

SCULPT: a Schema Language for Tabular Data on the Web

Wim Martens
Universität Bayreuth
wim.martens@uni-bayreuth.de

Frank Neven
Hasselt University and
transnational University of
Limburg
frank.neven@uhasselt.be

Stijn Vansummeren
Université Libre de Bruxelles
stijn.vansummeren@ulb.ac.be

ABSTRACT

Inspired by the recent working effort towards a recommendation by the World Wide Web Consortium (W3C) for tabular data and metadata on the Web, we present in this paper a concept for a schema language for tabular web data called SCULPT. The language consists of rules constraining and defining the structure of regions in the table. These regions are defined through the novel formalism of region selection expressions. We present a formal model for SCULPT and obtain a linear time combined complexity evaluation algorithm. In addition, we consider weak and strong streaming evaluation for SCULPT and present a SCULPT fragment for each of these streaming variants. Finally, we discuss several extensions of SCULPT including alternative semantics, types, complex content, and explore region selection expressions as a basis for a transformation language.

1. INTRODUCTION

Despite the availability of numerous standardized formats for semi-structured and semantic web data such as XML, RDF, and JSON, a very large percentage of data and open data published on the web, remains tabular in nature.¹ Tabular data is most commonly published in the form of comma separated values (CSV) files because such files are open and therefore processable by numerous tools, and tailored for all sizes of files ranging from a number of KBs to several TBs. Despite these advantages, working with CSV files is often cumbersome because they are typically not accompanied by a *schema* that describes the file's structure (i.e., “the second column is of integer datatype”, “columns are delimited by tabs”, ...) and captures its intended meaning. Such a description is nevertheless vital for any user trying to interpret the file and execute queries or make changes to it. In other data models, the presence of a schema is also important for query optimization (required for scalable query execution if

¹Jeni Tennison, one of the two co-chairs of the W3C CSV on the Web working group claims that “over 90% of the data published on data.gov.uk is tabular data” [28].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

the file is large), as well as other static analysis tasks. Finally, we strongly believe that schemas are a prerequisite for unlocking huge amounts of tabular data to the Semantic Web. Indeed, unless we have a satisfactory way of describing the structure of tabular data we cannot specify how its contents should be interpreted as RDF. Drawing parallels with relational databases, observe that R2RML mappings [8] (the W3C standard for mapping relational databases to RDF) inherently need to refer to the schema (structure) of the relational database in order to specify how database tuples can be represented as RDF.

In recognition of this problem, the *CSV on the Web* Working Group of the World Wide Web Consortium [31] argues for the introduction of a schema language for tabular data to ensure higher interoperability when working with datasets using the CSV or similar formats. In particular, their charter states [31]:

Whether converted to other formats or not, there is a need to describe the content of the CSV file: its structure, datatypes used in a specific column, language used for text fields, access rights, provenance, etc. This means that metadata should be available for the dataset, relying on standard vocabulary terms, and giving the necessary information for applications. The metadata can also be used for the conversion of the CSV content to other formats like RDF or JSON, it can enable automated loading of the data as objects, or it can provide additional information that search engines may use to gain a better understanding of the content of the data.

In the present paper, we introduce SCULPT as a concept for such a schema language for tabular data.²

The critical reader may wonder whether designing such a schema language isn't trivial. After all, doesn't it suffice to be able to specify, for each column, the column's name and the type of data allowed in its cells—similar to how relational database schemas are defined using the SQL data definition language? The answer is no. The reason is that there is a lot of variation in the tabular data available on the web and that there are examples abound of tabular data whose structure cannot be described by simple rules of the

²The name SCULPT for the language is in honour of Michelangelo, who allegedly said “*Every block of stone has a statue inside it and it is the task of the sculptor to discover it.*” Readers who like acronyms can read SCULPT as **S**chema for **U**n-**L**ocking and **P**rocessing **T**abular data.

form “column x has datatype y ”. Figures 3, 5, and 7, for example, show some tabular data sets drawn from the Use Cases and Requirements document drafted by the W3C CSV on the Web working group [27]. Notice how, in contrast to “standard” CSV files, Figure 3 has a header consisting of multiple lines. This causes the data in the first column to be non-uniform. Further notice how the `provenance` data in the Figure 5 is spread among multiple columns. Finally, notice how the shape of the rows in Figure 7 depends on the label in the first column of the column: `TITLE` rows have different structure than `AUTHOR` rows, which have a different structure than `ATOM` rows, and so on.

SCULPT schemas use the following idea to describe the structure of these tables. At their core, SCULPT schemas consist of rules of the form $\varphi \rightarrow \rho$. Here, φ selects a *region* in the input table (i.e., a subset of the table’s cells) and ρ constrains the allowed structure and content of this region. A table is valid with respect to a SCULPT schema if, for each rule in the schema, the region selected by φ satisfies the content constraints specified by ρ . It is important to note that SCULPT’s expressive power goes well beyond that of classical relational database schemas since SCULPT’s region selectors are not limited to selecting columns. In particular, the language that we propose for selecting regions is capable of navigating through a table’s cells bears much resemblance to the way XPath [4] navigates through the nodes of an XML tree. For tokenizing the content of single cells, we draw inspiration from XML Schema simple types ([9], Section 2.2). Both features combined will allow us to express the use cases of the W3C CSV on the Web Working Group.

We note that the W3C is also working on a schema language for tabular data [22]. At the moment, however, that schema language focuses on orthogonal issues like describing, for instance, *datatypes* and *parsing cells*. Also, it only provides facilities for the selection of *columns*, and is hence not able to express the schema of the more advanced use cases. SCULPT, in contrast, draws inspiration from well-established theoretical tools from logic and formal languages, which adds to the robustness of our approach. Due to the above mentioned orthogonality we expect that it is not difficult to integrate ideas from this paper in the W3C proposal.

In summary, we make the following contributions.

1. We illustrate the power of SCULPT, and its suitability as a schema language for tabular data on the web, by expressing several use cases drafted by the CSV on the Web W3C working group [27]. (Section 2)
2. We provide a formal model for the core of SCULPT. A key contribution in this respect is the introduction of the region selector language. (Section 3)
3. We show that, despite its rather attractive expressiveness, tables can be efficiently validated w.r.t. SCULPT schemas. In particular, when the table is small enough to be materialized in main memory, we show that validation can be done in linear time combined complexity (Section 4.1). For scenarios where materialization in main memory is not possible, we consider the scenario of streaming (i.e., incremental) validation. We formally introduce two versions of streaming validation: *weak streamability* and *strong streamability*. (Their differences are described in detail in Section 4.2.) We show in particular that the fragment of core-SCULPT where region selectors can only look “forward” and never “backward” in the CSV file is *weakly streamable*. If we further restrict region selectors to be

both forward-looking and *guarded* (a notion formalized in Section 4.2) validation becomes *strongly streamable*. All of the W3C Working group use cases considered here can be expressed using forward and guarded region selectors, hence illustrating the practical usefulness of this fragment.

4. While our focus in this paper is on introducing SCULPT as a means for specifying the structure of CSV files and related formats, we strongly believe that region selector expressions are a fundamental component in developing other features mentioned in the charter of the W3C CSV on the Web Working Group, such as a CSV transformation language (for converting tabular data into other formats such as RDF or JSON), the specification of the language used for text fields; access rights; provenance; etc. While a full specification of these features is out of this paper’s scope, we illustrate by means of example how SCULPT could be extended to incorporate them. (Section 5)

Note. Due to space restrictions, proofs of formal statements are only sketched. Proofs are provided in the Appendix.

Related Work. The present paper fits in the line of research, historically often published in the WWW conference, that aims to formalize and study the properties of various W3C working group drafts and standards (including XML Schema [5,6], SPARQL [3,18,21], and RDF [23,24]) with the aim of providing feedback and input to the working group’s activities.

Given the numerous benefits of schemas for data processing, there is a large body of work on the development, expressiveness, and properties of schema languages for virtually all data models, including the relational data model, XML [5,6,12,19,20], and, more recently, RDF [23,24]. SCULPT differs from the schema languages considered for XML and RDF in that it is specifically designed for tabular data, not tree-structured or graph-structured data. Nevertheless, the rule-based nature of SCULPT draws inspiration from our prior work on rule-based and pattern-based schema languages for XML [12,19,20].

As already mentioned, while traditional relational database schemas (formulated in e.g., the SQL data definition language) are specifically designed for tabular data, they are strictly less expressive than SCULPT schemas in the sense that relational schemas limit region selection expressions to those that select columns only. A similar remark holds for other recent proposals of CSV schema languages, including the CSV Schema language proposed by the UK National Archives [1], and Tabular Data Package [16]. The remark also applies to the part of Google’s Dataset Publishing Language (DSPL) [13] describing the contents of CSV files. In contrast, DSPL also has features to relate data from multiple CSV files, which SCULPT does not yet have.

