium Whitepaper*, 2016.
- [2] Konstantinos Mamouras, Mukund Raghothaman, Rajeev Alur, Zachary G. Ives, and Sanjeev Khanna. StreamQRE: Modular specification and efficient evaluation of quantitative queries over streaming data. 2017. manuscript.
- [3] Rajeev Alur, Dana Fisman, and Mukund Raghothaman. Regular programming for quantitative properties of data streams. In *Proceedings of the 25th European Symposium on Programming (ESOP '16)*, pages 15–40, 2016.
- [4] Shivnath Babu and Jennifer Widom. Continuous queries over data streams. *ACM Sigmod Record*, 30(3):109–120, 2001.
- [5] Daniel J. Abadi, Don Carney, Ugur Cetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. Aurora: A new model and architecture for data stream management. *The VLDB Journal*, 12(2):120–139, 2003.
- [6] Jin Li, David Maier, Kristin Tufte, Vassilis Papadimos, and Peter A. Tucker. Semantics and evaluation techniques for window aggregates in data streams. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, pages 311–322. ACM, 2005.
- [7] Arvind Arasu, Shivnath Babu, and Jennifer Widom. The CQL continuous query language: Semantic foundations and query execution. *The VLDB Journal*, 15(2):121–142, 2006.
- [8] Barzan Mozafari, Kai Zeng, and Carlo Zaniolo. High-performance complex event processing over XML streams. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, SIGMOD '12, pages 253–264. ACM, 2012.
- [9] Mohamed Ali, Badrish Chandramouli, Jonathan Goldstein, and Roman Schindlauer. The extensibility framework in Microsoft StreamInsight. In *Proceedings of the 27th IEEE International Conference on Data Engineering (ICDE '11)*, pages 1242–1253, 2011.
- [10] M. Hirzel, H. Andrade, B. Gedik, G. Jacques-Silva, R. Khandekar, V. Kumar, M. Mendell, H. Nasgaard, S. Schneider, R. Soul, and K. L. Wu. IBM Streams Processing Language: Analyzing big data in motion. *IBM Journal of Research and Development*, 57(3/4):7:1–7:11, 2013.
- [11] Mandana Vaziri, Olivier Tardieu, Rodric Rabbah, Philippe Suter, and Martin Hirzel. Stream processing with a spreadsheet. In *Proceedings of the 28th European Conference on Object-Oriented Programming (ECOOP '14)*, pages 360–384. Springer Berlin Heidelberg, 2014.

- [12] Margus Veanes, Pieter Hooimeijer, Benjamin Livshits, David Molnar, and Nikolaj Bjorner. Symbolic finite state transducers: Algorithms and applications. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '12)*, pages 137–150. ACM, 2012.
- [13] StreamQRE library. <http://www.seas.upenn.edu/~mamouras/StreamQRE/>.
- [14] Sanket Chintapalli, Derek Dagit, Bobby Evans, Reza Farivar, Thomas Graves, Mark Holderbaugh, Zhuo Liu, Kyle Nusbaum, Kishorkumar Patil, Boyang Jerry Peng, and Paul Poulosky. Benchmarking streaming computation engines: Storm, Flink and Spark streaming. In *First Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware*, 2016.
- [15] Pete Tucker, Kristin Tuft, Vassilis Papadimos, and David Maier. NEXMark: A benchmark for queries over data streams, 2002.
- [16] Brian Litt and Zachary Ives. The international epilepsy electrophysiology database. In *Proceedings of the Fifth International Workshop on Seizure Prediction*, 2011.
- [17] Ronald Book, Shimon Even, Sheila Greibach, and Gene Ott. Ambiguity in graphs and expressions. *IEEE Transactions on Computers*, C-20(2):149–153, 1971.
- [18] Richard Edwin Stearns and Harry B. Hunt III. On the equivalence and containment problems for unambiguous regular expressions, regular grammars and finite automata. *SIAM Journal on Computing*, 14(3):598–611, 1985.
- [19] Leonardo De Moura and Nikolaj Bjørner. Satisfiability modulo theories: Introduction and applications. *Communications of the ACM*, 54(9):69–77, 2011.
- [20] Java’s lambda expressions. <https://docs.oracle.com/javase/tutorial/java/java00/lambdaexpressions.html>.
- [21] Margus Veanes, Peli de Halleux, and Nikolai Tillmann. Rex: Symbolic regular expression explorer. In *Proceedings of the 3rd International Conference on Software Testing, Verification and Validation (ICST '10)*, pages 498–507. IEEE, 2010.
- [22] Jacques Sakarovitch. *Elements of Automata Theory*. Cambridge University Press, 2009.
- [23] Manfred Droste, Werner Kuich, and Heiko Vogler, editors. *Handbook of Weighted Automata*. Springer, 2009.
- [24] Albert Benveniste, Paul Caspi, Stephen A. Edwards, Nicolas Halbwachs, Paul Le Guernic, and Robert de Simone. The synchronous languages 12 years later. *Proceedings of the IEEE*, 91(1):64–83, 2003.
- [25] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [26] ReactiveX: An API for asynchronous programming with observable streams. <http://reactivex.io/>.
- [27] Esper for Java. <http://www.espertech.com/esper/>.
- [28] Apache Flink: Scalable batch and stream data processing. <https://flink.apache.org/>.