

Streamable Regular Transductions

Rajeev Alur^a, Dana Fisman^b, Konstantinos Mamouras^a, Mukund Raghothaman^a, Caleb Stanford^a

^a*University of Pennsylvania, Philadelphia, PA, USA*

^b*Ben-Gurion University, Be'er Sheva, Israel*

Abstract

Motivated by real-time monitoring and data processing applications, we develop a formal theory of quantitative queries for streaming data that can be evaluated efficiently. We consider the model of unambiguous Cost Register Automata (CRAs), which are machines that combine finite-state control (for identifying regular patterns) with a finite set of data registers (for computing numerical aggregates). The definition of CRAs is parameterized by the collection of numerical operations that can be applied to the registers. These machines give rise to the class of *streamable regular transductions* (SR), and to the class of *streamable linear regular transductions* (SLR) when the register updates are *copyless*, i.e. every register appears at most once the right-hand-side expressions of the updates. We give a logical characterization of the class SR (resp., SLR) using MSO-definable transformations from strings to DAGs (resp., trees) without backward edges. Additionally, we establish that the two classes SR and SLR are closed under operations that are relevant for designing query languages. Finally, we study the relationship with weighted automata (WA), and show that CRAs over a suitably chosen set of operations correspond to WA, thus establishing that WA are a special case of CRAs.

Keywords: Cost Register Automata, MSO transductions, regular functions, quantitative automata, stream processing

1. Introduction

Finite-state automata, and the associated class of regular languages, have many appealing properties. In particular, the class of regular languages has robust, and well understood, expressiveness with multiple characterizations,

and is closed under a variety of language operations. Furthermore, a finite-state automaton processes an input string in a single left-to-right pass using only constant space, in compliance with the *streaming* model of computation. This forms the basis of practical applications of automata in design and implementation of query languages for pattern matching in strings.

In a diverse range of applications such as financial tickers, network traffic monitoring, and click-streams of web usage, the core computational problem is to map a stream of data items to a numerical value. For example, suppose the monitoring software at a network router wants to compute the average number of packets per VoIP (Voice over IP) session in the incoming stream of packets. This requires detection of VoIP sessions, which can be characterized by regular patterns in the input stream, and computing numerical aggregates in a hierarchical manner, where the cost of a VoIP session is the number of packets it contains, and the cost of a stream is the average of costs of VoIP sessions it contains. Recent work on *Quantitative Regular Expressions* (QRE) provides a declarative query language for specifying such computation in a modular fashion, and has led to prototype implementations StreamQRE and NetQRE [1, 2, 3]. These quantitative queries involve a mix of regular patterns and numerical aggregates, and crucially, all queries can be evaluated in a streaming manner. In this paper we develop an understanding of the class of streaming transductions definable by such queries: the streamable regular transductions.

A natural starting point for this study is the model of *cost register automata* (CRA) and the associated class of *regular transductions* [4]. A CRA processes a data word, that is, a sequence over $\Sigma \times D$, where Σ is a finite set of tags and D is a set of data values, and outputs a value in D that is computed using a given set of operations over D (for example, D can be the set of natural numbers, and the set of operations can include addition, minimum, and maximum). A CRA has a finite-state control that is updated based only on the tag values of the input string, and a finite set of (write-only) registers that are updated at each step using the given operations. The partial functions from $(\Sigma \times D)^*$ to D computed by the CRAs are called regular with respect to the allowed set of operations. The computation of a regular transduction can be viewed as, first mapping a data word to a tree whose nodes are labeled with the allowed operations, and then evaluating the resulting term, where the former mapping is an MSO-definable string-to-tree transformation. This gives an alternative logical characterization of the class of regular transductions, and allows us to draw upon the well developed

theory of regular tree transformations [5, 6, 7].

The original CRA model is a deterministic machine that processes its input in a single pass. Its registers can hold data values as well as functions represented by terms with parameters, and the update at each step is required to be *copyless*, that is, a register can appear at most once in the right-hand-side expressions of the updates. The CRA model we study in section 2 differs in three ways. First, since our focus is on streamability, and terms can grow linearly with the size of the input stream, we require registers to hold only values. Second, we allow the finite-state control to be updated nondeterministically based on the input tags, but it is required to be *unambiguous* (that is, any accepted data word has exactly one accepting run). While it is known that when registers can hold terms, unambiguous machines are no more expressive than deterministic ones [7], we prove that when registers can hold only values that are updated in a copyless manner, unambiguous machines are strictly more expressive than their deterministic counterparts. The complexity of evaluation for unambiguous CRAs is only a constant (equal to the number of states of the machine) factor more than the deterministic ones. Finally, we consider both copyless and copyful (that is, unrestricted) updates since a CRA has a constant number of registers that hold data values and get updated by executing a constant number of operations while processing each data item. The number of bits needed to store data values can grow with the length of the input stream when the data domain is unbounded, and the exact complexity of the streaming algorithm for evaluation depends on the nature of operations and implementation details.

We call the class of transductions over data words defined by such unambiguous CRAs with registers containing only values and copyful updates as *streamable regular transductions* (SR), and the subclass when updates are required to be copyless as *streamable linear regular transductions* (SLR). Both these classes turn out to be closed under operations relevant for design of query languages. For instance, consider the `split` operation, the quantitative analog of the (unambiguous) language concatenation: given two partial transductions f and g over data words and a binary data operation op , `split`(f, g, op) is defined on an input string w if it can be split uniquely into two parts $w = w_1w_2$ such that both $f(w_1)$ and $g(w_2)$ are defined, and if so returns $op(f(w_1), g(w_2))$. Both classes SR and SLR are closed under this operation. It is worth noting that the class of general streaming algorithms is not closed under this operation (that is, it is possible that both f and g can be computed in a streaming fashion using memory that is logarithmic in the

size of the input string, but $\text{split}(f, g, op)$ requires memory that grows linearly with the length of the input string). This justifies the role of regularity in design of streaming algorithms in a modular fashion.

In section 3, we give a logical characterization of the streamable regular transductions. Recall that the class of regular transductions corresponds to MSO-definable string-to-tree transformations. We show that streamability corresponds to the requirement that in the output graph edges cannot go backwards: the output graph has multiple copies corresponding to an index position i in the input string, and the children of each such node must be copies corresponding to input positions less than or equal to i . Furthermore, allowing updates in a CRA to be copyful corresponds to allowing sharing in the output graph as long as no cycles are created. More precisely, the class **SR** corresponds to MSO-definable transformations from string graphs to directed acyclic graphs without backward edges, and the class **SLR** corresponds to MSO-definable transformations from string graphs to trees without backward edges. The proof builds on ideas developed in [5, 6, 7], but is self contained.

In section 4, we study the relationship with *weighted automata*, an extensively studied quantitative extension of automata [8]. A weighted automaton defines a (total) function from Σ^* to a set D equipped with two operations, product and sum, that form a semiring. Such an automaton is a nondeterministic finite-state machine whose each transition is labeled with a tag in Σ and a weight in D . The cost of a run of the automaton is the product of all the weights of the transitions in the run, and the cost associated with a string is the sum of the costs of all the corresponding runs of the automaton. If we choose the set of operations to contain only multiplication by constants, for each value in D , then we show that the two classes **SR** and **SLR** coincide, and correspond exactly to unambiguous weighted automata. If we choose the set of operations to contain multiplication by constants, for each value in D , and sum, then we show that the class **SR** corresponds exactly to transductions definable by weighted automata.

2. Cost Register Automata

Cost Register Automata (CRAs) were introduced in [4] with the goal of providing a machine-based characterization of the class of *regular transductions* (called *regular cost functions* in [4]) from strings to costs. To capture precisely the class of transductions considered in [4], it is sufficient to restrict the updates so that they are “copyless” (each variable is used at most once

on the right-hand side of a parallel update), but it is necessary to allow variables to hold terms with at most one “hole” (parameter that can be used for substitution). For the classes of functions that we consider here, we consider CRAs that can hold only values and the copylessness restriction is dropped to accommodate string-to-DAG transformations (defined later in section 3).

We will start by introducing the semantic objects of data transductions, which represent stream transformations, and several regular combinators on these objects. Then, we will present the formal computational model of Cost Register Automata (CRAs) and illustrate its use with several examples. We define the class of *streamable regular* (resp., *streamable linear regular*) transductions, denoted **SR** (resp., **SLR**), as the class of transductions that are computed by copyful (resp., copyless) unambiguous CRAs. The restriction of “copylessness” essentially says that it is not allowed to copy the contents of registers. A key feature of the CRA model is that it is parameterized by the operations on data values that are allowed when updating the registers.

We will also show that in the case of copyless CRAs, unambiguous machines are more expressive than their deterministic counterparts. It will then be established that both classes **SR** and **SLR** are closed under regular combinators that are straightforward analogs of the familiar regular combinators for word languages. We will conclude with a short discussion of the complexity of evaluation in the CRA model.

2.1. Data Transductions and Combinators

We fix a finite alphabet Σ whose elements are called *tags*, a data set D of *data values*. A *data word* is a sequence of tagged values, i.e. a word over the alphabet $\Sigma \times D$. A *data transduction* is a partial function of type

$$(\Sigma \times D)^* \rightarrow D.$$

For a data word $w \in (\Sigma \times D)^*$, we write $w|_{\Sigma}$ to denote the elementwise projection of w to the tag component. More formally,

$$(a_1, d_1) (a_2, d_2) \cdots (a_n, d_n) |_{\Sigma} = a_1 a_2 \cdots a_n.$$

We will only consider transductions f that satisfy the following property: for all input sequences u, v with $u|_{\Sigma} = v|_{\Sigma}$, $f(u)$ is defined iff $f(v)$ is defined. The *rate* of f is the language $R(f) \subseteq \Sigma^*$ defined as follows:

$$R(f) = \{\sigma \in \Sigma^* \mid \text{for every } \alpha \text{ and } w \text{ with } w|_{\Sigma} = \sigma, f(w) \text{ is defined}\}.$$

So, the domain of a transduction is equal to

$$\text{dom}(f) = \{w \in (\Sigma \times D)^* \mid w|_{\Sigma} \in \mathbf{R}(f)\}.$$

A transduction $f : (\Sigma \times D)^* \rightarrow D$ is said to be *value-oblivious* if for all input sequences u, v with $u|_{\Sigma} = v|_{\Sigma}$ it satisfies $f(u) = f(v)$. So, a value-oblivious transduction f can be represented by a function $\hat{f} : \Sigma^* \rightarrow D$ whose domain is equal to the range of f and which satisfies $f(w) = \hat{f}(w|_{\Sigma})$ for every input sequence w .

Now, we also fix a family \mathcal{O} of constants and operations on the data set D . We write \mathcal{O}_n to denote the set of n -ary operations that are contained in \mathcal{O} . In particular, \mathcal{O}_0 is the set of constants of \mathcal{O} . We will consider below several combinators for composing data transductions. Before we give the definitions of these combinators, we need to introduce some additional notation. The operation \odot is *unambiguous concatenation* of strings and $^{\circledast}$ is *unambiguous iteration*, which are defined as follows:

$$\begin{aligned} A \odot B &= \{w \mid \text{there are unique } u \in A \text{ and } v \in B \text{ with } w = uv\} \\ A^{\circledast} &= \{w \mid \text{there are unique } n \geq 0 \text{ and } w_1, \dots, w_n \in A \\ &\quad \text{with } w = w_1 \cdots w_n\} \end{aligned}$$

The *left fold* combinator $\text{fold} : (B \times A \rightarrow B) \times B \times A^*$ is defined as follows:

$$\begin{aligned} \text{fold}(f, b, \varepsilon) &= b \\ \text{fold}(f, b, a \cdot w) &= \text{fold}(f, f(b, a), w) \end{aligned}$$

for all $f : B \times A \rightarrow B$, $b \in B$, $a \in A$, and $w \in A^*$. For example,

$$\begin{aligned} \text{fold}(f, b, a_1) &= f(b, a_1) \\ \text{fold}(f, b, a_1 a_2) &= f(f(b, a_1), a_2) \\ \text{fold}(f, b, a_1 a_2 a_3) &= f(f(f(b, a_1), a_2), a_3) \end{aligned}$$

Table 1 contains the typing rules and the definitions of several combinators for data transductions: output combination, choice (**else**), quantitative concatenation (**split**), and quantitative iteration (**iter**). These combinators are quantitative analogs of the familiar combinators for string languages: intersection, union, concatenation, and iteration. These regular operations on data transductions are relevant for the design of query languages for stream

processing, since they enable the modular description of complex streaming computations.

The work [1] introduces a family of regular combinators that captures a class of transductions that is defined in terms of CRAs whose registers use terms with variables instead of values. The StreamQRE language of [2] extends the combinators of Table 1 with additional constructs (such as streaming composition and key-based partitioning) that are useful in practical stream processing applications.

2.2. Syntax and Semantics of CRAs

A Cost Register Automaton (CRA) is a machine that maps data words, i.e. strings over an input alphabet of tagged values, to output values. It uses a finite-state control and a finite set of registers that contain values. At each step, the machine reads an input tag and an input value, updates its control state, and updates its registers using a parallel assignment. The definition of the machine is parameterized by the set of expressions that can be used in the assignments. For example, a CRA with increments can use multiple registers to compute alternative costs, perform updates of the form $x := y + c$, where x and y are registers and c is a constant, at each step, and commit to the cost computed in one of the registers at the end.