The problem of streaming schema validation has been investigated in the XML context for DTDs and XML schemas [14,20,25,26]. In this work, the focus is on finding algorithms that can validate an XML document in a single pass using constant memory or, if this is not possible, a memory that is bounded by the depth of the document. Our notion of streaming, in contrast, is one where we can use a memory that is not constant but at most logarithmic in the size of the table (for strong streaming), or at most linear in the number of columns and logarithmic in the number of rows (for weak streaming). This allows us to restrict memory when

going from one row to the next and is essential to be able to navigate downwards in SCULPT region selection expressions.

While streaming validation is undoubtedly an important topic for all of the CSV schema languages mentioned above [1, 13, 16] (the National Archives Schema Language mentions it as an explicit design goal), no formal streaming validation algorithm has been proposed for them, to the best of our knowledge.

2. SCULPT BY EXAMPLE

In this section, we introduce SCULPT through a number of examples. The formal semantics of the examples is defined in Section 3. The syntax we use here is tuned for making the examples accessible to readers and is, of course, flexible.

SCULPT schemas operate on *tabular documents*, which are text files describing tabular data. SCULPT schemas consist of two parts (cf. Figure 2). The first part, *parsing information*, defines the row and column delimiters and further describes how words should be tokenized. This allows to parse the text file and build a table-like structure consisting of rows and columns. In this section we allow some rows to have fewer columns than others but we require them to be aligned to the left. That is, non-empty rows always have a first column. The second part of the schema consists of *rules* that interpret the table defined by the first part as a rectangular grid and enforce structure. In particular, rules are of the form $\varphi \rightarrow \rho$, where φ selects a *region* consisting of cells in the grid while ρ is a regular expression constraining the content of the selected region. We utilize a so-called *row-based semantics*: every row in the region selected by φ should be of a form allowed by ρ . We refer to φ as the *selector expression* and to ρ as the *content expression*.

Next, we illustrate the features of the language by means of examples. All examples are inspired by the use cases and requirements drafted by the CSV on the Web W3C working group [27].

EXAMPLE 2.1. Figure 1 contains a slightly altered fragment (we use a comma as a column separator) of a CSV file mentioned in Use Case 3, “*Creation of consolidated global land surface temperature climate databank*” [27]. The SCULPT schema, displayed as Figure 2, starts by describing parsing information indicating that the column delimiter is a comma while the row delimiter is a newline. Lines starting with a %-sign are comments. Tokens are defined based on regular expressions (regex for short).³ For instance, anything that matches the regex `[0-9]{4}."[0-9]{2}` follows the format *four digits, dot, two digits*, and is interpreted by the token `Timestamp` in the rules of the schema (similar for `Temperature`). Notice that we keep the regexes short (and sometimes imprecise) for readability, but they can of course be made arbitrarily precise if desired.

All the XML Schema primitive types like `xs:integer`, `xs:string`, `xs:date`, etc are pre-defined as tokens in a SCULPT schema. There is also a special pre-defined token `Empty` to denote that a certain cell is empty.

Notice that the schema in Figure 2 has three token definitions in which the regex defines only one character sequence (namely: `AURA`, `BOMBO`, `ENTEBBE AIR`). In the sequel, we will omit such rules for reasons of parsimony. For the

³For ease of exposition, we adopt the concise regex syntax popularized by scripting languages such as Perl, Python, and Ruby [11] in all of our examples.

	,	ARUA,	BOMBO,	ENTEBBE AIR
1935.04,	-99.00,	-99.00,	27.83	
1935.12,	-99.00,	-99.00,	25.72	
1935.21,	-99.00,	-99.00,	26.44	
1935.29,	-99.00,	-99.00,	25.72	
1935.37,	-99.00,	-99.00,	24.61	
1935.46,	-99.00,	-99.00,	24.33	
1935.54,	-99.00,	-99.00,	24.89	

Figure 1: Example tabular data inspired by Use Case 3 in [27].

```

% Parsing information
%% Delimiters
Col Delim = ,
Row Delim = \n

%% Tokens
%% left: token name
%% right: regex

Timestamp = [0-9]{4}."[0-9]{2}
Temperature = (-)?[0-9]{2}."[0-9]{2}
ARUA = ARUA
BOMBO = BOMBO
ENTEBBE AIR = ENTEBBE AIR

% Rules

row(1) -> Empty, ARUA, BOMBO, ENTEBBE AIR
col(1) -> Empty | Timestamp
col(ARUA) -> Temperature
col(BOMBO) -> Temperature
col(ENTEBBE AIR) -> Temperature

```

Figure 2: Schema for tabular data of the type in Figure 1.

same reason, we omit the explicit definition of column and row delimiters when they are a comma and newline character, respectively.

The rule

```
row(1) -> Empty, ARUA, BOMBO, ENTEBBE AIR
```

selects all cells in the first row and requires that the first is empty, the second contains `ARUA`, the third `BOMBO`, and the fourth `ENTEBBE AIR`. Next, `col(1)` selects the region consisting of all cells in the first column. As SCULPT assumes a row-based semantics per default,⁴ the rule

```
col(1) -> Empty | Timestamp
```

requires that every row in the selected region (notice that each such row consists of a single cell) is either empty (`Empty`) or contains data that matches the `Timestamp` token. The expression `col(ARUA)` selects all cells in the column below the cell containing `ARUA`. The rule

```
col(ARUA) -> Temperature
```

therefore requires that every row in the selected region matches the `Temperature` token. The two remaining rules are analogous. The fragment in Figure 1 satisfies the schema of

⁴We discuss an extension in Section 5.

```

QS601EW
Economic activity
27/03/2011

,      , Count  , Count
,      , Person , Person
,      , Activity, Activity
GeoID  , GeoArea, All    , Part-time
E92000001, England, 38881374, 27183134
W92000004, Wales  , 2245166  , 1476735

```

Figure 3: Fragment of a CSV-like-file, inspired by Use Case 2 in [27].

```

%% Tokens
%% left: token name
%% right: regex

name = QS[0-9]*EW
ctype = Economic Activity
geo_id = E[0-9]*

% Rules

row(1) -> name
row(2) -> ctype
row(3) -> Date
row(4) -> Empty
row(5) -> Empty, Empty, Count*
row(6) -> Empty, Empty, Person*
row(7) -> Empty, Empty, Activity*
row(8) -> GeoID, GeoArea, String*
col(GeoID) -> geo_id
col(GeoArea) -> String
down+(right+(GeoArea)) -> Number*

```

Figure 4: “Schema” for files of the type in Figure 3.

Figure 2. ■

Before moving on to some more advanced examples, we discuss in more detail the semantics of selector and content expressions. Each cell in a table is identified by its *coordinate*, which is a pair (k, ℓ) where k indicates the row number ($k \geq 1$) and ℓ the column number ($\ell \geq 1$). In each rule $\varphi \rightarrow \rho$, the selector expression φ returns a *set* of coordinates (a region) and ρ is a regular expression defining the allowed structure of each row in the region selected by φ . It is important to note that in each such row only the cells which are selected by φ are considered. Another way to interpret the row-based semantics is that of a ‘group by’ on the selected region per row.

The last rule we discussed in Example 2.1 uses a symbolic coordinate **ARUA** in its selector expression. Its semantics is as follows: a token τ returns the set of all coordinates (k, ℓ) whose cell contents matches τ . The operator **row** applied to a coordinate (k, ℓ) returns the set of coordinates $\{(k, \ell') \mid \ell' > \ell\}$. This corresponds to the row consisting of all elements to the right of (k, ℓ) . Note that coordinate (k, ℓ) itself is not included. Applying **row** to a *set* S of coor-

dinates amounts to taking the union of all $\text{row}((k, \ell))$ where $(k, \ell) \in S$. Similarly, the operator **col** applied to S returns the union of the regions $\{(k', \ell) \mid k' > k\}$ for each (k, ℓ) in S , corresponding to columns below elements in S . The selector expressions **row(1)** and **col(1)** that select the “first row” and “first column”, respectively, use syntactic sugar to improve readability. Formally, the notation **row(k)** and **col(l)** abbreviate **row** and **col** applied to $\{(k, 0)\}$ and $\{(0, \ell)\}$, respectively. Using the same principle as above, this means that **row(k)** selects the cells $\{(k, \ell) \mid \ell > 0\}$ and **col(l)** selects $\{(k, \ell) \mid k > 0\}$. Notice that we use the convention that the top left coordinate in tabular data bears the coordinate $(1, 1)$ — for *first row, first column*. While the value 0 does not refer to any cell in the table, it is used to define the semantics of expressions.

The next example illustrates the use of slightly more complex expressions for navigation and content.