Let X be a set of variables, and suppose that \mathcal{O} is the family of available constants and operations on the data set D . The set of *expressions* over X , denoted $\mathbb{E}_{\mathcal{O}}[X]$ is defined by the following rules:

$$\frac{\text{constant } c \in \mathcal{O}}{c : \mathbb{E}_{\mathcal{O}}[X]} \quad \frac{\text{variable } x \in X}{x : \mathbb{E}_{\mathcal{O}}[X]}$$

$$\frac{n\text{-ary operation } op \in \mathcal{O} \quad t_i : \mathbb{E}_{\mathcal{O}}[X] \text{ for all } i = 1, \dots, n}{op(t_1, \dots, t_n) : \mathbb{E}_{\mathcal{O}}[X]}$$

A function $\alpha : X \rightarrow D$ is called a *variable assignment*, and it extends uniquely to a homomorphism (i.e., operation-preserving mapping) $\hat{\alpha} : \mathbb{E}_{\mathcal{O}}[X] \rightarrow D$. An expression $t \in \mathbb{E}_{\mathcal{O}}[X]$ denotes a function $\llbracket t \rrbracket : D^X \rightarrow D$, defined as follows: for a variable assignment $\alpha : X \rightarrow D$, we put $\llbracket t \rrbracket(\alpha) = \hat{\alpha}(t)$. That is, $\llbracket t \rrbracket(\alpha)$ is the result of evaluating the variable-free expression that is obtained from t by replacing each variable x by $\alpha(x)$.

In the development that follows, we will also consider a special symbol **val**, which should be assumed to be always distinct from all variables in X . The symbol **val** is meant to be used as a reference to the value of the current data

$$\frac{op : D^n \rightarrow D \quad f_i : (\Sigma \times D)^* \rightarrow D \text{ for all } i}{op(f_1, \dots, f_n) : (\Sigma \times D)^* \rightarrow D} \text{ (output combination)}$$

$$\mathbf{R}(op(f_1, \dots, f_n)) = \mathbf{R}(f_1) \cap \dots \cap \mathbf{R}(f_n)$$

$$op(f_1, \dots, f_n)(w) = op(f_1(w), \dots, f_n(w))$$

$$\frac{f, g : (\Sigma \times D)^* \rightarrow D}{f \text{ else } g : (\Sigma \times D)^* \rightarrow D} \text{ (choice)}$$

$$\mathbf{R}(f \text{ else } g) = \mathbf{R}(f) \cup \mathbf{R}(g)$$

$$(f \text{ else } g)(w) = \begin{cases} f(w), & \text{if } w|_\Sigma \in \mathbf{R}(f) \\ g(w), & \text{otherwise} \end{cases}$$

$$\frac{f, g : (\Sigma \times D)^* \rightarrow D \quad op : D \times D \rightarrow D}{\text{split}(f, g, op) : (\Sigma \times D)^* \rightarrow D} \text{ (quantitative concatenation)}$$

$$\mathbf{R}(\text{split}(f, g, op)) = \mathbf{R}(f) \odot \mathbf{R}(g)$$

$$\text{split}(f, g, op)(uv) = op(f(u), g(v)), \text{ if } u|_\Sigma \in \mathbf{R}(f) \text{ and } v|_\Sigma \in \mathbf{R}(g)$$

$$\frac{f : (\Sigma \times D)^* \rightarrow D \quad c \in D \quad op : D \times D \rightarrow D}{\text{iter}(f, c, op) : (\Sigma \times D)^* \rightarrow D} \text{ (quantitative iteration)}$$

$$\mathbf{R}(\text{iter}(f, c, op)) = \mathbf{R}(f)^{\otimes}$$

$$\text{iter}(f, c, op)(w_1 \cdots w_n) = \text{fold}(op, c, [f(w_1), \dots, f(w_n)]),$$

$$\text{if } w_i|_\Sigma \in \mathbf{R}(f) \text{ for all } i = 1, \dots, n$$

Table 1: Some regular combinators for data transductions: output combination, choice, quantitative concatenation, and quantitative iteration.

value. An expression $t \in \mathbb{E}_{\mathcal{O}}[X \cup \{\text{val}\}]$ that potentially contains occurrences of val denotes a function $\llbracket t \rrbracket : D^X \times D \rightarrow D$, given by

$$\llbracket t \rrbracket(\alpha, d) = \hat{\beta}(t) \text{ where } \beta = \alpha[\text{val} \mapsto d] : X \cup \{\text{val}\} \rightarrow D$$

for every variable assignment $\alpha : X \rightarrow D$ and every value $d \in D$. We use the notation $\alpha[\text{val} \mapsto d]$ to mean the extension of α that maps val to d .

Definition 1 (CRA). A (nondeterministic, copyful) *Cost Register Automaton* (CRA) over the tag alphabet Σ , data values D , and data operations \mathcal{O} is a tuple

$$\mathcal{A} = (Q, X, \Delta, I, F),$$

where

- Q is a finite set of states
- X is a finite set of registers
- $\Delta \subseteq Q \times \Sigma \times U_{\mathcal{O}} \times Q$ is the set of transitions with $U_{\mathcal{O}}$ the set of register updates $X \rightarrow \mathbb{E}_{\mathcal{O}}[X \cup \{\text{val}\}]$
- $I : Q \rightarrow (X \rightarrow \mathcal{O}_0)$ is the initialization function
- $F : Q \rightarrow \mathbb{E}_{\mathcal{O}}[X]$ is the finalization function

The domain of I is the set of *initial states*, and the domain of F is the set of *final* or *accepting states*. A CRA is said to be *unambiguous* (resp., deterministic) if the underlying NFA is unambiguous (resp., deterministic). We use the abbreviations NCRA, UCRA, and DCRA to indicate that an automaton is nondeterministic, unambiguous, and deterministic respectively.

A CRA is said to be *copyless* if it satisfies the following restrictions: (1) for every transition $(p, a, \theta, q) \in \Delta$ and every register x , there is at most one occurrence of x in the list of expressions $\theta(x_1), \dots, \theta(x_n)$, where x_1, \dots, x_n is an enumeration of X , and (2) for every final state q and every register x , there is at most one occurrence of x in the expression $F(q)$.

A CRA is said to be *trim* if it satisfies the properties: (1) every state of the automaton is reachable from an initial state, and (2) some final state is reachable from every state.

Semantics of CRAs. A CRA computes like a classical finite-state automaton, but the configuration consists of both a current state and a register assignment $X \rightarrow D$. A transition specifies the next state, as well as the update of the variables using their current values and the current data value from the input. For an input sequence $w = (a_1, d_1) (a_2, d_2) \cdots (a_n, d_n)$ in $(\Sigma \times D)^*$, we define a w -run in \mathcal{A} to be a sequence

$$(q_0, \alpha_0) \xrightarrow{(a_1, d_1)} (q_1, \alpha_1) \xrightarrow{(a_2, d_2)} (q_2, \alpha_2) \xrightarrow{(a_3, d_3)} \cdots \xrightarrow{(a_n, d_n)} (q_n, \alpha_n)$$

with $q_i \in Q$ and $\alpha_i : X \rightarrow D$ for all i so that the following hold:

1. *Initialization:* q_0 is an initial state and $\alpha_0 = I(q_0)$.
2. *Transition:* for every transition $(p, \alpha) \xrightarrow{(a, d)} (q, \beta)$ of the run there is a register update function $\theta \in U_{\mathcal{O}}$ with $(p, a, \theta, q) \in \Delta$ s.t.

$$\beta(x) = \llbracket \theta(x) \rrbracket(\alpha, d) \text{ for every register } x \in X.$$

3. *Finalization:* q_n is a final state.

The *value* of the run is $\llbracket F(q_n) \rrbracket(\alpha_n)$. A nondeterministic CRA implements a multi-valued transduction $(\Sigma \times D)^* \rightarrow \mathcal{P}(D)$, where \mathcal{P} is the powerset operator. The value of this transduction on an input sequence w is the set of values of all w -runs. An unambiguous CRA has at most one run for each input sequence, and therefore implements a transduction $(\Sigma \times D)^* \rightarrow D$.

Definition 2 (Streamable Regular Transductions). Fix a tag alphabet Σ , a set D of values, and a family \mathcal{O} of operations over D . We define the class of *streamable regular transductions* (over \mathcal{O}), denoted $\text{SR}(\mathcal{O})$, to be the class of data transductions that can be computed by some UCRA (over \mathcal{O}). Similarly, we define the class of *streamable linear regular transductions*, denoted $\text{SLR}(\mathcal{O})$, to be the class of data transductions that can be computed by some copyless UCRA. Table 2 summarizes these definitions.

2.3. Examples


We present in this subsection several examples that illustrate how CRAs compute. We consider both deterministic and unambiguous variants, and we give examples of copyless and copyful (i.e., unrestricted) register updates.

CRA	Class of transductions
copyless over \mathcal{O}	SLR(\mathcal{O}): Streamable Linear Regular
copyful over \mathcal{O}	SR(\mathcal{O}): Streamable Regular

Table 2: Classes of transductions computed by unambiguous CRAs.

Whenever convenient, we will allow the use of ε -transitions for the unambiguous variants. This is w.l.o.g. because ε -transitions can be eliminated with a straightforward variant of the ε -elimination procedure for classical finite-state automata (assuming there are no ε -cycles).

Example 3 (Copyless DCRA). Suppose the tag alphabet is $\Sigma = \{a, b\}$, the type of values is $D = \mathbb{N}$ (the set of natural numbers), and \mathcal{O} consists of the constant 0 and the binary addition operation. Consider the following transduction $f : (\Sigma \times D)^* \rightarrow D$, which is defined on all nonempty sequences. If a sequence ends with an a -labeled value, then f outputs the sum of all a -labeled values in the sequence. Similarly, if a sequence ends with a b -labeled value, then f outputs the sum of all b -labeled values in the sequence. Then f is implemented by the following copyless CRA:




where the initialization θ and the register updates θ_a, θ_b are defined as follows:

$$\theta = \begin{cases} x := 0 \\ y := 0 \end{cases} \quad \theta_a = \begin{cases} x := x + \text{val} \\ y := y \end{cases} \quad \theta_b = \begin{cases} x := x \\ y := y + \text{val} \end{cases}$$

The register x holds the sum of all a -labeled values seen so far, and the register y holds the sum of all b -labeled values seen so far. The state p is active only upon initialization, and q_a (resp., q_b) is active when the input sequence ends with an a -labeled (resp., b -labeled) value. The CRA is deterministic.


Moreover, it is copyless, since every register is used at most once in the right-hand side of the updates θ_a and θ_b .

Example 4 (Copyless UCRA). In Example 3, it was shown how to implement the transduction f using a copyless DCRA with two registers. We will show now how to implement the same transduction using a copyless UCRA with one register. Informally, the idea is to employ unambiguous nondeterminism in order to guess whether the input data word will end with the tag a or the tag b . We assume additionally that ε -transitions are allowed.



The register x holds the sum of all a -labeled values at the states p_a and q_a , and the sum of all b -labeled values at the states p_b and q_b . The automaton is copyless, since the register x is used at most once in the right-hand side of the updates.

Example 5. Suppose the tag alphabet is $\Sigma = \{a, \#\}$, the type of values is $D = \mathbb{N}$ (the set of natural numbers), and \mathcal{O} consists of the constant 0, the binary addition operation, and \max . Consider the following transduction $f : (\Sigma \times D)^* \rightarrow D$, whose rate is given by the regex $(a^+ \cdot \#)^*$. A maximal subsequence of the input that is of the form $aa \dots a\#$ is called a *block*. The transduction f outputs that maximum cost over all input blocks, where the cost of a block is the sum of the a -labeled values. This is implemented by the following two-register copyless CRA:




where the initialization θ and the register updates $\theta_a, \theta_{\#}$ are given by:

$$\theta = \begin{cases} x := 0 \\ y := 0 \end{cases} \quad \theta_a = \begin{cases} x := x + \text{val} \\ y := y \end{cases} \quad \theta_{\#} = \begin{cases} x := 0 \\ y := \max(y, x) \end{cases}$$

The register x holds the sum of all a -labeled values in the current block, and the register y holds the maximum cost over all complete blocks seen so far. The CRA is deterministic and copyless (every register is used at most once in the right-hand side of the updates θ_a and $\theta_{\#}$).

Example 6 (Copyful UCRA). Suppose the tag alphabet is $\Sigma = \{a, b\}$, the type of values is $D = \mathbb{N}$ (the set of nonnegative integers), and \mathcal{O} consists of the constant 0 and the binary operations \max and \ominus (which is defined as $x \ominus y = \max(x - y, 0)$ for all $x, y \in \mathbb{N}$). Consider the following transduction $f : (\Sigma \times D)^* \rightarrow D$, which is defined on sequences that contain at least one a -labeled value. For these sequences, the value of f is the *maximum drawdown* (the maximum loss from a peak to a trough) in the input signal after the last occurrence of a b -labeled value. The transduction f is implemented by the following copyful CRA:




where the initialization θ_0 and the register update θ are defined as follows:

$$\theta_0 = \begin{cases} x := 0 \\ y := 0 \end{cases} \quad \theta = \begin{cases} x := \max(x, \text{val}) \\ y := \max(y, \max(x, \text{val}) \ominus \text{val}) \end{cases}$$


The numerical computation on the values concerns the part of the input after the last occurrence of a b -labeled value. The register x holds the maximum value, and the register y holds the maximum drawdown. The state p is always active, and q is active when at least one b -labeled value is seen in the input. The CRA is unambiguous. Moreover, it is not copyless, since the register x is used twice in the right-hand sides of the register update θ .

Example 7. Suppose the tag alphabet is $\Sigma = \{a\}$, the type of values is $D = \mathbb{Q}$ (the set of rational numbers), and \mathcal{O} consists of the constant 0 and

the binary operation op given by $op(x, y) = \lambda \cdot x + y$, where λ is a fixed rational constant in the open interval $(0, 1)$. The transduction f maps a data word w with $w|_D = d_1 d_2 \dots d_n \in D^*$ to $\lambda^{n-1} \cdot d_1 + \dots + \lambda \cdot d_{n-1} + d_n$. We write $w|_D$ to denote the elementwise projection of the data word to the value component. The transduction can be computed by the following CRA:



Define now the transduction g as $g(w_1 w_2 \dots w_n) = f(w_n \dots w_2 w_1)$. The transduction g cannot be computed by a CRA (over \mathcal{O}), because the registers can only hold values. However, if the registers were allowed to hold terms with a hole \square (a placeholder that can be substituted later in the computation), then g could be expressed.



In the machine shown above the register is meant to hold a term (with potential occurrences of \square), and the operation $x[t/\square]$ denotes the result of substituting t for \square in the contents of x . For example, if the input is $w = (a, d_1) (a, d_2) (a, d_3)$, then the successive contents of the register x (after 3 steps of computation) are shown below:

$$\square \quad op(\square, d_1) \quad op(op(\square, d_2), d_1) \quad op(op(op(\square, d_3), d_2), d_1)$$

So, the ability to use the holes \square and substitute terms into it provides additional computation power. The disadvantage of using terms is that in some cases their size can grow linearly in the size of the input.


2.4. Choice of Operations on Data Values

The definition of CRAs is parameterized by the choice of operations that are allowed when updating the registers. We will see now that seemingly inconsequential variations of the available operations can change what can be computed by a CRA. Suppose that $(D, \cdot, 1)$ is a monoid, i.e. \cdot is an associative binary operation on D and 1 is a left and right identity for the \cdot operation. Define \mathcal{O} to be the family of operations that contains 1 and \cdot , and \mathcal{O}' to contain the identity 1 and the unary operation $(- \cdot d)$ for every $d \in D$.

Lemma 8. The class of transductions $\text{SR}(\mathcal{O}')$ is equal to $\text{SLR}(\mathcal{O}')$.

Proof. It suffices to see that a copyful UCRA over \mathcal{O}' can be simulated by a copyless UCRA over \mathcal{O}' . The key observation is that in a copyful UCRA at most one register can contribute to the final output. So, at every step of the computation the machine can guess the register x that will contribute to the output (using unambiguous nondeterminism) and proceed by updating only x . Notice that this construction produces a UCRA with a single register. \square

In general, the class of transductions $\text{SR}(\mathcal{O})$ may be strictly larger than $\text{SLR}(\mathcal{O})$. For example, consider the case where $(D, \cdot, 1)$ is the free monoid generated by the alphabet Σ . That is, D is the set of finite words over Σ , \cdot is word concatenation, and 1 is the empty word. For all transductions on $\text{SLR}(\mathcal{O})$, the output is of size linear in the size of the input. This property does not hold, however, for the transductions of $\text{SR}(\mathcal{O})$. The copyful single-state single-register CRA



emits outputs that are of size exponential in the input. Lemma 8 and the preceding discussion establish that the precise choice of \mathcal{O} can affect the class of functions that can be computed.

Another important point about the set \mathcal{O} of operations is that it can incorporate tests on values. The definition of CRAs does not allow tests on registers as guards on the transitions, which means that the control of the automaton depends only on the observed sequence of input tags. It is possible, however, to compute a transduction such as

$$f(w) = \begin{cases} \text{sum of } a\text{-tagged values,} & \text{if } |w|_a = |w|_b \\ \text{sum of } b\text{-tagged values,} & \text{if } |w|_a \neq |w|_b \end{cases}$$

(where $|w|_a$ is the number of occurrences of the tag a in the data word w) by including in \mathcal{O} an *if-then-else* operation such as:

$$\text{ITE}(x, y, z, w) = \begin{cases} z, & \text{if } x = y \\ w, & \text{if } x \neq y \end{cases}$$

2.5. Unambiguity Versus Determinism

In this section we explore the relationship between unambiguous and deterministic CRAs. The first result is a kind of subset construction for transforming a copyful UCRA to an equivalent copyful DCRA. This means that unambiguous nondeterminism is not essential when the automata are allowed to copy values. In the case of copyless CRAs, however, we will see that unambiguity is essential. That is, there are data transductions that can be computed by copyless UCRA's but not by copyless DCRA's.

To get an intuitive understanding of why copyless DCRA's are weaker, suppose that a copyless DCRA has a register x which can contribute to the output in two different ways, depending on some regular property of the future input. Because the automaton is not allowed to copy the value of x , it cannot commit to any of the two different eventualities in order to update x appropriately. If this situation can only arise a fixed number of times in the computation, a solution can be given by blowing up the number of registers: create from the outset two copies of the register x , one for each possible eventuality and update them separately. We will see, however, that there are examples where this situation can arise an arbitrary number of times. Such examples witness the power of unambiguity in the model of copyless CRAs.

Theorem 9 (Copyful UCRA to Copyful DCRA). Every copyful unambiguous CRA over \mathcal{O} is equivalent to a copyful deterministic CRA over \mathcal{O} (assuming \mathcal{O} contains at least one constant).

Proof. Let $\mathcal{A} = (Q, X, \Delta, I, F)$ be an arbitrary copyful unambiguous CRA. W.l.o.g. we assume that \mathcal{A} is trim. We define the deterministic CRA \mathcal{B} using a modified subset construction. The automaton \mathcal{B} has registers $Q \times X$ and its state space is equal to

$$\{\Delta(\text{dom}(I), w) \mid w \in \Sigma^*\},$$

where $\Delta(P, w)$ for $P \subseteq Q$ is the set of states reachable from a state in P via a path that is labeled with the word w .

The initial state of \mathcal{B} is equal to the set of initial states of \mathcal{A} , and a register (q, x) of \mathcal{B} is initialized to $I(q)(x)$ if q is initial. The initialization of the rest of the registers can be chosen arbitrarily.

A state P of \mathcal{B} is final if P contains a final state q of \mathcal{A} . Since \mathcal{A} is unambiguous, each state of \mathcal{B} contains at most one final state of \mathcal{A} . The

finalization term for P is defined to be the term that results from $F(q)$ by substituting (q, x) for every register x .

Consider now a state P of \mathcal{B} and a tag $a \in \Sigma$. Then, $R = \Delta(P, a) = \bigcup_{p \in P} \Delta(p, a)$ is also a state of \mathcal{B} . We include a transition $P \xrightarrow{a} R$ in \mathcal{B} . It remains to specify the register update function η for this transition. Since \mathcal{A} is unambiguous and trim, for every $q \in R$ there is a unique $p \in P$ with $q \in \Delta(p, a)$. For a register (q, x) of \mathcal{B} we put

$$\eta(q, x) = \text{the term resulting from } t \text{ by replacing every register } y \text{ by } (p, y),$$

where $t = \theta(x)$ and θ is the register update function of the unique transition $p \xrightarrow{a, \theta} q$ in the automaton \mathcal{A} . The update for registers not covered by the above description can be set arbitrarily.

In order to show that \mathcal{B} is equivalent to \mathcal{A} it suffices to establish that for a given input sequence $w \in (\Sigma \times D)^*$, the configurations $\{(q, \alpha_q) \mid q \in X\}$, that are w -reachable from some initial state of \mathcal{A} are encoded in the configuration (X, β) of \mathcal{B} that results after consuming w : $\beta(q, x) = \alpha_q(x)$ for all $q \in X$. \square

Theorem 9 establishes the expressive equivalence of copyful UCRA's and copyful DCRA's. Now, we will show (Theorem 13) that copyless UCRA's are strictly more expressive than copyless DCRA's (for a fixed family of data operations). To establish this separation result, it suffices to focus on CRA's that do not make use of the input data values in their computations. We say that a CRA is *value-oblivious* if its register update function contains no occurrence of the `val` symbol. In this case the denotation of the CRA is a value-oblivious transduction, since the only the input tags are used and the input data values do not affect the computation. For the rest of this subsection, we will be dealing exclusively with value-oblivious CRA's. For this reason, a data transduction is taken here to be a function of type $\Sigma^* \rightarrow D$, and a CRA is implicitly assumed here to not contain any occurrence of the `val` symbol. For the special case where $D = \Gamma^*$ for some finite alphabet Γ , a transduction $\Sigma^* \rightarrow D$ is also called a *string transduction* (the term *word transduction* is often used in the relevant literature).

As an example of this, consider the following string transduction $f : \Sigma^* \rightarrow \Sigma^*$ with domain $\text{dom}(f) = \{a, b\}^+\#$, where the alphabet is $\Sigma = \{a, b, \#\}$:

$$f(u\#) = \begin{cases} a^i\#, & \text{if } |u| = i \text{ and } u \text{ ends with } a \\ b^i\#, & \text{if } |u| = i \text{ and } u \text{ ends with } b \end{cases}$$

for all $u \in \{a, b\}^+$. We assume that the only data operation that is allowed is to concatenate string constants to the registers. An unambiguous CRA can compute this transduction with a single register by guessing the letter that will precede the $\#$ symbol. A deterministic CRA, on the other hand, requires at least two registers: one register for the case where the input ends in $a\#$, and one register for the case where the input ends in $b\#$.

This situation extends to the case where the computation of f is iterated several times. That is, let us define the string transduction $f^k : \Sigma^* \rightarrow \Sigma^*$ that iterates f sequentially k times. The domain of f^k is equal to $\text{dom}(f^k) = (\{a, b\}^+ \#)^k$, and f^k is given by

$$f^k(u_1\# \cdots u_k\#) = f(u_1\#) \cdots f(u_k\#)$$

for all $u_1, \dots, u_k \in \{a, b\}^+$. In order to compute f^k , an unambiguous CRA requires just one register: for every block of the input (where a block is a sequence of letters that ends with a $\#$ symbol), the automaton guesses the last letter of the block and extends the unique register with the appropriate letter. A copyless DCRA, however, requires at least 2^k registers: a separate register is needed for each possible combination of ending letters for the k input blocks. For example, if $k = 2$ then there are $2^k = 4$ different cases for the form of the input strings in the domain of f^k :

$$\dots a\# \dots a\# \quad \dots a\# \dots b\# \quad \dots b\# \dots a\# \quad \dots b\# \dots b\#$$

A copyless DCRA can only compute f^k by using one separate register for each one of these 2^k possibilities.

In the previous paragraph, we discussed how the iteration of f a fixed number k of times creates the need for 2^k registers in the model of copyless DCRA. An immediate consequence of that is that in this model, the iteration of f a finite (but arbitrary) number of times cannot be computed with a finite number of registers. This establishes the result that UCRA are more expressive than DCRA in the copyless case. Establishing this result formally is rather involved and requires a sequence of constructions on CRA:

1. Every CRA over *unary* data operations is equivalent to the disjoint union of single-register CRA.
2. A copyless DCRA can be modified in a way that suppresses a prefix of the output.

3. If a string transduction g requires at least k registers to be computed by a copyless DCRA, then the string transduction that computes f (defined previously) and then g in sequence requires at least $2k$ registers in this model.

Lemma 12 is the heart of the argument. It establishes that the computation of the sequential iteration f^k (defined previously) requires at least 2^k registers in the model of copyless DCRA.

We start the technical development with Lemma 10 below, which describes a construction for simplifying CRAs that use only unary operation on the values. For this special case, it is shown that the register updates are w.l.o.g. of the form $x := \delta(c)$ or $x := \delta(x)$, where x is a register, c is a constant, and δ is a unary operation. In other words, the CRA can be decomposed into the disjoint union of several single-register CRAs.

Lemma 10 (Separable Updates). Suppose that the collection of operations \mathcal{O} consists only of unary operations and constants. Then, every copyless NCRA (resp., UCRA/DCRA) over \mathcal{O} is equivalent to some copyless NCRA (resp., UCRA/DCRA) over \mathcal{O} (with the same set of registers) whose register updates are of one of the following forms: $x := \delta(c)$ or $x := \delta(x)$, where c is a constant and δ is the composition of unary operations.

Proof. We use the symbol δ in order to indicate an arbitrary composition of unary operation symbols. In a CRA $\mathcal{A} = (Q, X, \Delta, I, F)$ over \mathcal{O} the register updates can be of any of the forms: $x := \delta(c)$, or $x := \delta(x)$, or $x := \delta(y)$ for $x \neq y$. We will employ a product construction in order to eliminate updates of the form $x := \delta(y)$ with $x \neq y$. If the assignment $x := \delta(y)$ is part of a parallel update, then y cannot flow into another register due to the copyless restriction.

We will describe now the construction of the CRA \mathcal{B} , which has the same registers as \mathcal{A} . The idea is that \mathcal{B} computes like \mathcal{A} , but also maintains some extra information in the finite control that allows to “rename” the registers. A renaming function σ is a bijection on the set X of registers, and the state space of \mathcal{B} is the cartesian product of the state space of \mathcal{A} with the function space of bijections $X \rightarrow X$. The intuition is that when the computation is in a state (q, σ) of \mathcal{B} , the value that \mathcal{A} would hold in the register x is contained in the register $\sigma(x)$ of \mathcal{B} .

If q is an initial state of \mathcal{A} , then (q, id_X) is an initial state of \mathcal{B} with the same initialization. The renaming function id_X is the identity, which

essentially means that the registers are not renamed.

If q is a final state of \mathcal{A} , then (q, σ) is a final state of \mathcal{B} . The finalization term for (q, σ) in \mathcal{B} is defined to be equal to $\sigma(F(q))$.

Since \mathcal{A} is copyless, every register update function θ induces *at least one* bijection $v : X \rightarrow X$, where $v(x)$ is the register that appears on the right-hand side of the assignment to x (if such a register exists). If no register is used (i.e., x is set to $\delta(c)$ for a constant c), then $v(x)$ can be set arbitrarily to one of the registers that does not appear in the right-hand side of θ at all. Thus, there will be multiple possible bijections v , and we choose one arbitrarily. For a transition from state p to state q in \mathcal{A} with label λ and register update θ , we put for all renaming functions σ a transition from (p, σ) to (q, τ) in \mathcal{B} with label λ and register update η , where:

$$\tau = v; \sigma \qquad \eta(x) = \sigma(\theta(\tau^{-1}(x)))$$

and $;$ is function composition in diagrammatic order. In other words, if $x := \delta(y)$ is a register update of \mathcal{A} , then we put $\tau(x) := \delta(\sigma(y))$ for the corresponding register update in \mathcal{B} . Moreover, notice that $v(x) = y$ and therefore $\tau(x) = (v; \sigma)(x) = \sigma(v(x)) = \sigma(y)$. So, the update in \mathcal{B} has the same register appearing in both the left-hand and right-hand side.

The main claim is that the possible runs in \mathcal{B} correspond to the runs of \mathcal{A} modulo the variable renaming described by the renaming functions. \square

For an alphabet Σ that contains the symbol $\#$, we define the *prefix removal* operation $\partial_{\#} : \Sigma^* \rightarrow \Sigma^*$ as follows:

$$\partial_{\#}(w) = \begin{cases} \varepsilon, & \text{if } \# \text{ does not appear in } w; \\ v, & \text{if } w = u\#v. \end{cases}$$

That is, $\partial_{\#}$ removes the prefix of the string until the first occurrence of a $\#$.