EXAMPLE 2.2. Figure 3 displays a (slightly altered) fragment of a CSV-like-file inspired by Use Case 2 (“*Publication of National Statistics*”) in [27]. This fragment originates from the Office for National Statistics (UK) and refers to the dataset “QS601EW Economic activity” derived from the 2011 Census. The file starts with three lines of metadata, referring to the name of the file and the census date, continues with a blank line, before listing the actual data separated by commas. Notice that this file is, strictly speaking, not a comma-separated-value file because not all rows have an equal number of columns.⁵ Indeed, the first four rows have only (at most) one column and the later rows have four columns. Figure 4 depicts the SCULPT schema describing the structure of such tables.

The schema starts by describing parsing information, analogous to Example 2.1. The first four rules are very basic and are similar to those of Example 2.1. We first describe the fifth rule of the schema:

```
row(5) -> Empty, Empty, Count*
```

selects all cells in the fifth row, requiring the first two to be **Empty** and the remaining non-empty cells to contain **Count**. We note that the original data fragment from [27] contains 16 such columns. The remaining rules constraining rows are similar.

The rule **col(GeoID) -> geo_id** selects all cells below cells containing the word **GeoID**. The content expression says that this column contains values that match the **geo_id** token. The last rule is the most interesting one:

```
down+(right+(GeoArea)) -> Number*.
```

This rule selects all cells appearing strictly downward and to the right of **GeoArea** and requires them to be of type **Number**. More precisely, **GeoArea** is a symbolic coordinate selecting all cells containing the word **GeoArea**. The navigational operators **right** and **down** select cells one step to the right and one step down, respectively, from a given coordinate. The operator **+** indicates an arbitrary strictly positive number of applications of the navigational operator to which it is applied. In particular, as on the table given in Figure 3, **GeoArea** is the singleton cell with coordinate $(8, 2)$, **right(GeoArea)** returns $\{(8, 3)\}$, while **right+(GeoArea)** is the region $\{(8, \ell) \mid \ell > 2\}$. Likewise, **down(right+(GeoArea))** is the region $\{(9, \ell) \mid \ell > 2\}$ and, finally, **down+(right+(GeoArea))** is the

⁵Actually CSV does not have a standard, but the informative memo RFC4180 (<http://tools.ietf.org/html/rfc4180>) states rectangularity in paragraph 2.4.

subject	predicate	object	provenance
:e4	type	PER	
:e4	mention	"Bart"	D00124 283-286
:e4	mention	"JoJo"	D00124 145-149 0.9
:e4	per:siblings	:e7	D00124 283-286 173-179 274-281
:e4	per:age	"10"	D00124 180-181 173-179 182-191 0.9
:e4	per:parent	:e9	D00124 180-181 381-380 399-406 D00101 220-225 230-233 201-210

Figure 5: Fragment of a CSV-like file, inspired by Use Case 13 in [27].

region downward and to the right of the `GeoArea` coordinate, that is, $\{(k, \ell) \mid k > 8 \text{ and } \ell > 2\}$. ■

Example 2.2 uses more refined navigation than just selecting a row or a column. SCULPT has four navigational axes: `up`, `down`, `left`, `right` which navigate one cell upward, downward, leftward, or rightward. These axes can be applied to a set S of coordinates and add a vector v to it. More formally, an axis A , when applied to a set S of coordinates, returns $A(S) := \{c + v_A \mid c \in S\}$. Here,

- $v_A = (-1, 0)$ when $A = \text{up}$,
- $v_A = (1, 0)$ when $A = \text{down}$.
- $v_A = (0, 1)$ when $A = \text{right}$, and
- $v_A = (0, -1)$ when $A = \text{left}$.

Furthermore, there is also an axis `cell` that does not navigate away from the current cells, i.e., $\text{cell}(S) = S$. When applying an axis to a set of coordinates, we always return only the coordinates that are valid coordinates in the table. For example, `left({1,1})` returns the empty set because $(1, 0)$ is not a cell in the table.

While the just discussed features of SCULPT are sufficient to describe the structure of almost all CSV-like data on the Web Working group use cases [27], we extend in Section 3 SCULPT to include XPath-like navigation. These features will be useful for annotations and transformations, see Section 5. We now showcase SCULPT by illustrating it on the most challenging of the W3C use cases.

EXAMPLE 2.3. Figure 5 contains a fragment of a CSV-like file, inspired by Use Case 13 in [27] (“*Representing Entities and Facts Extracted From Text*”). Figure 6 depicts the SCULPT schema. Compared to the previous examples, the most interesting rule is

```
down+(right*(provenance))
-> (prov-book, prov-pos*, prov-node?)*
```

which states that every row that starts with a coordinate of the form $(k, 4)$ (as `provenance` only occurs in column 4) with $k > 1$ should match `(prov-book, prov-pos*, prov-node?)*`. Notice that the empty row starting at $(2, 4)$ also matches this expression. Here, `*` denotes an arbitrary number (including zero) of applications of the navigational operator. ■

3. FORMAL MODEL

In this section, we present a formal model for the logical core of SCULPT. We refer to this core as `core-SCULPT` and discuss extensions in Section 5. We first define the data model.

Tables. For a number $n \in \mathbb{N}$, we denote the set $\{1, \dots, n\}$ by $[n]$. By \perp we denote a special distinguished null value. For any set \mathcal{V} , we denote the set $\mathcal{V} \cup \{\perp\}$ by \mathcal{V}_\perp . The W3C formalizes tabular documents through *tables*, which can be defined as follows.

```
% Tokens
%% left: token name
%% right: regex

rdf-id    = [a-zA-Z0-9]*:[a-zA-Z0-9]*
rdf-lit   = "[a-zA-Z0-9]*"
prov-book = D[0-9]{5}
prov-pos  = [0-9]{3}-[0-9]{3}
prov-node = [0-9].[0.9]
word      = [a-z]*

% Rules
row(1) -> subject,predicate,object, provenance
col(subject) -> rdf-id
col(predicate) -> word | rdf-id
col(object) -> rdf-lit | rdf-id
down+(right*(provenance))
-> (prov-book, prov-pos*, prov-node?)*
```

Figure 6: Schema for files of the type in Figure 5.

DEFINITION 3.1 (CORE TABULAR DATA MODEL, [29]). Let \mathcal{V} be a set. A *table* over \mathcal{V} is an $n \times m$ matrix T (for some $m, n \in \mathbb{N}$) in which each cell carries a value from \mathcal{V}_\perp . We say that T has n rows and m columns. A (table) *coordinate* is an element of $[n] \times [m]$. A *cell* is determined by coordinate $(k, \ell) \in [n] \times [m]$ and its *content* is the value $T_{k,\ell} \in \mathcal{V}_\perp$ at the intersection of row k and column ℓ . We denote the *coordinates* of T by $\text{coords}(T)$.

Tabular documents. Notice that tables are always rectangular⁶ whereas, in Section 2, this was not the case for some of the use cases. We model this by padding shorter rows by \perp . More precisely, we see the correspondence between *tabular documents*, i.e., text files that describe tabular data (like CSV files), and tables as follows. Let Σ be a finite alphabet and let D be a set of *delimiters*, disjoint from Σ . We assume that D contains two designated elements which we call *row delimiter* and *column delimiter*, which, as the name indicates, separate cells vertically or horizontally. (We discuss other delimiters in Section 5.) Therefore, a sequence of symbols in $(D \cup \Sigma)^*$ can be seen as a table over Σ^* : every row delimiter induces a new row in the table, every column delimiter a new column, and the Σ -strings between delimiters define the cell contents. In the case that some rows have fewer columns than others, missing columns are expanded to the right and filled with \perp . Conversely, a table over Σ^*

⁶Tables are required to be rectangular by Section 2.1 of [29]; as by paragraph 2.4 of the memo RFC4180 on CSV (<http://tools.ietf.org/html/rfc4180>).

can also be seen as a string over $(D \cup \Sigma)^*$ by concatenating all its cell values in top-down left-to-right order and inserting cell delimiters and row delimiters in the correct places; we do not insert column delimiters next to \perp -cells. As such, when we convert a tabular document into a table and back; we obtain the original tabular document.

We consider both representations in the remainder of the paper. In particular we view the table representation as a structure that allows efficient navigation in all directions and the string representation as structure for streaming validation. **core-SCULPT schemas.** Abstractly speaking, a core-SCULPT schema S is a 4-tuple (D, Δ, Θ, R) where D is the set of delimiters; Δ is a finite set of tokens; Θ is a mapping that associates a regular expression over Σ to each token $\tau \in \Delta$; and R is a *tabular schema*, a set of rules that constrain the admissible table contents (further defined below).

Checking whether a tabular document σ in $(D \cup \Sigma)^*$ satisfies S proceeds conceptually in three phases. In the first phase, the delimiters are used to parse σ into a table T^{raw} over Σ^* , as described above. In the second phase, the token definitions Θ are used to transform T^{raw} into a *tokenized* table T , which is a table where each cell contains a set of tokens (i.e., each cell contains a subset of Δ , namely those tokens that match the cell). Formally, T is the table of the same dimension as T^{raw} such that

$$T_{k,\ell} = \{\tau \in \Delta \mid T_{k,\ell}^{\text{raw}} \in \mathcal{L}(\Theta(\tau))\}.$$

Here $\mathcal{L}(\cdot)$ denotes the language of a regular expression. Finally, the rules in R check validity of the tokenized table T (and not of the raw table T^{raw}), as explained next.