Suppose now that Σ and Γ are finite alphabets, and that Γ contains the $\#$ symbol. For a string transduction of type $\Sigma^* \rightarrow \Gamma^*$, we define the *output prefix removal* operation $\partial_{\#}$ as follows:

$$\partial_{\#}(f)(w) = \begin{cases} \text{undefined,} & \text{if } f(w) \text{ is undefined;} \\ \partial_{\#}(f(w)), & \text{if } f(w) \text{ is defined.} \end{cases}$$

We show below that there is a construction on CRAs that corresponds to the output prefix removal operation on string transductions.

Lemma 11 (Remove Output Prefix). Suppose that Σ and Γ are tag alphabets with $\# \in \Gamma$, Γ^* is the set of values, and \mathcal{O} consists of the constant ε and the unary operations $(- \cdot a)$ for every $a \in \Gamma$. If the string transduction $f : \Sigma^* \rightarrow \Gamma^*$ is implementable by a UCRA (resp., DCRA) over \mathcal{O} , then $\partial_{\#}(f)$ is implementable by a UCRA (resp. DCRA) with the same number of registers.

Proof. Suppose that the UCRA $\mathcal{A} = (Q, X, \Delta, I, F)$ implements f . We will describe a UCRA \mathcal{B} that implements $\partial_{\#}(f)$. The idea is to employ a product construction that records some additional information in the finite control: whether the first $\#$ symbol would have already been added to a register in the execution of \mathcal{A} or not. Additionally, we modify the updates of \mathcal{A} so that no tags are appended to registers until after the first $\#$ would have been appended in the execution of \mathcal{A} .

The state space of \mathcal{B} is $Q \times (X \rightarrow \{0, 1\})$. A run of \mathcal{A} that ends in a configuration $(q, \alpha) \in Q \times (X \rightarrow \Gamma^*)$ corresponds to a run of \mathcal{B} that ends in a configuration $(q, \rho, \beta) \in Q \times (X \rightarrow \{0, 1\}) \times (X \rightarrow \Gamma^*)$ that satisfies

$$\beta(x) = \partial_{\#}(\alpha(x)) \quad \text{and} \quad \rho(x) = \begin{cases} 0, & \text{if } \alpha(x) \text{ does not contain } \# \\ 1, & \text{if } \alpha(x) \text{ contains } \# \end{cases}$$

for every register x . Notice that the construction that is outlined here preserves determinism. \square

Lemma 12. Suppose that $\Sigma = \{a, b, \#\}$ and Γ is the tag alphabet, Γ^* is the set of values, and \mathcal{O} consists of the constant ε and the unary operations $(- \cdot a)$ and $(- \cdot b)$. Let $f : \Sigma^* \rightarrow \Sigma^*$ be a string transduction, and define the string transduction $g : \Sigma^* \rightarrow \Sigma^*$ as follows:

$$\begin{aligned} \text{dom}(g) &= \{a, b\}^+ \cdot \# \cdot \text{dom}(f) \\ g(u\#v) &= \begin{cases} a^i \cdot \# \cdot f(v) \text{ where } i = |u|, & \text{if } u \text{ ends with } a \\ b^i \cdot \# \cdot f(v) \text{ where } i = |u|, & \text{if } u \text{ ends with } b \end{cases} \end{aligned}$$

for all $u \in \{a, b\}^+$ and $v \in \text{dom}(f)$. If the transduction g can be implemented by a copyless DCRA with k registers, then f can be implemented by a copyless DCRA with at most $k/2$ registers.

Proof. Let \mathcal{A} be a DCRA with k registers that implements g . By Lemma 10, we can assume without loss of generality that every register update of \mathcal{A} is

of the form $x := u$ or $x := x \cdot u$ for some $u \in \Sigma^*$. Consider all register updates that are enabled when the automaton \mathcal{A} consumes the input sequences $\{a, b\}^*$. Let R_a (resp., R_b) be the set of registers that are eventually only contain and are extended with a tags (resp., b tags) in these updates. The registers that are not among R_a and R_b cannot contribute to the output of the computation for sufficiently long input prefixes before the first occurrence of $\#$. One of R_a, R_b is of cardinality at most $k/2$. W.l.o.g. we can assume that $|R_a|$ is at most $k/2$. Then, consider the product \mathcal{B} of \mathcal{A} with an automaton that accepts $a^+ \cdot \# \cdot \text{dom}(f)$. In the automaton \mathcal{B} , only the registers of R_a contribute nontrivially to the output, so we can assume that R_a is the register set of \mathcal{B} . Notice that \mathcal{B} computes the transduction $h : \Sigma^* \rightarrow \Sigma^*$, given by:

$$\text{dom}(h) = a^+ \cdot \# \cdot \text{dom}(f) \qquad h(a^i \# v) = a^i \cdot \# \cdot f(v)$$

for all $i \geq 1$ and $v \in \text{dom}(f)$. Now, Lemma 11 implies that a DCRA \mathcal{C} with at most $k/2$ registers can compute the transduction $k : \Sigma^* \rightarrow \Sigma^*$, given by:

$$\text{dom}(k) = a^+ \cdot \# \cdot \text{dom}(f) \qquad k(a^i \# v) = f(v)$$

for all $i \geq 1$ and $v \in \text{dom}(f)$. Let (q, α) be the configuration of \mathcal{C} after consuming $a\#$. The DCRA that results from \mathcal{C} by setting the start state to be q and the initialization function to α computes the transduction f and has at most $k/2$ registers. \square

Theorem 13. There is a string transduction that can be computed by a copyless UCRA but not by a copyless DCRA.

Proof. Suppose that $\Sigma = \{a, b, \#\}$ and $f : \Sigma^* \rightarrow \Sigma^*$ is the string transduction with domain $\text{dom}(f) = \{a, b\}^+ \cdot \#$ that is defined as follows:

$$f(u\#) = \begin{cases} a^i \#, & \text{if } |u| = i \text{ and } u \text{ ends with } a \\ b^i \#, & \text{if } |u| = i \text{ and } u \text{ ends with } b \end{cases}$$

for all $u \in \{a, b\}^+$. Any copyless DCRA that computes f needs at least two registers. Indeed, if there is only one register, then on any input string it either contains an a or it contains a b or it contains neither; if we take the input string to be sufficiently large and then read in a b or an a , there is not enough information stored in the state to know the value of $|u| = i$

and output b^i or a^i . Define f^n to be the string transduction with domain $\text{dom}(f^n) = (\{a, b\}^+ \cdot \#)^n$ and

$$f^n(u_1\#u_2\#\cdots u_n\#) = f(u_1\#)f(u_2\#)\cdots f(u_n\#)$$

where every u_i is an element of $\{a, b\}^+$. A consequence of Lemma 12 is that any copyless DCRA that computes f^n requires at least 2^n registers.

Now, consider the string transduction f^* with domain $(\{a, b\}^+ \cdot \#)^*$ defined as $f^*(u_1\#\cdots u_n\#) = f^n(u_1\#\cdots u_n\#)$. The transduction f^* can be computed by a copyless UCRA that guesses the last letter of each block (maximal segment of letters that ends with a $\#$ symbol).

Assume for the sake of contradiction that there exists a copyless DCRA \mathcal{A} with $m \geq 1$ registers that computes f^* . The product of \mathcal{A} with a DFA that accepts the language $(\{a, b\}^+ \cdot \#)^k$ has m registers and computes the transduction f^k . But we have already discussed that such a DCRA would need at least 2^k registers, hence taking k large enough that $2^k > m$ gives us the desired contradiction. \square

2.6. Closure Under Regular Combinators

We will now show both classes **SR** (streamable regular) and **SLR** (streamable linear regular) of transductions are closed under operations that are relevant for the design of streaming query languages. In particular, we already presented a collection of some regular combinators for data transductions, in Table 1.

Theorem 14. The classes of transductions **SR** and **SLR** (see Definition 2 and Table 2) are closed under the regular combinators of Table 1.

Proof. The proof of the theorem relies on constructions that are variants of well-known constructions on classical finite-state automata. We will present the proofs only for the cases $op(f, g)$ and $\text{iter}(f, c, op)$ and leave the rest of the cases as exercises for the reader. For the cases that we present, we will only give the main ingredients of the construction.

Suppose that \mathcal{A} and \mathcal{B} are CRAs. We will construct a CRA \mathcal{C} that computes $op(\llbracket \mathcal{A} \rrbracket, \llbracket \mathcal{B} \rrbracket)$, where op is a binary operation on D . The state space of \mathcal{C} is the product of the state space of \mathcal{A} and the state space of \mathcal{B} . The registers of \mathcal{C} is the union of the registers of \mathcal{A} and the registers of \mathcal{B} . The automaton \mathcal{C} simulates the parallel execution of \mathcal{A} and \mathcal{B} , where the registers of the two machines are updated separately. When a state (p, q) of

\mathcal{C} consists of a final state p of \mathcal{A} and a final state q of \mathcal{B} , the finalization term is defined so that it combines the output values of the two automata \mathcal{A} and \mathcal{B} . This construction preserves the “copylessness” property: if the automata \mathcal{A} and \mathcal{B} are copyless, then so is \mathcal{C} .

Suppose that \mathcal{A} is a CRA, c is a data value, and op is a binary operation on values. We will construct a CRA \mathcal{C} that computes $\text{iter}(\llbracket \mathcal{A} \rrbracket, c, op)$. W.l.o.g. we assume that we can use ε -transitions, and therefore \mathcal{A} has a unique initial state and a unique final state. We will also assume that $\varepsilon \notin R(\llbracket \mathcal{A} \rrbracket)$, because this special case can be handled separately to avoid ε -cycles. We will describe the construction of \mathcal{C} as a sequence of modifications on \mathcal{A} .

- First, add a new register y to \mathcal{A} that is meant to hold the aggregate value computed in the iteration. The register y stays unchanged in the old transitions of \mathcal{A} .
- Add an ε -transition from the unique final state to the unique initial state. In this transition, update the aggregate value of y to reflect one more step of the iteration.
- Add a new initial and final state for computing the initial aggregate c .
- Since the NCRA that has resulted from the previous steps could potentially be *ambiguous*, take the product with a (register-free) DFA that accepts the language $R(\llbracket \mathcal{A} \rrbracket)^*$.

The last step of the construction removes the ambiguity. So, the resulting CRA is unambiguous and computes the transduction $\text{iter}(f, c, op)$. This construction preserves the property of “copylessness”: if \mathcal{A} is copyless then so is \mathcal{C} . \square

One can show that the class of *all streaming algorithms* with an appropriate streaming complexity bound is not closed under these regular combinators. For instance, suppose the algorithm must have $\log(n)$ space in the input stream and $\log(n)$ time to process each element. Let $\Sigma = \{\mathbf{a}, \mathbf{b}\}$. Define data transductions f, g where f is equal to 1 on all sequences with the same number of as and bs , and undefined otherwise, and g is equal to 1 on sequences with twice as many as as bs , and undefined otherwise. Then, f and g can be given as streaming algorithms, but $\text{split}(f, g, op)$ cannot, thus justifying the role of regularity in guaranteeing modular composition, as we mentioned in the introduction.

2.7. Complexity of Evaluation

We are interested in evaluation in the streaming model of computation. The main resources in this model are: 1) the total space required by the algorithm, and 2) the time needed to process each data item. Both these resources are given as a function of the length of the stream that has been seen so far. Ideally, both these key parameters should be constant or logarithmic in the length of the stream.

Suppose that $\mathcal{A} = (Q, X, \Delta, I, F)$ is a trim UCRA. The computation of the semantics of \mathcal{A} in the streaming model of computation is similar to the deterministic simulation of classical NFAs, in which the set of *active states* is maintained and updated at every step of the computation. The difference for CRAs is that we need to maintain for every active state q a register assignment $\alpha_q : X \rightarrow D$. Unambiguity guarantees that at every step of the computation we need to store *at most one* register assignment for every state. This property does not hold for ambiguous machines, as there are several different paths that could lead to the same state. So, the total number of values that are needed to store is bounded above by $|Q| \cdot |X|$. For deterministic CRAs, the corresponding upper bound is $|X|$. Analogous upper bounds can be obtained for the time needed per element for both the unambiguous and the deterministic case.

A more precise accounting of the used computational resources requires looking at the number of bits that are needed to store each value during the computation. In order to make such finer distinctions, we need to take into account the operations that are allowed in \mathcal{O} . For example, suppose that the values are the natural numbers and that \mathcal{O} contains numerical constants, binary addition, max and min. In the copyless case, the contents of registers can grow only linearly in the length of the input and therefore a logarithmic number of bits is required for each register. In the copyful case, however, even if we restrict \mathcal{O} to contain only 0 and binary addition, the contents of registers can grow exponentially and therefore a linear number of bits may be needed.

In some applications, the data values range over a finite domain (e.g., 32-bit integers, 64-bit floating-point numbers, and so on), and therefore the space needed to store each register is constant. Moreover, the operations on such bounded values require constant time. In this setting, the space and time-per-element resource requirements for evaluating CRAs is constant.

3. MSO-Definable Transductions

In this section, we give logical characterizations of the classes of transductions defined by the Cost Register Automata of Definition 1:

- (1) streamable regular or $\text{SR}(\mathcal{O})$ (computed by copyful UCRA's),
- (2) streamable linear regular or $\text{SLR}(\mathcal{O})$ (computed by copyless UCRA's).

These logical characterizations are in terms of word-to-DAG transformations which are given by formulas written in the language of monadic second-order logic (MSO).

The MSO-definable transductions that we consider here specify output graphs *with no backward edges*. This restriction for the output DAGs to only have forward edges is crucial so that the function is implementable by an automaton which maintains values (as opposed to maintaining terms which may have “holes”, as in [9]). Indeed, we will show that the forward-only MSO-definable transductions are exactly the transformations that are implementable by the value-based UCRA's of Definition 1.