Tabular schema. The *tabular schema* R describes the structure of the tokenized table. Intuitively, a tabular schema is a set of rules $s \rightarrow c$ in which s *selects* a region in the table and c describes what the *content* of the selected region should be. More formally, a *region* z of an $n \times m$ table T is a subset of $[n] \times [m]$. A *region selection language* \mathcal{S} is a set of expressions such that every $s \in \mathcal{S}$ defines a region in every table T . More precisely, $s[T]$ is always a (possibly empty) region of T . A *content language* \mathcal{C} is a set of expressions such that every $c \in \mathcal{C}$ maps each region z of T to true or false. We denote by $T, z \models c$ that c maps z to true in T and say that z *satisfies* c in T .

DEFINITION 3.2 (TABULAR SCHEMA). A (*tabular*) *schema* (over \mathcal{S} and \mathcal{C}) is a set R of rules $s \rightarrow c$ for which $s \in \mathcal{S}$ and $c \in \mathcal{C}$. A table T *satisfies* R , denoted $T \models R$, when for every rule $s \rightarrow c \in R$ we have that $T, s[T] \models c$.

The above definition is very general as it allows arbitrary languages for selecting regions and defining content. We now propose concrete languages for these purposes.

Region selection expressions. Our region selection language is divided into two sorts of expressions: *coordinate expressions* (ranged over by φ, ψ) and *navigational expressions* (ranged over by α, β), defined by the following syntax:

$$\begin{aligned} \varphi, \psi &:= a \mid \text{root} \mid \text{true} \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \neg \varphi \mid \langle \alpha \rangle \mid \alpha(\varphi) \\ \alpha, \beta &:= \varepsilon \mid \text{up} \mid \text{down} \mid \text{left} \mid \text{right} \mid [\varphi] \mid (\alpha \cdot \beta) \mid (\alpha|\beta) \mid (\alpha^*) \end{aligned}$$

Here, a ranges over tokens in Δ and root is a constant referring to coordinate $(1, 1)$. When evaluated over an $n \times m$ table T over 2^Δ , a coordinate expression φ defines a region

$\llbracket \varphi \rrbracket_T \subseteq \text{coords}(T)$, whereas a navigational expression α defines a function $\llbracket \alpha \rrbracket : 2^{\text{coords}(T)} \rightarrow 2^{\text{coords}(T)}$, as follows.

$$\begin{aligned} \llbracket a \rrbracket_T &:= \{(i, j) \in [n] \times [m] \mid a \in T_{i,j}\} \\ \llbracket \text{root} \rrbracket_T &:= \{(1, 1)\} \\ \llbracket \text{true} \rrbracket_T &:= [n] \times [m] \\ \llbracket (\varphi \vee \psi) \rrbracket_T &:= \llbracket \varphi \rrbracket_T \cup \llbracket \psi \rrbracket_T \\ \llbracket (\varphi \wedge \psi) \rrbracket_T &:= \llbracket \varphi \rrbracket_T \cap \llbracket \psi \rrbracket_T \\ \llbracket (\neg \varphi) \rrbracket_T &:= ([n] \times [m]) \setminus \llbracket \varphi \rrbracket_T \\ \llbracket \langle \alpha \rangle \rrbracket_T &:= \{c \in \text{coords}(T) \mid \llbracket \alpha(\{c\}) \rrbracket_T \neq \emptyset\} \\ \llbracket \alpha(\varphi) \rrbracket_T &:= \llbracket \alpha(\llbracket \varphi \rrbracket_T) \rrbracket_T \end{aligned}$$

Furthermore, for every set of coordinates $C \subseteq \text{coords}(T)$,

$$\begin{aligned} \llbracket \varepsilon(C) \rrbracket_T &:= C \\ \llbracket \text{up}(C) \rrbracket_T &:= \{(i-1, j) \mid (i, j) \in C, i > 1\} \\ \llbracket \text{down}(C) \rrbracket_T &:= \{(i+1, j) \mid (i, j) \in C, i < n\} \\ \llbracket \text{left}(C) \rrbracket_T &:= \{(i, j-1) \mid (i, j) \in C, j > 1\} \\ \llbracket \text{right}(C) \rrbracket_T &:= \{(i, j+1) \mid (i, j) \in C, j < m\} \\ \llbracket [\varphi](C) \rrbracket_T &:= C \cap \llbracket \varphi \rrbracket_T \\ \llbracket (\alpha \cdot \beta)(C) \rrbracket_T &:= \llbracket \beta(\llbracket \alpha(C) \rrbracket_T) \rrbracket_T \\ \llbracket (\alpha|\beta)(C) \rrbracket_T &:= \llbracket \alpha(C) \rrbracket_T \cup \llbracket \beta(C) \rrbracket_T \\ \llbracket (\alpha^*)(C) \rrbracket_T &:= \bigcup_{i \geq 0} \llbracket \alpha^i(C) \rrbracket_T \end{aligned}$$

Here, $\alpha^i(C)$ abbreviates the i -fold composition $\alpha \cdots \alpha(C)$. We also use this abbreviation in the remainder. Notice that every coordinate (k, ℓ) of T can be expressed as $\text{down}^{k-1} \cdot \text{right}^{\ell-1}(\text{root})$. For navigational expressions α , we abbreviate $\alpha \cdot \alpha^*$ by α^+ and $\alpha|\varepsilon$ by $\alpha?$. One can read $\alpha(\varphi)$ as “apply the regular expression α to φ ”. The definition of the semantics of $\alpha \cdot \beta$ is conform with this view.

EXAMPLE 3.3. Region selection expressions navigate in tables, similar to how XPath expressions navigate on trees. For example, assuming `dummy` to be a token for `-99.00` in Figure 1, the expression

$$\text{right}^+(\text{root}) \wedge \neg(\text{up}^*(\text{dummy}))$$

selects the top cells of columns that do not contain a dummy value anywhere. In the excerpt of Figure 1, this expression hence selects the cell containing `ENTEbbe AIR`.

Assuming the token `literal` for cells with quotation marks (`regex \\"[a-zA-Z0-9]\\"`) in Figure 5, the expression

$$\text{down}^+ \cdot [\text{literal}] \cdot \text{right}^+(\text{object})$$

selects all provenance information for rows in which the `object` is between quotes, like `"Bart"`, `"JoJo"`, and `"10"`. Notice in particular that the semantics of the operator “[]” in navigational expressions is the same as filter-expressions in XPath. ■

Readers familiar with propositional dynamic logic [10] (PDL for short) will recognize that the above language is nothing more than propositional dynamic logic, tweaked to navigate in tables.

There are some differences between the syntax of core-SCULPT and the region selection expressions used in the examples of Section 2. In particular, the latter examples use the following abbreviations.

REMARK 3.4. (i) As already observed, absolute coordinates in Section 2 are syntactic sugar for navigations that start at the root. For example, the coordinate $(2, 2)$ would be unfolded to $\text{down} \cdot \text{right}(\text{root})$ in core-SCULPT.

(ii) The keywords `row` and `col` in Section 2 are syntactic sugar for right^+ and down^+ in core-SCULPT, respectively. So, $\text{col}((2, 2))$, which denotes *the column below the cell* $(2, 2)$ in Section 2, is syntactic sugar for $\text{down}^+(\text{down} \cdot \text{right}(\text{root}))$.

(iii) The only exception to rule (ii) above are `row` and column expressions of the form $\text{row}(k)$ and $\text{col}(\ell)$. These abbreviate $\text{right}^*(k, 1)$ and $\text{down}^*(1, \ell)$, respectively. (Where $(k, 1)$ and $(1, \ell)$ need to be further unfolded themselves.)

As an example, the selection expression $\text{row}(1)$ of Figure 6 can be written as $\text{right}^*(\text{root})$ or, equivalently, right^* and the expression $\text{col}(\text{subject})$ as $\text{down}^+(\text{subject})$. ■

Content expressions. A *content expression* is simply a regular expression ρ over the set of tokens Δ . To define when a region in a tokenized table T is valid with respect to content expression ρ , let us first introduce the following order on coordinates. We say that coordinate (k, ℓ) precedes coordinate (k', ℓ') if we visit (k, ℓ) earlier than (k', ℓ') in a left-to-right top-down traversal of the cells of T , i.e., it precedes it in lexicographic order. Formally, $(k, \ell) < (k', \ell')$ if $k < k'$ or if $k = k'$ but $\ell < \ell'$.