This class of transductions has previously arisen in the study of attribute grammars [5] under the name of *direction preserving* MSO graph transductions (page 11 in [5]). An important difference here compared to [5] is the characterization of the transduction classes in terms of automata ([5] uses attribute grammars) whose efficient evaluation in the streaming model of computation is obvious. Another technical difference is that the transductions in the present paper involve data values, and we use a special symbol `val` and the position of a `val`-labeled vertex in order the semantics of an MSO transduction. Our proof is self-contained, and exhibits a pleasant uniformity due to the translation from forward-only MSO transductions to *unambiguous* CRAs. Recall that unambiguity is essential in order to capture the class SLR in the copyless CRA model (Theorem 13). In the word-to-DAG case, the output graphs can have more than one edge emanating from a vertex, which corresponds to copying a register in the machine model. In the word-to-tree case, each vertex of the output graph has at most one edge emanating from it, which implies that registers are used in a copyless manner in the machine model.

An *MSO transduction* specifies the output DAG using MSO formulas for the vertices and edges that are interpreted over the entire input stream. The

challenge is to convert this global description of the output into local computation rules for the automaton. In particular, the existence of a vertex or edge depends on a regular property of the *future* of the input (regular look-ahead). To deal with this regular look-ahead, we make use of *unambiguously nondeterministic* CRAs. The unambiguous nondeterminism conveniently allows us to make guesses regarding a regular property of the future of the input.


The crucial intermediate result for obtaining equivalence between MSO transductions and CRAs asserts that every MSO transduction can be transformed equivalently into a *single-step* MSO transduction, namely one where the edges of the output DAG extend over at most one input element. To implement a single-step MSO transduction it is then sufficient to maintain as many data registers as the maximum number of DAG vertices that appear in some position. In this section, we will allow the use of ε -transitions for unambiguous CRAs. This is w.l.o.g. because ε -transitions can be eliminated with a straightforward variant of the ε -elimination procedure for classical finite-state automata.

Let Σ be the finite alphabet of tags, D a (possibly infinite) set of data values, and \mathcal{O} a collection of constants and operations on D . Recall that we only consider transductions $f : (\Sigma \times D)^* \rightarrow D$ that satisfy the following condition: whether $f(w)$ is defined or not depends only on the sequence of tags $w|_{\Sigma}$.

3.1. MSO Transductions

An MSO-definable graph transduction [10] specifies a partial function from labeled graphs to labeled graphs. The nodes, edges and labels of the output graph are described in terms of MSO formulas over the nodes, edges and labels of the input graph.

The tag projection of a data word $w = (a_1, d_1) (a_2, d_2) \dots (a_n, d_n)$ of length n over $\Sigma \times D$ can be represented as a labeled directed graph $G(w)$ with $n + 1$ vertices $\{0, 1, \dots, n\}$ and n Σ -labeled edges $(i - 1, a_i, i)$ for all $i = 1, \dots, n$. For example, we visualize the graph $G(w)$ of the word $w = (\mathbf{A}, 1) (\mathbf{B}, 2) (\mathbf{A}, 3) (\mathbf{B}, 4) (\mathbf{A}, 1)$ over the alphabet with tags $\Sigma = \{\mathbf{A}, \mathbf{B}\}$ and values that are natural numbers as follows:



A graph $G(w)$ specifies a relational structure whose domain is the set of vertices and which has a binary predicate edge_a^w for every tag $a \in \Sigma$ consisting


of the \mathbf{a} -labelled edges. We use the *monadic second-order language* $\text{MSO}(\Sigma)$ to express properties of words when represented as graphs. For every tag \mathbf{a} of the alphabet Σ , the language has a binary predicate symbol $\text{edge}_{\mathbf{a}}$. Informally, $\text{edge}_{\mathbf{a}}(x, y)$ means that there is an edge from x to y labeled with the tag \mathbf{a} . The language also contains the equality predicate $=$ (i.e., $x = y$ means that the vertices x and y are equal) and the containment predicate \in (i.e., $x \in X$ means that the vertex x belongs to the set of vertices X). The formulas of $\text{MSO}(\Sigma)$ are built from atomic formulas using the Boolean connectives, first-order quantification (over vertices of $\mathbf{G}(w)$), and second-order monadic quantification (over sets of vertices of $\mathbf{G}(w)$). We do not need to include \leq , because we can introduce it as an abbreviation using $\text{edge}_{\mathbf{a}}$.

Note that we use binary predicates $\text{edge}_{\mathbf{a}}(x, y)$ to encode tags, instead of the more commonly used unary predicates [11]. That is, $\mathbf{G}(w)$ has labeled edges instead of labeled vertices. The reason for this encoding is that we want to allow transductions that map the empty input word into some output graph (as is done, for example, in page 232 of [12]). In any case, this difference is inconsequential for the definable class of functions.

We write $\text{Word}(\Sigma)$ for the class of graph structures over the signature of $\text{MSO}(\Sigma)$ that correspond to finite words over Σ . If ϕ is a sentence of $\text{MSO}(\Sigma)$ and w is a word over Σ , then we write $w \models \phi$ to denote that the graph $\mathbf{G}(w)$ satisfies ϕ . The abbreviation $\text{edge}(x, y) := \bigvee_{\mathbf{a} \in \Sigma} \text{edge}_{\mathbf{a}}(x, y)$ says there is an edge from x to y . In turn, we abbreviate \leq as

$$(x \leq y) := \forall X. x \in X \wedge (\forall u \forall v. u \in X \wedge \text{edge}(u, v) \rightarrow v \in X) \rightarrow y \in X.$$

A *ranked alphabet* Γ is a finite set with assignment of a natural number to each element, call its *arity*. A *directed acyclic graph* (DAG) over a set of ranked symbols Γ has Γ -labeled vertices, a single sink vertex (its *root*), and edges indexed between 1 and the maximum arity, such that the graph respects the arity of the symbols of Γ . For example, suppose $\Gamma = \{a, b, f\}$, where a and b are constants (arity 0) and f is binary (arity 2). The following graph is a DAG over Γ , where we use arrows \rightarrow for the first argument of f , and arrows \twoheadrightarrow for the second argument of f :



Definition 15 (Forward-only word-to-DAG MSO transductions). Let $\Gamma = \mathcal{O} \cup \{\text{val}\}$ be the (finite) set of ranked operation symbols that includes the operations \mathcal{O} on the data values D as well as a fresh constant symbol val . Let i_{\max} denote the maximum arity of a symbol in Γ . A *forward-only MSO-definable word-to-DAG transduction* from $(\Sigma \times D)^*$ to D consists of:

1. An $\text{MSO}(\Sigma)$ sentence ϕ_{dom} , called the *domain sentence*.
2. A finite *copy set* C .
3. For every copy element $c \in C$ and every symbol $\gamma \in \Gamma$, an $\text{MSO}(\Sigma)$ *vertex formula* $\phi_{\gamma}^c(x)$ with one free variable x .
4. For all copy elements $c, d \in C$ and for all argument indices i from 1 to i_{\max} , an $\text{MSO}(\Sigma)$ *edge formula* (for the i -th argument) $\psi_i^{c,d}(x, y)$ with two free variables x and y s.t. $\models \forall x \forall y. \psi_i^{c,d}(x, y) \rightarrow x \leq y$.

such that all of the well-formedness conditions shown in Observation 16 hold. These conditions assert that the formulas define a properly structured DAG over the symbols Γ in the output structure (which consists of C copies of the vertices of $\mathbf{G}(w)$).

Semantics: The MSO transduction *denotes* a partial function $\tau : (\Sigma \times D)^* \dashrightarrow D$ whose domain is the set of models $\llbracket \phi_{\text{dom}} \rrbracket$ of the domain sentence. For every data word w satisfying $\llbracket \phi_{\text{dom}} \rrbracket$, the transduction specifies a single DAG over the symbols $\mathcal{O} \cup \{\text{val}\}$. Every symbol in \mathcal{O} is interpreted as an operation on D , and val is a special symbol for the data value of the input string at the corresponding position. The DAG therefore can be evaluated to a value in D , which is defined to be the value of τ on w .

Note that while the formulas above *define* an output-graph, they are only logically *interpreted over* an input graph $\mathbf{G}(w)$. This is why they define a single output structure (not, e.g., a set of output structures or a relation between input and output structures). The well-formedness conditions then guarantee this output structure has the correct form. The conditions are stated as MSO formulas, but it is not actually necessary for our results that they are expressible in MSO; rather, we only use the fact that the output of every forward-only MSO transduction is a properly-structured DAG, and we have baked this condition into the definition of such transductions. However, the fact that they are $\text{MSO}(\Sigma)$ formulas means the conditions are effectively checkable (decidable).

Observation 16 (Well-formedness). We list here the well-formedness conditions for the forward-only MSO-definable transductions of Definition 15. Note that all conjunctions and disjunctions (\bigvee, \bigwedge) are finite, and there are finitely many formulas overall. Logical validity is over all structures in $\text{Word}(\Sigma)$. I.e., for every input graph $G(w) \in \text{Word}(\Sigma)$ (and all interpretations of the free variables), the formula should be true.

- (1) **At most one vertex label:** $\bigwedge_{\gamma} \bigwedge_{\delta \neq \gamma} \neg \phi_{\gamma}^c(x) \vee \neg \phi_{\delta}^c(x)$ is valid for every $c \in C$.
- (2) **Edges connect active vertices:** $\psi_i^{c,d}(x, y) \rightarrow \phi^c(x) \wedge \phi^d(y)$ is valid for all $c, d \in C$ and $i = 1, 2, \dots, i_{\max}$, where $\phi^c(x) := \bigvee_{\gamma \in \Gamma} \phi_{\gamma}^c(x)$ abbreviates that the c -vertex at x is active.
- (3) **(Non)existence of output:** The formulas $\phi_{\text{dom}} \rightarrow \exists x \bigvee_c \phi^c(x)$ and $\neg \phi_{\text{dom}} \rightarrow \forall x \bigwedge_c \neg \phi^c(x)$ are valid.
- (4) **No local cycle at any position:** For every potential cycle $c_1 \xrightarrow{i_1} c_2 \xrightarrow{i_2} \dots \xrightarrow{i_{k-1}} c_k \xrightarrow{i_k} c_1$ of length $k \leq |C|$, where $c_j \in C$ and $1 \leq i_j \leq i_{\max}$ are edge labels, the following formula is valid:

$$\neg \exists x. \psi_{i_1}^{c_1, c_2}(x, x) \wedge \psi_{i_2}^{c_2, c_3}(x, x) \wedge \dots \wedge \psi_{i_k}^{c_k, c_1}(x, x).$$

- (5) **The arity of labels is respected:** For every symbol $\gamma \in \Gamma$ of arity n , every $i = 1, \dots, n$, and every $c \in C$, the following formula is valid:

$$\phi_{\gamma}^d(y) \rightarrow \exists x. \left(\bigvee_c (\psi_i^{c,d}(x, y) \wedge \bigwedge_{\hat{c} \neq c} \neg \psi_i^{\hat{c},d}(x, y)) \wedge \forall \hat{x} \bigwedge_{\hat{c}} (\psi_i^{\hat{c},d}(\hat{x}, y) \rightarrow \hat{x} = x) \right).$$

- (6) **Global uniqueness of sink:** For all $c \in C$, the formula

$$\text{sink}_c(x) \rightarrow \bigwedge_{d \neq c} \neg \text{sink}_d(x) \wedge \forall y. (y \neq x \rightarrow \bigwedge_d \neg \text{sink}_d(y))$$


is valid, where $\text{sink}_c(x) := \forall y \bigwedge_d \bigwedge_i \neg \psi_i^{c,d}(x, y)$ abbreviates that there is no outgoing edge from copy c at position x .

- (7) **No value at first position:** $\text{first}(x) \rightarrow \bigwedge_c \neg \phi_{\text{val}}^c(x)$ is valid, where $\text{first}(x) := \forall y. x \leq y$ abbreviates the very first position. The constant symbol val , which is meant to refer to the input data value immediately before the point of occurrence, should not appear in this position.

Output graph	Class of transductions
Tree	Streamable Linear Regular
DAG	Streamable Regular

Table 3: Classes of MSO-definable transductions with forward-only edges.

Example 17 (MSO Transduction). The tag alphabet is $\Sigma = \{A, B\}$ and the data domain D is the set of natural numbers. The collection \mathcal{O} of operations on D contains the constant 0 and addition.



We will describe the transduction that sums up all data values that are tagged with A. The copy set is taken to be $C = \{V, S\}$. The domain sentence is $\phi_{\text{dom}} = \text{true}$ and the vertex formulas are:

$$\phi_{\text{val}}^V(x) = \phi_+^S(x) = \text{edge}_A(x-1, x) \quad \phi_0^S(x) = \text{first}(x)$$


where $\text{edge}_A(x-1, x)$ is notation for $\exists y. \text{edge}_A(y, x)$. The edge formulas are given as follows:

$$\begin{aligned} \psi_2^{V,S}(x, y) &= (x = y) \wedge \text{edge}_A(x-1, x) \\ \psi_1^{S,S}(x, y) &= (\text{first}(x) \vee \text{edge}_A(x-1, x)) \wedge \text{edge}_A(y-1, y) \wedge \\ &\quad \forall z. x < z < y \rightarrow \neg \text{edge}_A(z-1, z) \end{aligned}$$

All omitted vertex and edge formulas are equal to false.


As we will see later, the data transductions that are definable by forward-only word-to-DAG MSO-transductions are exactly the streamable regular transductions. Similarly, the data transductions that are definable by forward-only word-to-tree MSO-transductions will be shown to be exactly the streamable linear regular transductions. See Table 3. The condition that the output graph is a tree can be guaranteed by requiring that every vertex of the graph has at most one outgoing edge. This condition can be expressed by an MSO formula, and therefore its validity can be effectively checked.