Now, let T be a tokenized table, let z be a region of T , and let ρ be a content expression. Then (T, z) satisfies the content expression ρ under the *region-based semantics*, denoted $T, z \models_{\text{region}} \rho$ if there exist tokens $a_1, \dots, a_n \in \Delta$ such that $a_1 \dots a_n \in \mathcal{L}(\rho)$ and $a_i \in T_{c_i}$, where c_1, \dots, c_n is the enumeration in table order of all coordinates in z .

To define the row-based semantics we used in Section 2, we require the following notions. Let z be a region of T . We say that subregion $z' \subseteq z$ is a *row of* z if there exists some k such that $z' = \{(k, \ell) \mid (k, \ell) \in z\}$. Now, (T, z) satisfies the content expression ρ under the *row-based semantics*, denoted $T, z \models \rho$, if for every row z' of z , we have $T, z' \models_{\text{region}} \rho$.

REMARK 3.5. Recall that, for ease of exposition, we allowed tables to be non-rectangular in Section 2 whereas in our formal model, tables are always rectangular. In particular, shorter rows are padded with \perp to obtain rectangularity. This implies that, some content expressions of Section 2 need to be adapted in our formal model. For example, the rule $\text{row}(1) \rightarrow \text{name}$ of Figure 4 needs to be adapted to $\text{row}(1) \rightarrow \text{name}, \perp, \perp, \perp$ to take the padding into account. ■

4. EFFICIENT VALIDATION

In this section we consider the *validation* (or *evaluation*) problem for tabular schemas. This problem asks, given a tokenized table or tabular document T and a tabular schema R , whether T satisfies R . We consider the problem in a *main-memory* and *streaming* variant. Intuitively, T is given as a table in the former and as a tabular document in the latter setting.

4.1 Validation in Linear Time

When T is given as a tokenized table, we can essentially assume that we can navigate from a cell (i, j) to any of its four neighbours $\text{up}(\{(i, j)\})$, $\text{down}(\{(i, j)\})$, $\text{left}(\{(i, j)\})$, and $\text{right}(\{(i, j)\})$ in constant time. Under these assumptions we show that T can be validated against a tabular

schema in linear time combined complexity.⁷ The proof strongly relies on the linear time combined complexity for PDL model checking.

THEOREM 4.1. *The valuation problem for a tabular document T and a schema R is in linear time combined complexity, that is, time $O(|T||R|)$.*

4.2 Streaming Validation

Even though Theorem 4.1 implies that SCULPT schemas can be efficiently validated, the later claim only holds true when the tabular document can be fully loaded in memory and multiple passes can be made through the document. However, when the input data is large it is sometimes desirable to have a *streaming validation algorithm* that makes only a single pass over the input tabular document and uses only limited memory. In this section we identify several fragments of core-SCULPT that admits such streaming validation algorithms.

Streaming model. Let us begin by defining when an algorithm validates in a streaming fashion. In this respect, we draw inspiration from the SAX Streaming API for XML: we can view a tokenized table T as a sequence of events generated by visiting the cells of T in table order. Here, whenever we visit a new cell, an event $\langle \text{cell } \Gamma \rangle$ is emitted, with Γ the set of tokens in the visited cell. Whenever we move to a new row, an event of type $\langle \text{new row} \rangle$ is emitted.

Note that the tokenized event stream can easily be generated “on the fly” when parsing a tabular document: we start reading the tabular document, one character at a time, until we reach a delimiter. All non-delimiter characters are used as input to, e.g., a finite state automaton that allows us to check which tokens match the current cell’s content. When we reach a delimiter, a $\langle \text{cell } \Gamma \rangle$ event is emitted with the corresponding set of matching tokens. If the delimiter is a row delimiter, then also a $\langle \text{new row} \rangle$ is emitted. We repeat this until the end of the file.

EXAMPLE 4.2. Consider the tabular document from Figure 1 together with the corresponding SCULPT schema S in Figure 2. The tokenized table of this document according to S yields the event stream

```

⟨cell ∅⟩⟨cell {ARUA}⟩⟨cell {BOMBO}⟩⟨cell {ENTEBBE AIR}⟩
⟨new row⟩⟨cell {Timestamp}⟩⟨cell {Temperature}⟩
⟨cell {Temperature}⟩⟨cell {Temperature}⟩⟨new row⟩...

```

DEFINITION 4.3 (STREAMABILITY). A tabular schema R is said to be *weakly streamable*, if there exists a Turing Machine M that

- can only read its input tape once, from left to right;
- for every tokenized table T , when started with the event stream of T on its input tape, accepts iff $T \models R$; and
- has an auxiliary work tape that can be used during processing, but it cannot use more than $O(m \log(n))$ of space on this work tape, where n is the total number of cells in T , and m the number of columns.

We say that R is *strongly streamable* if the Turing Machine M only requires $O(\log(n))$ space on its work tape.

Here, strong streamability corresponds to the commonly studied notions of streaming evaluation. We consider weak

⁷Combined complexity is a standard complexity measure introduced by Vardi; see [30].

streamability to be very relevant as well because, based on the W3C use cases, tabular data often seems to be similar in spirit to relational tables and, in these cases, is very narrow and deep. In particular, $m = O(\log n)$ in these cases.

Weak streamability. To enable streaming validation, we restrict our attention to so-called *forward* coordinate and navigational expressions which are expressions where $\langle \alpha \rangle$ is not allowed, and we never look up or left. That is, a coordinate or navigational expression is *forward* if it is generated by the following syntax.

$$\begin{aligned} \varphi, \psi &:= a \mid \text{root} \mid \text{true} \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \neg \varphi \mid \alpha(\varphi) \\ \alpha, \beta &:= \varepsilon \mid \text{down} \mid \text{right} \mid [\varphi] \mid (\alpha \cdot \beta) \mid (\alpha + \beta) \mid (\alpha^*) \end{aligned}$$

We do not consider the operator $\langle \alpha \rangle$ in the forward fragment because it can be seen as a backward operator: $\langle \text{right} \cdot [a] \rangle$ is equivalent to $\text{left}(a)$.

A core-SCULPT schema is *forward* if it mentions only forward coordinate expressions.

THEOREM 4.4. *Forward core-SCULPT is weakly streamable.*

PROOF SKETCH. Consider a rule $\varphi \rightarrow \rho$ with φ a forward coordinate expression and ρ a content expression. We will show that coordinate expressions φ can be evaluated in a streaming fashion by constructing a special kind of finite state automaton (called *coordinate automaton*) that allows us to decide, at each position in the event stream, if the currently visited cell is in $\llbracket \varphi \rrbracket_T$. Whenever we find that this is the case, we apply the current cell contents to ρ (which we also evaluate by means of a finite state automaton). Now observe that $T \models \varphi \rightarrow \rho$ iff (1) under the row based semantics, whenever we see $\langle \text{new row} \rangle$, the automaton for ρ is in a final state and (2) under the region-based semantics, when we reach the end of the event stream, the automaton for ρ is in a final state. We then obtain weak streamability by showing that coordinate automata for φ can be simulated in space $O(m \log(n))$, whereas it is known that the finite state automaton for ρ can be simulated in constant space. \square

Strong streamability. Forward core-SCULPT is not strongly streamable as no schema with a rule that contains subexpressions of the form $\text{col}(a)$ (which are prevalent in Section 2) can be strongly streamable. This can be seen using a simple argument from communication complexity. Indeed, assume that the first row has k cells, some of which have the token a and some of which do not. If we want to evaluate $\text{col}(a)$ in a streaming fashion, we need to identify the cells in the second row that are in the same columns as the a -tokens in the first row. But, this is precisely the equality of two k -bit strings problem, which requires $\Omega(k)$ bits in deterministic communication complexity (Example 1.21 in [15]). These $\Omega(k)$ bits are what we need to store when going from the first to the second row. Since k can be $\Theta(n)$, this amount of space is more than we allow for strongly streamable tabular schemas, and hence no schema containing $\text{col}(a)$ is strongly streamable.

The underlying reason why $\text{col}(a)$ is not strongly streamable is because, in general, the token a can occur arbitrarily often. However, in all such cases in Section 2 and in the W3C use cases, the occurrences of a are very restricted. We could therefore obtain strong streamability for such expressions by adding constructs in the language that restrict how certain tokens can appear:

unique(a) **unique-per-row(a)**

The former asserts that token a should occur only once in the whole table and the latter that a occurs at most once in each row. More formally, the former predicate holds in a table T if $\llbracket a \rrbracket_T$ contains at most one element and the latter holds in table T if $\llbracket a \rrbracket_T$ contains at most one element of the form (r, c) for each row number r . Notice that a strong streaming algorithm can easily check whether these predicates hold.

We use the above predicates to define two notions of *guardedness* for region selection expressions. Guarded formulas will be strongly streamable. We say that token a is *row-guarded* if **unique-per-row(a)** appears in the schema. If **unique(a)** appears in the schema it is, in addition, also *guarded*. The two notions of guardedness capture the following intuition: if φ is row-guarded, then $\text{down}(\varphi)$ is strongly streamable and if it is guarded, then $\text{down}^*(\varphi)$ is strongly streamable. The main idea is that, in both cases, the number of cells we need to remember when going from one row to the next does not depend on the width of the table. We now define (row)-guardedness inductively on the forward language:

- **root** and **true** are guarded and row-guarded;
- **right** $^*(\varphi)$ is guarded and row-guarded for every φ that does not contain a navigational subexpression;
- if $\varphi, \psi, \alpha(\varphi), \beta(\varphi)$ are guarded (resp., row-guarded), then
 - $\varphi \wedge \psi, \varphi \vee \psi, \varepsilon(\varphi), \text{down}(\varphi), \text{right}(\varphi),$
 - $(\alpha \cdot \text{down})(\varphi), (\alpha \cdot \text{right})(\varphi),$ and $(\alpha + \beta)(\varphi)$
 are guarded (resp., row-guarded);
- if φ and $\alpha(\psi)$ are guarded then $\text{down}^*(\varphi)$ and $(\alpha \cdot \text{down}^*)(\psi)$ are guarded; and
- if φ and $\alpha(\psi)$ are row-guarded then **right** $^*(\varphi)$ and $(\alpha \cdot \text{right}^*)(\psi)$ are guarded.

DEFINITION 4.5. A forward core-SCULPT schema is *guarded*, if all region selection expressions that use the **down**-operator are row-guarded and all region selection expressions that use **down** * are guarded.

Notice that guardedness of a SCULPT schema can be tested in linear time. Furthermore notice that every SCULPT schemas in this paper becomes strongly streamable if we add the predicates **unique(a)** for tokens a that we use in expressions using **col**, **down**, or **down** * .

THEOREM 4.6. *Guarded forward core-SCULPT is strongly streamable.*

5. SCULPT EXTENSIONS

Next, we describe a number of extensions to SCULPT. These include alternative grouping semantics, types, complex content cells, and a concept for a transformation language.

5.1 Region semantics

The examples in Section 2 all use a row-based semantics of SCULPT where the content expression is matched over every row in the selected region. That is, the cells in the selected region are ‘grouped by’ the row they occur in. There are of course other ways to group cells, by column, for instance, or by not grouping them at all. The latter case is already defined in Section 3 as *region-based semantics*. In SCULPT, we indicate rules using this semantics with a double arrow \Rightarrow rather than a single arrow. Notice the difference between

the rules `col(2) -> Null | Number` and `col(2) => (Null | Number)*`. Both require each cell in the second column to be empty or a number but express this differently. (The former way is closer to how one defines the schema of a table in SQL, which is why we chose it as a default.) Example 5.1 describes a more realistic application of `=>`-rules. This example corresponds to use case 12 in [27], is called “*Chemical Structures*” and aims to interpret Protein Data Bank (PDB) files as tabular data. This particular use case is interesting because it illustrates that the view of W3C on tabular data is not restricted to traditional comma-separated values files. We note that Theorems 4.1, 4.4, and 4.6 still hold if SCULPT schemas contain both rules under row-based and region-based semantics.

EXAMPLE 5.1. Figure 7 displays a slightly shortened version of the PDB file mentioned in use case 12 in [27]. The corresponding SCULPT schema could contain the following rules:

```
row(1) -> HEADER, Type, Date, ID
col(1) => HEADER, TITLE*, dots, EXPDATA, AUTHOR*,
         dots, REMARK*, dots, SEQRES*, dots, ATOM*
```

The last rule employs the region semantics and specifies the order in which tokens in the first column should appear. ■

5.2 Token types

The PDB fragment in Figure 7 contains cells that have the same content but seem to have a different meaning. It can be convenient to differentiate between cells by using *token types* as follows:

```
%% Token types
%% left: name of the token type
%% right: region selection expression for token type
REMARK-Header <= down*[dots].down[REMARK]
REMARK-Comment <= down*[dots].down[REMARK].down
REMARK-Rest <=
  down*[dots].down[REMARK].down.(down[REMARK])*
```

Note that we abbreviated rules of the form $\alpha(\text{root})$ by α . We denoted the concatenation operator of navigational expressions by “.”. `REMARK-Header` is the topmost cell containing `REMARK` in Figure 7, `REMARK-Comment` is the one immediately below, and `REMARK-Rest` is the rest. We can now use token types to write rules such as

```
row(REMARK-Header) -> ...
row(REMARK-Comment) -> ...
row(REMARK-Rest) -> ...
```

Token types do not add additional expressiveness to the language since one can simply replace `REMARK-HEADER` by `down*[dots].down[REMARK](root)` in the rule. But the ability to use different names for fields with the same content may be useful for writing more readable schemas. In this case, the names suggest that the block of remarks is divided into a header, some comment, and the rest.

5.3 Transformations and Annotations

While it is beyond the scope of this document to develop a transformation language for tables, we argue that region selection expressions can be easily employed as basic building blocks for a transformation language aimed at transforming tables into a variety of formats like, for instance, RDF, JSON, or XML (one of the scopes expressed in [31]). Region

selection expressions are then used to identify relevant parts of a table.

Basic Transformations. Consider Figure 1 (of Example 2.1) again, where we see that several columns have the value `-99.00`. Since winter does not get this extreme in Uganda, this value is simply a dummy which should not be considered when computing, e.g., the average temperature in Uganda in 1935. Instead, for the fragment of Figure 1, it would be desirable to only select the columns that do not contain `-99.00`. To do this, we can simply define a new token and a new token type for the region of the table we are interested in.

```
Useless-Temp = -99.00
```

```
%% Token type
Useful <= col(1) or
         (Temperature and not Useless-Temp) or
         (row(1) and not up*(Useless-Temp))
```

The region defined by `Useful` contains

```
, ENTEBBE AIR
1935.04,      27.83
1935.12,      25.72
1935.21,      26.44
[...]
```

which could then be exported. Using simple for-loops we can iterate over rows, columns, or cells, and compute aggregates. For example,

```
Useful-values <= (Temperature and not Useless-Temp)
```

```
For each column c in Useful-values {
  print Average(c)
}
```

would output `25.65`, the average of the values below `ENTEBBE AIR` in Figure 1. The region defined by `Useful-values` is a set of table cells, with coordinates. These coordinates can be used to handle information column-wise in the for-loop: It simply iterates over all column coordinates that are present in the region. Iteration over rows or single cells would work analogously.

Namespaces, Annotations and RDF. Assume that we want to say that certain cells in Figure 3 are geographical regions. To this end, the SCULPT schema could contain a definition of a default namespace:

```
namespace default = http://foo.org/nationalstats.csv
namespace x = [...]
```

Region selection expressions can then be used to specify which cells should be treated as objects in which namespace. For example, the code fragment

```
For each cell c in col(GeoArea) {
  c.namespace = default
}
```

could express that each cell below `GeoArea` is an entity in namespace `http://foo.org/nationalstats.csv`. So, the cell containing `England` represents the entity

```
http://foo.org/nationalstats.csv:England,
```

similar for the cell containing `Wales`, etc. (Here we assume that `.namespace` is a predefined operation on cells.)

We can also annotate cells with meta-information (as is currently being considered in Section 2.2 of [29]). The code fragment

```

HEADER      EXTRACELLULAR MATRIX                22-JAN-98  1A3I
TITLE       X-RAY CRYSTALLOGRAPHIC DETERMINATION OF A COLLAGEN-LIKE
TITLE       2 PEPTIDE WITH THE REPEATING SEQUENCE (PRO-PRO-GLY)
...
EXPDTA     X-RAY DIFFRACTION
AUTHOR     R.Z.KRAMER,L.VITAGLIANO,J.BELLA,R.BERISIO,L.MAZZARELLA,
AUTHOR     2 B.BRODSKY,A.ZAGARI,H.M.BERMAN
...
REMARK 350 BIOMOLECULE: 1
REMARK 350 APPLY THE FOLLOWING TO CHAINS: A, B, C
REMARK 350  BIOMT1   1  1.000000  0.000000  0.000000          0.00000
REMARK 350  BIOMT2   1  0.000000  1.000000  0.000000          0.00000
...
SEQRES     1  A     9  PRO PRO GLY PRO PRO GLY PRO PRO GLY
SEQRES     1  B     6  PRO PRO GLY PRO PRO GLY
SEQRES     1  C     6  PRO PRO GLY PRO PRO GLY
...
ATOM       1  N     PRO A   1           8.316  21.206  21.530  1.00  17.44          N
ATOM       2  CA    PRO A   1           7.608  20.729  20.336  1.00  17.44          C
ATOM       3  C     PRO A   1           8.487  20.707  19.092  1.00  17.44          C
ATOM       4  O     PRO A   1           9.466  21.457  19.005  1.00  17.44          O
ATOM       5  CB    PRO A   1           6.460  21.723  20.211  1.00  22.26          C

```

Figure 7: Fragment of a PDB file.

```

For each cell c in col(GeoArea) {
  annotate c with "rdf:type dbpedia-owl:Place"
  annotate c with "owl:sameAs fbase:" + c.content
}

```

(assuming appropriate namespace definitions for `rdf`, `owl`, etc.) could express that each cell below `GeoArea` should be annotated with `rdf:type dbpedia-owl:Place` and, in addition, the `England` cell with `owl:sameAs fbase:England`, the `Wales` cell with `owl:sameAs fbase:Wales`, etc. We assume that `annotate`, `with`, and `.content` are reserved words or operators in the language.