Definition 18 (Models of Open Formulas). Let $\phi(x)$ be an $\text{MSO}(\Sigma)$ formula with one free individual variable. The set of *models* of $\phi(x)$, denoted $\llbracket \phi(x) \rrbracket$,



is defined to be the set of pairs $(u, v) \in \Sigma^* \times \Sigma^*$ such that the word $u \cdot v$ satisfies $\phi(x)$ under the variable assignment $x \mapsto |u|$.

Let $\psi(x, y)$ be an $\text{MSO}(\Sigma)$ formula with two free individual variables such that $\models \psi(x, y) \rightarrow x \leq y$. The set of *models* of $\psi(x, y)$, denoted $\llbracket \psi(x, y) \rrbracket$,



is defined to be the set of triples $(u, v, w) \in \Sigma^* \times \Sigma^* \times \Sigma^*$ such that the word $u \cdot v \cdot w$ satisfies $\psi(x, y)$ under the variable assignment $x \mapsto |u|$, $y \mapsto |uv|$.

Lemma 19 ([13], Lemma 5). Let $\phi(x)$ be an $\text{MSO}(\Sigma)$ formula with one free individual variable. The set of models $\llbracket \phi(x) \rrbracket$ is a finite disjoint union of sets of the form $\llbracket r_1 \rrbracket \times \llbracket r_2 \rrbracket$, where r_1 and r_2 are regular expressions over Σ .

Similarly, suppose that $\psi(x, y)$ is an $\text{MSO}(\Sigma)$ formula with two free individual variables such that $\models \psi(x, y) \rightarrow x \leq y$. The set of models $\llbracket \psi(x, y) \rrbracket$ is a finite disjoint union of sets of the form $\llbracket r_1 \rrbracket \times \llbracket r_2 \rrbracket \times \llbracket r_3 \rrbracket$, where r_1 , r_2 and r_3 are regular expressions over Σ .

Proof sketch. This is a folklore result for regular languages (see, for example, Lemma 5 of [13]). For the first part of the lemma, the idea is to construct the DFA \mathcal{A} that accepts the models of $\phi(x)$. The annotation with x is performed by the automaton by allowing some states to be labeled with x . For every state q of \mathcal{A} that is labeled with x , we include the regular expression for the set $L_q \cdot R_q$, where L_q is the set of strings that lead from the initial state to q , and R_q is the set of strings that lead from q to a final state of \mathcal{A} . The construction for the second part of the lemma is similar. \square

Using Lemma 19 an MSO transduction can be transformed into a set of rules that are expressed equivalently with regular expressions instead of MSO formulas. The domain sentence ϕ_{dom} can be given as a regular expression. A vertex formula $\phi_\gamma^c(x)$ can be given as a finite set of pairwise disjoint *vertex rules* of the form

$$\text{label } \gamma \text{ at } c : r_1; r_2,$$

where r_1 and r_2 are regular expressions. The rule specifies a prefix (from start to x) with r_1 and a suffix (from x to end) with r_2 so that the copy c of the output DAG is active at position x . Similarly, an edge formula $\psi_i^{c,d}(x, y)$ of the transduction can be presented as a finite set of pairwise disjoint *edge rules* of the forms

$$\text{edge } c \rightarrow_i d : r_1; r_2; r_3 \quad [\text{forward edges}]$$

where r_1, r_2 and r_3 are regular expressions. The rule specifies a prefix (from start to x) with r_1 , a middle part (from x to y) with r_2 , and a suffix (from y to end) with r_3 so that the output DAG has an edge from copy c at position x to copy d at position y . We have required in the definition that edges be forward-only, so there are no models in which y comes before x in $\psi_i^{c,d}(x, y)$.

Example 20. Using Lemma 19 we will present the MSO transduction of Example 17 equivalently with regular expressions. We have: $\phi_{\text{dom}} = \Sigma^*$ and

$$\begin{aligned} \phi_{\text{val}}^{\text{V}} = \phi_+^{\text{S}} = (\Sigma^* \mathbf{A}); \Sigma^* & & \psi_2^{\text{V,S}} = (\Sigma^* \mathbf{A}); \varepsilon; \Sigma^* \\ \phi_0^{\text{S}} = \varepsilon; \Sigma^* & & \psi_1^{\text{S,S}} = (\varepsilon \cup \Sigma^* \mathbf{A}); (\mathbf{B}^* \mathbf{A}); \Sigma^* \end{aligned}$$

The vertex rule ϕ_+^{S} , for example, says that the copy S is active with label $+$ at every position immediately after the occurrence of a letter \mathbf{A} in the input string. Note that here, each formula becomes a single rule; in general, each formula will become a set of pairwise disjoint rules.

A MSO transduction is said to be *single-step* if the edges that it specifies are restricted so that they refer to positions that are at most 1 apart: the formula

$$\psi_i^{c,d}(x, y) \rightarrow (x = y) \vee \text{edge}(x, y)$$

is valid for all $c, d \in C$ and every argument index i .

Lemma 21 (Single-Step Decomposition). Let Σ be a finite tag alphabet, and D be a set of data values with operations \mathcal{O} . A forward-only MSO-definable word-to-DAG transduction over Σ, D, \mathcal{O} can be transformed into an equivalent single-step transduction over $\Sigma, D, \mathcal{O} \cup \{\text{id}\}$, where $\text{id} : D \rightarrow D$ is the identity function on D .

Proof. Suppose that the transduction is presented in the form using regular expressions instead of MSO formulas, as described earlier. We will show how

to eliminate rules that violate the single-step constraint by replacing them with rules that are either ε -transitions or single-letter transitions.

Consider a rule $\psi_i^{c,d} = c \rightarrow_i d : r_1; r_2; r_3$ for i -edges from copy c to copy d which is not an ε -transition, i.e. $r_2 \neq \varepsilon$. The idea is to decompose an edge specified by the rule, which spans a substring matching r_2 , by simulating the execution of an automaton for the expression r_2 . Take the minimal DFA for r_2 and prune it into an equivalent partial DFA so that every state has a path to some final state. Suppose this automaton is $(Q, \Sigma, I, \Delta, F)$. For a state $q \in Q$ define the regular set

$$L_q = \{w \in \Sigma^* \mid \Delta(I, w) = q\},$$

that is, L_q is the set of strings over Σ that lead from the initial state I to q . Moreover, define the regular set

$$R_q = \{w \in \Sigma^* \mid \Delta(q, w) \in F\},$$

that is, R_q is the set of strings over Σ that lead from q to some final state. Notice that $L_q \cdot R_q \subseteq r_2$ for every state q , and additionally $r_2 = \bigcup \{L_q \mid q \in F\}$.


For every state q of the automaton, we introduce a fresh copy element v_q . We can replace the rule $\psi_i^{c,d}$ by the following set of vertex and edge rules:

$$\begin{aligned} &\text{label id at } v_q : (r_1 \cdot L_q); (R_q \cdot r_3) \\ &\text{edge } c \rightarrow_1 v_I : r_1; \varepsilon; (r_2 \cdot r_3) \\ &\text{edge } v_q \rightarrow_i d : (r_1 \cdot L_q); \varepsilon; r_3, \text{ for every } q \in F \\ &\text{edge } v_p \rightarrow_1 v_q : (r_1 \cdot L_p); \mathbf{a}; (R_q \cdot r_3), \text{ for every } (p, \mathbf{a}, q) \in \Delta. \end{aligned}$$


Since the automaton is deterministic, we are guaranteed that all of these rules are disjoint, and moreover that the modified transduction satisfies all the well-formedness conditions. In addition, it preserves the property of the output DAG being a tree.

The construction is iterated until there is no edge rule left that violates the single-step condition. The resulting transduction can then be equivalently presented using MSO formulas. \square

Example 22 (Single-step). We will apply now the construction of the proof of Lemma 21 to the transduction of Example 17. Using the presentation of Example 20 we see that the only edge rule that violates the single-step condition is $\psi_1^{s,s}$. Now, the pruned partial DFA for B^*A is the following:



The initial state is indicated with a gray background. We introduce fresh copy elements p and q , one for every state of the automaton for B^*A . So, the new copy set is $C' = \{V, S, p, q\}$. The following figure illustrates how the edges defined by $\psi_1^{S,S}$ are decomposed into ε -transitions and single-letter transitions.



So, the transduction is specified by the following rules:

$$\begin{array}{ll}
\phi_{\text{dom}} = \Sigma^* & \psi_2^{V,S} = (\Sigma^*A); \varepsilon; \Sigma^* \\
\phi_{\text{val}}^V = (\Sigma^*A); \Sigma^* & \psi_1^{S,p} = (\varepsilon \cup \Sigma^*A); \varepsilon; B^*A\Sigma^* \\
\phi_0^S = \varepsilon; \Sigma^* & \psi_1^{p,p} = (\varepsilon \cup \Sigma^*A)B^*; B; B^*A\Sigma^* \\
\phi_+^S = (\Sigma^*A); \Sigma^* & \psi_1^{p,q} = (\varepsilon \cup \Sigma^*A)B^*; A; \Sigma^* \\
\phi_{\text{id}}^p = (\varepsilon \cup \Sigma^*A)B^*; B^*A\Sigma^* & \psi_1^{q,S} = (\varepsilon \cup \Sigma^*A)B^*A; \varepsilon; \Sigma^* \\
\phi_{\text{id}}^q = (\varepsilon \cup \Sigma^*A)B^*A; \Sigma^* &
\end{array}$$

The ε -transitions $\psi_1^{S,p}$, $\psi_1^{q,S}$ and the single-letter transitions $\psi_1^{p,p}$, $\psi_1^{p,q}$ expand the edges of $\psi_1^{S,S}$ according to the automaton for B^*A .

Theorem 23. Every single-step word-to-dag (resp., word-to-tree) forward-only MSO transduction can be implemented by a copyful (resp., copyless) UCRA.

Proof. Suppose we are given an MSO transduction, presented in the form with regular expressions instead of formulas. Given an input string, a position x in the string and a copy element c , in order to decide whether the vertex (c, x) of the output DAG is labeled with f we have to check if the suffix

(from x to the end) satisfies the regular expression r_2 , where $\phi_f^c = r_1$; r_2 is a vertex rule of the transduction. Similarly, to decide whether an edge $c \rightarrow_i d$ emanates from copy c at position x we have to check if the suffix (from x to end) satisfies the expressions $r_2 \cdot r_3$, where $\psi_i^{c,d} = r_1$; r_2 ; r_3 is an edge rule of the transduction. Finally, to decide whether an edge $d \rightarrow_i c$ has the vertex (c, x) as destination we have to check if the suffix (from x to end) satisfies the expression r_3 , where $\psi_i^{d,c} = r_1$; r_2 ; r_3 is an edge rule of the transduction.

Let \mathcal{R} be the collection of all regular sets needed for these suffix tests, that is, $\llbracket r_2 \rrbracket$ from every vertex rule r_1 ; r_2 , as well as both $\llbracket r_2 \cdot r_3 \rrbracket$ and $\llbracket r_3 \rrbracket$ from every edge rule r_1 ; r_2 ; r_3 . Now, define \mathcal{R}' to be the smallest collection of regular sets that contains \mathcal{R} and is closed under derivatives. It can be seen that $\mathcal{R}' = \bigcup_{S \in \mathcal{R}} \mathcal{R}'_S$, where \mathcal{R}'_S is the collection of regular sets obtained in the following way: take the collection of all derivatives of S , which correspond to the unique minimal automaton for S . Finally, define $\text{At}_{\mathcal{R}}$ to be the collection of atoms (nonempty minimal elements) of the Boolean algebra generated by \mathcal{R}' . The elements of $\text{At}_{\mathcal{R}}$ partition Σ^* (they are pairwise disjoint and their union is equal to Σ^*), and they define an *unambiguous NFA*, where the language accepted from every state $S \in \text{At}_{\mathcal{R}}$ is equal to S . To see how the transitions of the NFA $\text{At}_{\mathcal{R}}$ are calculated, consider a state (atom) $S = S_1 \cap \sim S_2 \cap S_3$ where each of S_i is an element of \mathcal{R}' . For a letter $\mathbf{a} \in \Sigma$ the derivative of S is


$$\begin{aligned} D_{\mathbf{a}}(S) &= D_{\mathbf{a}}(S_1) \cap D_{\mathbf{a}}(\sim S_2) \cap D_{\mathbf{a}}(S_3) \\ &= D_{\mathbf{a}}(S_1) \cap \sim D_{\mathbf{a}}(S_2) \cap D_{\mathbf{a}}(S_3), \end{aligned}$$

and it is partitioned by some atoms T_1, \dots, T_ℓ of $\text{At}_{\mathcal{R}}$ because every $D_{\mathbf{a}}(S_i)$ belongs to \mathcal{R}' . So, $S \rightarrow^{\mathbf{a}} T_i$ are all the transitions from S on letter \mathbf{a} in the unambiguous NFA $\text{At}_{\mathcal{R}}$. The automaton $\text{At}_{\mathcal{R}}$ thus constructed is called the *future automaton* for the transduction, because every state specifies the possible future suffixes.

Similarly, the application of the rules of the MSO transduction depends on regular properties of the input prefix seen so far. For a vertex rule r_1 ; r_2 , an ε -transition rule r_1 ; ε ; r_2 , and a letter transition rule r_1 ; \mathbf{a} ; r_2 , we need to check whether the input prefix satisfies the pattern r_1 . To perform all these tests we construct the *past automaton*, which simulates the parallel execution of automata for all these prefix patterns.

The product of the future automaton and the past automaton is an unambiguous NFA, which we call the *future-past automaton* where every state

specifies exactly which rules of the MSO transduction apply at the current position. So, every state of the future-past automaton specifies precisely the *shape* of a position: which copy elements are active, which ε -edges are enabled, and which outgoing letter-edges are enabled. A state only needs to specify values for variables whose values are going to be used in the next step of the computation. The expressions for these outgoing variables can be read off immediately from the shape. For example, suppose the state T of the future-past automaton specifies the following active vertices and edges:



We know that any incoming transition $S \rightarrow T$ from some state S must specify a value for the variable x_q (for the copy q), because it has to be used. The state T has to specify a new value for x_q , because the only outgoing letter-edge points to the copy q . So, the variable update for the transition $S \rightarrow T$ must be $x_q := x_q + \text{val}$. If a state is initial, then it must be possible to calculate the values for the variables without relying on any incoming value from a previous state. If the shape has an active copy without an incoming/outgoing edge (on the right), then its value is a potential final value for the computation and is stored in a special variable x_F . For the particular case where no vertex is active, we have to propagate the value of the final variable with the update $x_F := x_F$.