These ingredients also seem useful for exporting to RDF. We could write, e.g.,

```
print "@prefix : <http://foo.org/nationalstats.csv>"
```

```

For each cell c in col(GeoArea) {
  print ":"+c.content+"owl:sameAs fbase:"+c.content
}

```

to produce an RDF file that says that `:England` in the default namespace is the same as `fbase:England`. Looking at Figure 5, one can also imagine constructs like

```
RDF <= col(subject) or col(predicate) or col(object)
```

```

For each row r in RDF {
  print r.cells[1] +" "+ r.cells[2] +" "+ r.cells[3]
}

```

to facilitate the construction of RDF triples taking content from several cells.

5.4 Complex content

The CSV on the Web WG is considering allowing complex content (such as lists) in cells (Section 3.8 in [22]). SCULPT can be easily extended to reason about complex content. Our formal definition of tabular documents already considers (Section 3) a finite set of delimiters, which goes beyond

the two delimiters (row- and column-) that we used until now.

In a spirit similar to region-based semantics, one can also imagine a subcell-based semantics, for example, a rule of the form

```
col(1) .> (String)*
```

could express that each cell in the first column contains a list of Strings. Notice the use of `.>` instead of `->` to denote that we specify the contents of each individual cell in the region, instead of each row. The statement `List Delim = ;` in the beginning of the schema could say that the semicolon is the delimiter for lists within a cell.

6. CONCLUSIONS

We presented the schema language SCULPT for tabular data on the Web and showcased its flexibility and usability through a wide range of examples and use cases. While region selection expressions are at the very center of SCULPT, we think they can be more broadly applied. Region selection expressions can be used, for instance, as a cornerstone for annotation- and transformation languages for tabular data and thus for a principled approach for integrating such data into the Semantic Web. The whole approach of SCULPT is strongly rooted in theoretical foundations and, at the same time, in well established technology such as XPath. For these reasons, we expect the language to be very robust and, at the same time, highly accessible for users. Two prominent directions for future work are the following: (1) expand the usefulness of SCULPT by further exploring the extensions in Section 5; and, (2) study static analysis problems related to SCULPT and region selector expressions leveraging on the diverse box of tools from formal language theory and logic.

Acknowledgments

We thank Marcelo Arenas for bringing [29] to our attention.

7. REFERENCES

- [1] R. W. Adam Retter, David Underdown. Csv schema 1.0: A language for defining and validating csv data. <http://digital-preservation.github.io/csv-schema/csv-schema-1.0.html>.
- [2] N. Alechina and N. Immerman. Reachability logic: An efficient fragment of transitive closure logic. *Logic Journal of the IGPL*, 8(3):325–337, 2000.
- [3] M. Arenas, S. Conca, and J. Pérez. Counting beyond a yottabyte, or how SPARQL 1.1 property paths will prevent adoption of the standard. In *International World Wide Web Conference (WWW)*, pages 629–638, 2012.
- [4] A. Berglund, S. Boag, D. Chamberlin, M. F. Fernández, M. Kay, J. Robie, and J. Siméon. XML Path Language (XPath) 2.0. Technical report, World Wide Web Consortium, January 2007. W3C Recommendation, <http://www.w3.org/TR/2007/REC-xpath20-20070123/>.
- [5] G. J. Bex, W. Gelade, F. Neven, and S. Vansumeren. Learning deterministic regular expressions for the inference of schemas from XML data. In *International World Wide Web Conference (WWW)*, pages 825–834, 2008.
- [6] G. J. Bex, W. Martens, F. Neven, and T. Schwentick. Expressiveness of XSDs: from practice to theory, there and back again. In *International World Wide Web Conference (WWW)*, pages 712–721, 2005.
- [7] R. Cleaveland and B. Steffen. A linear-time model-checking algorithm for the alternation-free modal mu-calculus. *Formal Methods in System Design*, 2(2):121–147, 1993.
- [8] S. Das, S. Sundara, and R. Cyganiak. R2RML: RDB to RDF mapping language. <http://www.w3.org/TR/r2rml/>. W3C Recommendation 27 September 2012.
- [9] D. Fallside and P. Walmsley. XML Schema Part 0: Primer (second edition). Technical report, World Wide Web Consortium, October 2004. <http://www.w3.org/TR/2004/REC-xmlschema-0-20041028/>.
- [10] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 18(2):194–211, 1979.
- [11] J. E. F. Friedl. *Mastering Regular Expressions*. O’Reilly Media, 3rd edition edition, 2006.
- [12] W. Gelade and F. Neven. Succinctness of pattern-based schema languages for XML. *J. Comput. Syst. Sci.*, 77(3):505–519, 2011.
- [13] Google. DSPL: Dataset publishing language. <https://developers.google.com/public-data/>. Last accessed 04/11/2014.
- [14] V. Kumar, P. Madhusudan, and M. Viswanathan. Visibly pushdown automata for streaming XML. In *International World Wide Web Conference (WWW)*, pages 1053–1062, 2007.
- [15] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [16] O. K. F. Labs. Tabular data package. <http://dataprotocols.org/tabular-data-package/>. Version 1.0-beta-2. Last accessed 04/11/2014.
- [17] L. Libkin, W. Martens, and D. Vrgoc. Querying graph databases with XPath. In *International Conference on Database Theory (ICDT)*, pages 129–140, 2013.
- [18] K. Losemann and W. Martens. The complexity of evaluating path expressions in SPARQL. In *International Symposium on Principles of Database Systems (PODS)*, pages 101–112, 2012.
- [19] W. Martens, F. Neven, M. Niewerth, and T. Schwentick. Developing and analyzing xsds through bonxai. *PVLDB*, 5(12):1994–1997, 2012.
- [20] W. Martens, F. Neven, T. Schwentick, and G. Bex. Expressiveness and complexity of XML Schema. *ACM Transactions on Database Systems*, 31(3):770–813, 2006.
- [21] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. *ACM Trans. Database Syst.*, 34(3), 2009.
- [22] R. Pollock and J. Tennison. Metadata vocabulary for tabular data. Technical report, World Wide Web Consortium (W3C), July 2014. www.w3.org/TR/2014/WD-tabular-metadata-20140710/.
- [23] E. Prud’hommeaux, J. E. L. Gayo, and H. Solbrig. Shape expressions: An RDF validation and transformation language. In *International Conference on Semantic Systems*, 2014.
- [24] A. G. Ryman, A. L. Hors, and S. Speicher. OSLC resource shape: A language for defining constraints on linked data. In *WWW Workshop on Linked Data on the Web*, 2013.
- [25] L. Segoufin and C. Sirangelo. Constant-memory validation of streaming XML documents against DTDs. In *International Conference on Database Theory (ICDT)*, pages 299–313, 2007.
- [26] L. Segoufin and V. Vianu. Validating streaming XML documents. In *International Symposium on Principles of Database Systems (PODS)*, pages 53–64, 2002.
- [27] J. Tandy, D. Ceolin, and E. Stephan. CSV on the Web: Use cases and requirements. Technical report, World Wide Web Consortium (W3C), October 2014. <http://w3c.github.io/csvw/use-cases-and-requirements/>.
- [28] J. Tennison. 2014: The year of CSV. <http://theodi.org/blog/2014-the-year-of-csv>. last accessed 04/11/2014.
- [29] J. Tennison and G. Kellogg. Model for tabular data and metadata on the web. Technical report, World Wide Web Consortium (W3C), July 2014. www.w3.org/TR/2014/WD-tabular-data-model-20140710/.
- [30] M. Y. Vardi. The complexity of relational query languages (extended abstract). In *ACM Symposium on Theory of Computing (STOC)*, pages 137–146, 1982.
- [31] W3C. CSV on the web working group charter. <http://www.w3.org/2013/05/lcsv-charter.html>.