It is easy to check that if the transduction is word-to-tree (that is, every vertex in a shape has at most one outgoing edge), then the constructed CRA is copyless. \square

Example 24 (Future Automaton). For the single-step MSO transduction of Example 22 we will construct the future automaton, as described in the proof of Theorem 23. The collection of regular sets needed for suffix tests is:

$$\mathcal{R} = \{\Sigma^*, B^*A\Sigma^*, B^+A\Sigma^*, A\Sigma^*\}.$$

After closing under derivatives we obtain the collection

$$\mathcal{R}' = \{\Sigma^*, B^*A\Sigma^*, B^+A\Sigma^*, A\Sigma^*, \emptyset\}.$$


The set $B^*A\Sigma^*$ consists of all words that contain at least one A , so its complement is B^* . The set $B^+A\Sigma^*$ consists of all words that start with B and contain at least one A , so its complement is $A\Sigma^* \cup B^*$. The complement of $A\Sigma^*$ is $\varepsilon \cup B\Sigma^*$. The elements of $\text{At}_{\mathcal{R}}$ are therefore the minimal nonempty intersections of elements from $B^*A\Sigma^*$, $B^+A\Sigma^*$, $A\Sigma^*$, B^* , and $\varepsilon \cup B\Sigma^*$. We claim that the Boolean atoms generated by this list are

$$\text{At}_{\mathcal{R}} = \{A\Sigma^*, B^+A\Sigma^*, B^*\}.$$

Indeed, $\text{At}_{\mathcal{R}}$ partitions Σ^* and moreover we have:

$$\begin{aligned} B^*A\Sigma^* &= A\Sigma^* \cup B^+A\Sigma^* \\ \varepsilon \cup B\Sigma^* &= B^* \cup B^+A\Sigma^* \end{aligned}$$

The unambiguous NFA $\text{At}_{\mathcal{R}}$ is illustrated below:




All the states are marked as initial (with the gray background) because the domain of the transduction is Σ^* .

Example 25 (Past Automaton). The expression $(\varepsilon \cup \Sigma^*A)B^*$ is equivalent to Σ^* , hence the rules of the single-step transduction of Example 22 can be simplified:


$$\begin{array}{ll} \phi_{\text{val}}^V = \Sigma^*A; \Sigma^* & \psi_2^{V,S} = \Sigma^*A; \varepsilon; \Sigma^* \\ \phi_0^S = \varepsilon; \Sigma^* & \psi_1^{S,p} = (\varepsilon \cup \Sigma^*A); \varepsilon; B^*A\Sigma^* \\ \phi_+^S = \Sigma^*A; \Sigma^* & \psi_1^{p,p} = \Sigma^*; B; B^*A\Sigma^* \\ \phi_{\text{id}}^p = \Sigma^*; B^*A\Sigma^* & \psi_1^{p,q} = \Sigma^*; A; \Sigma^* \\ \phi_{\text{id}}^q = \Sigma^*A; \Sigma^* & \psi_1^{q,S} = \Sigma^*A; \varepsilon; \Sigma^* \end{array}$$

The regular expressions Σ^*A , ε , Σ^* , $\varepsilon \cup \Sigma^*A$ are relevant for checking whether the input prefix seen so far satisfies the left-hand expression of a rule. The automaton that simulates the parallel execution of DFAs for these expressions

is the following:



The labels are computed by taking derivatives of the expressions. It is more convenient to label a state q with the regular expression that denotes the set of strings that lead from the initial state to q .



The label of each state determines if the automaton accepts for each one of the “past” expressions Σ^*A , ε , Σ^* , $\varepsilon \cup \Sigma^*A$.

Example 26 (Future-Past Automaton, CRA). The product of the future automaton from Example 24 and the past automaton from Example 25 is the *future-past* automaton of Figure 1. Every state q is labeled with L_q (top), the set of strings that lead from some initial state to q , and with R_q (bottom), the set of strings that lead from q to some final state. The purpose of these labelings is that they specify exactly which rules of the transaction are enabled at a particular state. In Figure 2 every state of the future-past automaton is annotated with a shape that indicates the active copies and the edges (ε -edges and letter-edges) that are enabled. From these shapes we can immediately read off the variable updates for the transitions, and the initialization of the variables in the initial states. The resulting CRA is shown in Figure 3. Every state q_i of the CRA is annotated with the set of variables that have to be set when the computation reaches q_i .

Corollary 27. Forward-only word-to-DAG (resp., word-to-tree) MSO transductions can be implemented by copyful (resp., copyless) UCRA.

Proof. Suppose we are given an MSO transduction. By virtue of Lemma 19 we can present it equivalently using regular expressions instead of MSO formulas. Lemma 21 says that the transduction can be transformed into an




Figure 1: The future-past automaton for the single-step transduction of Example 22.

equivalent one, with potentially more copy elements, in which edges are either ϵ -edges or single-letter edges. From this modified transduction we can construct an unambiguous CRA that computes the transduction, as shown in Theorem 23. \square

Theorem 28. The transductions of the class **SR** (resp., **SLR**) can be defined in forward-only word-to-DAG (resp., word-to-tree) MSO.

Proof. With a construction that is a simple variant of the classical transformation of an automaton into a MSO formula defining the same language, we can convert an unambiguous CRA into a single-step MSO transduction. We need one copy element per register, as well as sufficient copies to construct the algebraic terms that appear in all variable updates of the CRA. The resulting MSO transduction may specify vertices labeled with the identity function id . Sequences of edges of the form $f \rightarrow_1 \text{id} \rightarrow_1 \dots \rightarrow_1 \text{id} \rightarrow_i g$ can be specified with an MSO formula, and can therefore be replaced by a single edge of the form $f \rightarrow_i g$. This results in a (not necessarily single-step) MSO transduction that specifies the desired function. This construction handles copyful and copyless UCRA uniformly: i.e., on a copyless UCRA it produces a word-to-tree MSO transduction. The well-formedness conditions for MSO transductions hold, including that all edges are forward-only. \square




Figure 2: The future-past automaton for the single-step transduction of Example 22, annotated with “shapes” that show the enabled rules.




Figure 3: Unambiguous CRA for the single-step transduction of Example 22.

4. Relation to Weighted Automata

In this section we investigate the relationship between weighted automata and CRAs. We will show, in particular, that there are CRAs for suitable choice of \mathcal{O} that correspond precisely to unambiguous and nondeterministic weighted automata respectively. This means that CRAs are a more general model of quantitative automata.

The model of weighted automata is a widely studied quantitative extension of nondeterministic finite-state automata [8] that describe (total) functions from Σ^* to a set of data values D . The set D is typically equipped with two operations, product and sum, that form a semiring. In a weighted automaton every transition is labeled with a tag in Σ and a weight in D . The weight of a run of the automaton is the product of all transition weights in the run. The weight associated with an input string w is the sum of the weights of all the runs that are labeled with w .

Formally, a *semiring* (also called a *rig*) is $\mathcal{D} = (D, +, \cdot, 0, 1)$ where $+$

is a commutative, associative binary operation on D with identity 0, \cdot is an associative binary operation on D with identity 1 and absorbing element 0, and \cdot distributes over $+$. A *weighted automaton* over \mathcal{D} is a tuple $\mathcal{A} = (Q, \Sigma, \Delta, I, F)$, where Q is the set of states, Σ is the input alphabet, $\Delta : Q \times \Sigma \times Q \rightarrow D$ is the transition function, $I : Q \rightarrow D$ is the initial weight function, and $F : Q \rightarrow D$ is the final weight function. For a word $u = a_1 a_2 \cdots a_n \in \Sigma^*$, a u -path in \mathcal{A} is a sequence

$$\pi = q_0 \xrightarrow{a_1/d_1} q_1 \xrightarrow{a_2/d_2} q_2 \cdots \xrightarrow{a_n/d_n} q_n,$$

where $\Delta(q_{i-1}, a_i, q_i) = d_i$ for all $i = 1, 2, \dots, n$. The *weight* of the path π is $\text{weight}(\pi) = d_1 \cdot d_2 \cdots d_n$. We denote by $\text{first}(\pi)$ and $\text{last}(\pi)$ the first state q_0 and the last state q_n of the path π respectively. Finally, the *weight* of a word $u \in \Sigma^*$ in \mathcal{A} is given by

$$\text{weight}(u) = \sum_{u\text{-path } \pi} I(\text{first}(\pi)) \cdot \text{weight}(\pi) \cdot F(\text{last}(\pi)).$$

The weighted automaton \mathcal{A} computes the function $\llbracket \mathcal{A} \rrbracket : \Sigma^* \rightarrow D$, where $\llbracket \mathcal{A} \rrbracket(u)$ is the weight of u in \mathcal{A} .

The set of *initial states* of \mathcal{A} is $\{q \in Q \mid I(q) \neq 0\}$, and similarly the set of *final states* of \mathcal{A} is $\{q \in Q \mid F(q) \neq 0\}$. A path π in \mathcal{A} is *successful* if $\text{first}(\pi)$ is initial and $\text{last}(\pi)$ is final. A path is *unsuccessful* if it is not successful. Observe that the weight of an unsuccessful path is equal to 0. A weighted automaton \mathcal{A} is said to be *unambiguous* if for every word $u \in \Sigma^*$ there is at most one successful u -path in \mathcal{A} .

An unambiguous weighted automaton can be over a monoid, instead of a semiring: modify I and F to be partial functions instead of total, where their domains are the initial and final states, respectively. Then removing $+$ and 0, we are left a single associative operation \cdot on D with identity 1. The following theorem shows a correspondence between unambiguous weighted automata over the monoid $(D, \cdot, 1)$ and *copyless* UCRA that can use the multiplication operation \cdot in a certain restricted way: only multiplying by constants on the right.

Theorem 29. Suppose that $(D, \cdot, 1)$ is a monoid, i.e. \cdot is an associative binary operation on D and 1 is a left and right identity for \cdot . Moreover, assume that \mathcal{O} consists of the constant 1 and the family of unary operations $(- \cdot d)$ where $d \in D$. Then, the class of total transductions in $\text{SLR}(\mathcal{O})$ is equal to the class of functions computed by unambiguous weighted automata over $(D, \cdot, 1)$.

Proof. It is easy to see that every unambiguous weighted automaton over $(D, \cdot, 1)$ can be simulated by a copyless single-register UCRA over \mathcal{O} .

It remains to show that the computation of a copyless UCRA \mathcal{A} over \mathcal{O} can be simulated by an unambiguous weighted automaton \mathcal{B} over $(D, \cdot, 1)$. The idea is that at every step of the computation of \mathcal{A} , the content of at most one register can contribute to the final output. So, \mathcal{B} is defined to have state space $Q \times X$, where Q is the state space of \mathcal{A} and X is the set of registers of \mathcal{A} . We also may need additional states to compute the finalization function. \square

Similarly to Theorem 29 we now establish a correspondence between (potentially ambiguous) weighted automata over a semiring $(D, +, \cdot, 0, 1)$ and copyful CRAs that can use the multiplication operation \cdot only to right-multiply by constants.

Theorem 30. Suppose that D is the type of data values and $\mathcal{D} = (D, +, \cdot, 0, 1)$ is a semiring. Moreover, assume that \mathcal{O} consists of the constants $0, 1$, the binary operation $+$, and the family of unary operations $(- \cdot d)$ for every $d \in D$. The class of total transductions of $\text{SR}(\mathcal{O})$ is equal to the class of transductions computed by weighted automata over \mathcal{D} .

Proof. Every weighted automaton $(Q, \Sigma, \Delta, I, F)$ can be simulated by a copyful single-state CRA by considering one register x_q for each state q of the weighted automaton. The registers are used to encode the configuration of the weighted automaton, which is a mapping from the state space to D . For every tag a , we include an a -labeled transition in the CRA that updates the variable x_q with the assignment $x_q := \sum_{p \in Q} x_p \cdot \Delta(p, a, q)$.

It remains to show that the computation of a (copyful, total) CRA can be simulated by a weighted automaton. Given a CRA \mathcal{A} , let us describe now a weighted automaton \mathcal{B} that computes the same transduction as \mathcal{A} . We may algebraically simplify every update of \mathcal{A} to be of the form

$$x = x_1 \cdot d_1 + \cdots + x_n \cdot d_n + d,$$

where x_1, \dots, x_n is an enumeration of the registers of \mathcal{A} . The state space of \mathcal{B} consists of states q and pairs (q, x) , where q is a state of \mathcal{A} and x is a register of \mathcal{A} . The register update function of a transition $q \xrightarrow{a} q'$ of \mathcal{A} can

be presented in the following form:

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}^\top = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}^\top \cdot \begin{bmatrix} d_{11} & d_{12} & \cdots & d_{1n} \\ d_{21} & d_{22} & \cdots & d_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{n1} & d_{n2} & \cdots & d_{nn} \end{bmatrix} + \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{bmatrix}^\top.$$

The above matrix equation can be reformulated as follows:

$$x_j = \left(\sum_{i=1}^n x_i \cdot d_{ij}\right) + d_j \text{ for all } j = 1, \dots, n.$$

The update function is simulated in the weighted automaton \mathcal{B} by defining the following transitions:

$$\Delta((q, x_i), a, (q', x_j)) = d_{ij} \quad \Delta(q, a, (q', x_j)) = d_j \quad \Delta(q, a, q') = 1$$

If q is a final state of \mathcal{A} with finalization function $x_1 \cdot d_1 + \cdots + x_n \cdot d_n + d$, we define the final weight function F of \mathcal{B} by $F(q) = d$ and $F((q, x_i)) = d_i$. If q is not a final state, then we put $F(q) = 0$ and $F(q, x) = 0$. \square

It is important that we have restricted the allowed operations of the CRA to include only right-multiplication by a constant, a unary operation. If a (copyful) CRA can use the binary \cdot operation of the semiring $(D, +, \cdot, 0, 1)$, then it can compute the function $a^i \mapsto d^{2^i}$ (where $a \in \Sigma$ and $d \in D$), which is not implementable by a weighted automaton.

5. Related Work

Weighted automata and extensions. Weighted Automata, which were introduced in 1961 by Schützenberger [14] (see also the more recent monograph [8]), extend classical nondeterministic automata by annotating transitions with *weights* and can be used for the computation of simple quantitative properties on finite or infinite strings of symbols [15]. Weighted automata have found applications in speech and language processing [16], and they have been extended for modeling systems and verifying quantitative properties of these systems [17]. The studied models for the latter purpose include *Nested Weighted Automata* [18], a two-level variant of weighted automata for infinite strings. The computational problems that are relevant for quantitative verification are analysis questions such as universality and equivalence. These

questions are decidable only when the weights and the operations used on them are very simple [19, 20], so the studied models are usually equipped with a very limited set of primitive operations that are insufficient for expressing realistic computations over streaming data. We have shown in Section 4 that CRAs generalize the weighted automata. In particular, when the family \mathcal{O} of available operations is chosen appropriately, CRAs correspond precisely to unambiguous and nondeterministic weighted automata respectively.

Register automata. Another approach to augment classical automata with quantitative features has been with the addition of *registers* that can store values from a potentially infinite set. These models are typically varied in two aspects: by the choice of data types and operations that are allowed for register manipulation, and by the ability to perform tests on the registers for control flow.

The literature on data words, data/register automata and their associated logics [21, 22, 23, 24, 25] studies models that operate on words over an infinite alphabet, which is typically of the form $\Sigma \times \mathbb{N}$, where Σ is a finite set of tags and \mathbb{N} is the set of the natural numbers. They allow comparing data values for equality, and these equality tests can affect the control flow. Due to this feature many interesting decision problems become undecidable [22].

Cost register automata. Cost Register Automata (CRAs) [4, 26] and Streaming String Transducers (SSTs) [27, 28, 9] have recently been proposed as models to permit more complex operations over the cost domain. Unlike the case of register automata, which traditionally only allow checking for equality, CRAs are parameterized by the set of permissible operations \mathcal{O} over the values of the registers. On the other hand, CRAs and SSTs also enforce a strict separation between data and control, so that data values from the input sequence can only affect the contents of the registers but not determine the state transitions. These models have robust expressive power [27, 4], including equivalent characterizations in MSO and closure under various transformations, and have been used in applications such as verifying list processing programs [28]. Work has been devoted to study the circuit complexity of the functions being computed [29, 30], and on analysis problems such as determining the register complexity [26, 31]. Unlike the present paper, these papers allow registers whose contents may include *syntactic terms* rather than concrete values from the data domain, but limit themselves to the copyless setting [32].

MSO-definable graph transformations. An important goal for language theorists has been to characterize operational models such as automata and transducers using properties and functions expressible in various logical theories. The earliest of these results includes the seminal work of Büchi [33] and Elgot [34], establishing the equivalence between regular languages and properties expressible using monadic second-order logic (MSO). Courcelle considered the notion of MSO-definable graph transductions [10]: Engelfriet and Hoogeboom [13] showed the equivalence of string-to-string transformations expressible in this framework with those computable using two-way finite state transducers, and Alur and Černý equated their expressive power with the deterministic one-way model of streaming string transducers (SSTs) [27]. These have also been extended to transformations of infinite strings [35]. Streaming tree transducers [9] and CRAs [4] are similar characterizations of tree-to-tree and string-to-tree transformations respectively. With this background, the present paper shows that copyful CRAs over \mathcal{O} can compute exactly those functions which can be expressed as word-to-DAG transformations in MSO, i.e. the class of streamable regular transductions, $\text{SR}(\mathcal{O})$.

Quantitative regular expressions. The connection between finite automata and regular expressions has motivated research into identifying similar linguistic characterizations of the transductions computable by CRAs and SSTs [1, 36]. The resulting formalisms, *Quantitative Regular Expressions* (QREs) and DReX respectively, have subsequently been applied in processing streaming data (StreamQRE) [2], expressing string transformations [37], and monitoring network traffic (NetQRE) [3]. The closure of the transduction classes SR and SLR under the regular combinators naturally facilitates the design of these declarative languages. These papers are then largely concerned with twin problems of expressiveness—providing translation algorithms to and from CRAs and SSTs—and fast query evaluation. As with previous research on CRAs, they focus on copyless models, and exploit these for time- and memory-efficient one-pass query evaluation. An important technical challenge is the design of a type system which constrains the rate of the transduction and the composition rules to permit efficient modular evaluation. The combinators of Table 1, for example, have been refined in [2] with a type system that eliminates sources of complexity. When the registers of CRAs maintain terms, an important consideration is term *simplification*, so that the contents of each register may be maintained in constant

space, regardless of the length of the input stream seen so far [1]. In this context, developing a regular expression-like characterization of the class of streamable regular transductions, $\text{SR}(\mathcal{O})$, which we consider in this paper, is a problem for future work.

6. Conclusion

We have studied the class of *streamable regular transductions* (SR), which are partial functions from input streams of tagged values to output values that can be computed by unambiguous Cost Register Automata (UCRAs). The subclass of *streamable linear regular transductions* (SLR) consists of those transductions that are computed by *copyless* UCRAs, where the register updates are restricted so that each register appears at most once in the right-hand side of a parallel update. We have shown that the classes SR and SLR have appealing logical characterizations: SR (resp., SLR) corresponds to MSO-definable transformations from strings to DAGs (trees) without backward edges. A precise relationship with the classical model of weighted automata has also been established.

Acknowledgements

This work was supported by NSF award CCF 1763514.

References

- [1] R. Alur, D. Fisman, M. Raghothaman, Regular programming for quantitative properties of data streams, in: Proceedings of the 25th European Symposium on Programming (ESOP '16), 2016, pp. 15–40. doi:10.1007/978-3-662-49498-1_2.
- [2] K. Mamouras, M. Raghothaman, R. Alur, Z. G. Ives, S. Khanna, StreamQRE: Modular specification and efficient evaluation of quantitative queries over streaming data, in: Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '17, ACM, New York, NY, USA, 2017, pp. 693–708. doi:10.1145/3062341.3062369.

- [3] Y. Yuan, D. Lin, A. Mishra, S. Marwaha, R. Alur, B. T. Loo, Quantitative network monitoring with NetQRE, in: Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17, ACM, 2017, pp. 99–112. doi:10.1145/3098822.3098830.
- [4] R. Alur, L. D'Antoni, J. Deshmukh, M. Raghothaman, Y. Yuan, Regular functions and cost register automata, in: Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '13), 2013, pp. 13–22. doi:10.1109/LICS.2013.65.
- [5] R. Bloem, J. Engelfriet, A comparison of tree transductions defined by monadic second order logic and by attribute grammars, Journal of Computer and System Sciences 61 (1) (2000) 1–50. doi:10.1006/jcss.1999.1684.
- [6] J. Engelfriet, S. Maneth, Macro tree transducers, attribute grammars, and MSO definable tree translations, Information and Computation 154 (1) (1999) 34–91. doi:10.1006/inco.1999.2807.
- [7] R. Alur, L. D'Antoni, Streaming tree transducers, Journal of the ACM 64 (5) (2017) 31:1–31:55. doi:10.1145/3092842.
- [8] M. Droste, W. Kuich, H. Vogler (Eds.), Handbook of Weighted Automata, Springer, 2009. doi:10.1007/978-3-642-01492-5.
- [9] R. Alur, L. D'Antoni, Streaming tree transducers, in: Proceedings of the 39th International Colloquium on Automata, Languages, and Programming (ICALP '12), 2012, pp. 42–53. doi:10.1007/978-3-642-31585-5_8.
- [10] B. Courcelle, Monadic second-order definable graph transductions: A survey, Theoretical Computer Science 126 (1) (1994) 53–75. doi:10.1016/0304-3975(94)90268-2.
- [11] W. Thomas, Languages, Automata, and Logic, Springer, 1997, pp. 389–455. doi:10.1007/978-3-642-59126-6_7.
- [12] J. Engelfriet, H. J. Hoogeboom, MSO definable string transductions and two-way finite-state transducers, ACM Transactions on Computational Logic 2 (2) (2001) 216–254. doi:10.1145/371316.371512.

- [13] J. Engelfriet, H. J. Hoogeboom, MSO definable string transductions and two-way finite-state transducers, *ACM Transactions on Computational Logic* 2 (2) (2001) 216–254. doi:10.1145/371316.371512.
- [14] M. P. Schützenberger, On the definition of a family of automata, *Information and control* 4 (2) (1961) 245–270. doi:10.1016/S0019-9958(61)80020-X.
- [15] K. Chatterjee, L. Doyen, T. A. Henzinger, Quantitative languages, *ACM Transactions on Computational Logic (TOCL)* 11 (4) (2010) 23. doi:10.1145/1805950.1805953.
- [16] M. Mohri, Finite-state transducers in language and speech processing, *Computational Linguistics* 23 (2) (1997) 269–311.
- [17] K. Chatterjee, T. A. Henzinger, J. Otop, Quantitative monitor automata, in: X. Rival (Ed.), *Proceedings of the 23rd International Symposium on Static Analysis (SAS '16)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 23–38. doi:10.1007/978-3-662-53413-7_2.
- [18] K. Chatterjee, T. A. Henzinger, J. Otop, Nested weighted automata, in: *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '15)*, IEEE Computer Society, 2015, pp. 725–737. doi:10.1109/LICS.2015.72.
- [19] D. Krob, The equality problem for rational series with multiplicities in the tropical semiring is undecidable, *International Journal of Algebra and Computation* 4 (3) (1994) 405–425. doi:10.1142/S0218196794000063.
- [20] S. Almagor, U. Boker, O. Kupferman, What’s decidable about weighted automata?, in: T. Bultan, P.-A. Hsiung (Eds.), *Proceedings of the 9th International Symposium on Automated Technology for Verification and Analysis (ATVA '11)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 482–491. doi:10.1007/978-3-642-24372-1_37.
- [21] M. Kaminski, N. Francez, Finite-memory automata, *Theoretical Computer Science* 134 (2) (1994) 329–363. doi:10.1016/0304-3975(94)90242-9.

- [22] F. Neven, T. Schwentick, V. Vianu, Finite state machines for strings over infinite alphabets, *ACM Transactions on Computational Logic (TOCL)* 5 (3) (2004) 403–435. doi:10.1145/1013560.1013562.
- [23] S. Demri, R. Lazić, LTL with the freeze quantifier and register automata, *ACM Transactions on Computational Logic (TOCL)* 10 (3) (2009) 16:1–16:30. doi:10.1145/1507244.1507246.
- [24] H. Björklund, T. Schwentick, On notions of regularity for data languages, *Theoretical Computer Science* 411 (4) (2010) 702–715, *fundamentals of Computation Theory*. doi:10.1016/j.tcs.2009.10.009.
- [25] M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, L. Segoufin, Two-variable logic on data words, *ACM Transactions on Computational Logic (TOCL)* 12 (4) (2011) 27:1–27:26. doi:10.1145/1970398.1970403.
- [26] R. Alur, M. Raghothaman, Decision problems for additive regular functions, in: F. V. Fomin, R. Freivalds, M. Kwiatkowska, D. Peleg (Eds.), *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming, Part II (ICALP 2013)*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 37–48. doi:10.1007/978-3-642-39212-2_7.
- [27] R. Alur, P. Cerný, Expressiveness of streaming string transducers, in: K. Lodaya, M. Mahajan (Eds.), *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, Vol. 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2010, pp. 1–12. doi:10.4230/LIPIcs.FSTTCS.2010.1.
- [28] R. Alur, P. Černý, Streaming transducers for algorithmic verification of single-pass list-processing programs, in: *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '11*, ACM, New York, NY, USA, 2011, pp. 599–610. doi:10.1145/1926385.1926454.
- [29] E. Allender, I. Mertz, Complexity of regular functions, in: A.-H. Dediu, E. Formenti, C. Martín-Vide, B. Truthe (Eds.), *Language and Automata Theory and Applications*, Springer International Publishing, Cham, 2015, pp. 449–460.

- [30] E. Allender, A. Krebs, P. McKenzie, Better Complexity Bounds for Cost Register Automata, in: K. G. Larsen, H. L. Bodlaender, J.-F. Raskin (Eds.), 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017), Vol. 83 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017, pp. 24:1–24:14. doi:10.4230/LIPIcs.MFCS.2017.24.
- [31] L. Daviaud, P.-A. Reynier, J.-M. Talbot, A generalised twinning property for minimisation of cost register automata, in: Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, ACM, New York, NY, USA, 2016, pp. 857–866. doi:10.1145/2933575.2934549.
- [32] F. Mazowiecki, C. Riveros, On the expressibility of copyless cost register automata, CoRR abs/1504.01709.
- [33] B. J. Richard, Weak secondorder arithmetic and finite automata, *Mathematical Logic Quarterly* 6 (16) 66–92. doi:10.1002/malq.19600060105.
- [34] C. C. Elgot, Decision problems of finite automata design and related arithmetics, *Transactions of the American Mathematical Society* 98 (1) (1961) 21–51.
- [35] R. Alur, E. Filiot, A. Trivedi, Regular transformations of infinite strings, in: Proceedings of the 2012 27th Annual IEEE/ACM Symposium on Logic in Computer Science, LICS '12, IEEE Computer Society, Washington, DC, USA, 2012, pp. 65–74. doi:10.1109/LICS.2012.18.
- [36] R. Alur, A. Freilich, M. Raghothaman, Regular combinators for string transformations, in: Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, ACM, New York, NY, USA, 2014, pp. 9:1–9:10. doi:10.1145/2603088.2603151.
- [37] R. Alur, L. D'Antoni, M. Raghothaman, Drex: A declarative language for efficiently evaluating regular string transformations, in: Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles

of Programming Languages, POPL '15, ACM, New York, NY, USA,
2015, pp. 125–137. doi:10.1145/2676726.2676981.