APPENDIX

A. PROOF OF THEOREM 4.1

LEMMA A.1 (SEE FACT 5.1, [17]). *For every coordinate expression φ , we can compute $\llbracket \varphi \rrbracket_T$ in time $O(|T||\varphi|)$. For every set C of coordinates in T and every navigational expression α , we can compute $\llbracket \alpha(C) \rrbracket_T$ in time $O(|T||\alpha|)$.*

PROOF OF THEOREM 4.1 (SKETCH). As we have already mentioned, our region expressions are essentially a variant of Propositional Dynamic Logic (PDL), tweaked to navigate in tables. It is known that PDL has linear time combined complexity for global model checking [2, 7]. That is, given a PDL formula φ and a Kripke structure (essentially, a graph) G , one can decide in time $O(|\varphi||G|)$ whether $G \models \varphi$. One can simply view our tables as Kripke structures that are shaped like a grid. The result follows by a straightforward adaptation of the algorithm in [2]. \square

B. PROOF OF THEOREM 4.4

The crux behind Theorem 4.4 is the following. Consider a rule $\varphi \rightarrow \rho$ with φ a forward coordinate expression and ρ a content expression. We will show that we are able to evaluate coordinate expressions φ in a streaming fashion by constructing a special kind of finite state automaton (called *coordinate automaton*) that allows us to decide, at each position in the event stream, if the currently visited cell is in $\llbracket \varphi \rrbracket_T$. Whenever we find that this is the case, we apply the current cell contents to ρ (which we also evaluate by means of a finite state automaton). Now observe that $T \models \varphi \rightarrow \rho$ iff

1. Under the row based semantics, the automaton for ρ is in a final state whenever we see $\langle \text{new row} \rangle$.
2. Under the region-based semantics, when we reach the end of the event stream, the automaton for ρ is in a final state.

B.1 Coordinate automata

Let us first introduce the kinds of finite state automata that we will use to evaluate coordinate expressions. A *coordinate automaton* (CA for short) A over Λ is a tuple (Q, q_0, F, δ) where:

- Q is a finite set of states;
- $q_0 \in Q$ is an initial state;
- $F \subseteq Q$ is a set of *final states*;
- δ is a finite transition relation consisting of triples, each having one of the forms (q, ε, q') , $(q, [\varphi], q')$, (q, right, q') , (q, down, q') , $(q, \langle \text{new row} \rangle, q')$, with φ a *forward coordinate expression*.

A CA processes event streams of tokenized tables. It largely behaves like a normal non-deterministic finite state automaton, but has a number of special features. First, it is equipped with a register that stores the coordinate of the current cell being processed. Second, in each step during execution, the automaton can either be in an *active state*, or in a *suspended state*. Only active states can fire new transitions; suspended states remain dormant until they become active again. Syntactically, a suspended state is simply a pair of the form $q: \ell$ with $q \in Q$ and $\ell \in \mathbb{N}$, $\ell \geq 1$.

The semantics of a suspended state $q: \ell$ is that the automaton keeps reading events from the input stream (updating the coordinate register, but remaining in the suspended state) until the next time that we visit a cell in column ℓ , when the state becomes active again. This way, the coordinate automaton can simulate moving downwards in the table. In particular, when following a transition of the form (q, down, q') from state q , the automaton records the column number of the current position (say, ℓ); moves to state q' ; but immediately suspends q' until the following cell in column ℓ is visited. Third and finally, a CA can check whether the current cell is selected by a forward coordinate expression φ by means of a transition of the form $(q, [\varphi], q')$. The CA moves from q to q' if the corresponding cell in the table is selected by φ . (Note that the CA can hence use coordinate expressions as oracles.) In contrast to standard automata, these transitions do not cause the CA to move to the next event, however. Moving to the next event is done by either transitions of the form (q, right, q') or $(q, \langle \text{new row} \rangle, q')$.

Note. Coordinate automata, as defined above, are very general and powerful. Obviously, the fact that a coordinate automaton can use coordinate expressions as oracles makes them very powerful. Indeed, as we will see below, they are almost trivially able to express the semantics of coordinate expressions. In order to obtain Theorem 4.4, however, we will compile coordinate expressions into CA whose oracles are restricted.

Semantics. Formally, a *configuration* of A is a tuple (i, k, ℓ, ς) where i is a pointer to the current event being read from the input tape; (k, ℓ) is a coordinate; and ς is either an element of Q (an active state), or a pair of the form $q: j$ with $q \in Q$ and $j \in \mathbb{N}$, $j \geq 1$ (a suspended state).

Let T be a table and let $s = \sigma_1 \cdots \sigma_n$ be an event stream for T , each σ_i being either $\langle \text{cell } \Gamma \rangle$ with Γ a set of tokens, or $\langle \text{new row} \rangle$. Let $c_0 = (i, k, \ell, \varsigma)$ be a configuration, with $1 \leq i \leq n$ and (k, ℓ) the coordinate of the cell visited in T when σ_i is emitted. A *run* ρ of A on s starting at c_0 is a sequence c_0, \dots, c_m of configurations such that, for all j with $1 \leq j < m$ one of the following holds for $c_j = (i_j, k_j, \ell_j, \varsigma_j)$ and $c_{j+1} = (i_{j+1}, k_{j+1}, \ell_{j+1}, \varsigma_{j+1})$:


- Transition from active state: $\varsigma_j \in Q$
 1. $(\varsigma_j, \varepsilon, \varsigma_{j+1}) \in \delta$, $i_{j+1} = i_j$, $k_{j+1} = k_j$, and $\ell_{j+1} = \ell_j$ (epsilon transition)
 2. $(\varsigma_j, [\varphi], \varsigma_{j+1}) \in \delta$, $(k_j, \ell_j) \in \llbracket \varphi \rrbracket_T$, $i_{j+1} = i_j$, $k_{j+1} = k_j$, and $\ell_{j+1} = \ell_j$ (ordinary transition)
 3. $(\varsigma_j, \text{right}, \varsigma_{j+1}) \in \delta$, $\sigma_{i_{j+1}} \neq \langle \text{new row} \rangle$, $i_{j+1} = i_j + 1$, $k_{j+1} = k_j$, and $\ell_{j+1} = \ell_j + 1$ (move right transition)
 4. $(\varsigma_j, \langle \text{new row} \rangle, \varsigma_{j+1}) \in \delta$, $\sigma_{i_{j+1}} = \langle \text{new row} \rangle$, $i_{j+1} = i_j + 2$, $k_{j+1} = k_j + 1$, and $\ell_{j+1} = 1$ (new row transition)
 5. $(\varsigma_j, \text{down}, q) \in \Delta$, $i_{j+1} = i_j$, $k_{j+1} = k_j$, $\ell_{j+1} = \ell_j$, and $\varsigma_{j+1} = q: \ell_j$, (down transition)
- Transition from suspended state: $\varsigma_j = q: \ell$ for some $q \in Q$ and $\ell \in \mathbb{N}$ with $\ell \geq 1$.
 1. $\sigma_{i_{j+1}} \neq \langle \text{new row} \rangle$, $i_{j+1} = i_j + 1$, $k_{j+1} = k_j$, $\ell_{j+1} = \ell_j + 1$, and either (a) $\ell \neq \ell_{j+1}$ and $\varsigma_{j+1} = \varsigma_j$ (suspended right transition) or (b) $\ell = \ell_{j+1}$ and $\varsigma_{j+1} = q$ (wake up).
 2. $\sigma_{i_{j+1}} = \langle \text{new row} \rangle$, $i_{j+1} = i_j + 2$, $k_{j+1} = k_j + 1$, $\ell_{j+1} = 1$, and either (a) $\ell \neq 1$ and $\varsigma_{j+1} = \varsigma_j$

(suspended new row transition) or (b) $\ell = 1$ and $\varsigma_{j+1} = q$ (wake up).

DEFINITION B.1. Let T be a tokenized table. A CA A over Λ is said to select coordinate $(k, \ell) \in \text{coords}(T)$, denoted $T, (k, \ell) \models A$, if, there exists a run ρ of A on the event stream for T starting from configuration $(1, 1, 1, q_0)$ with q_0 the initial state of A such that there is a configuration c in ρ with $c = (i, k, \ell, \varsigma)$ for some i and ς such that $\varsigma \in F$. (In particular, ς is not suspended.)

We write $\llbracket A \rrbracket_T$ for the set of all coordinates selected by A on T .

EXAMPLE B.2. The following CA selects the same coordinates as region expression $\text{down}^* \text{right}^+(\text{ARUBA})$



■

Note that, in the definition above, a CA always selects coordinates when started from the root coordinate $(1, 1)$ of the table (i.e., the beginning of the event stream). In what follows, it will be convenient for technical reasons to say that a CA A selects coordinate (k, ℓ) in T when started at coordinate (k', ℓ') . This is formally defined as follows.

DEFINITION B.3. Let $s = \sigma_1 \cdots \sigma_n$ be the event stream of T . Let σ_i with $1 \leq i \leq n$ be the symbol in this stream that corresponds to the cell in T with coordinate (k', ℓ') .⁸ Then A selects (k, ℓ) when started from (k', ℓ') in T if there exists a run ρ of A on $\sigma_1 \cdots \sigma_n$ starting with configuration (i, k', ℓ', q_0) such that there is a configuration c in ρ with $c = (j, k, \ell, \varsigma)$ for some j and some $\varsigma \in F$. (In particular, ς is not suspended.)

We write $\llbracket A \rrbracket_{T,c}$ for the set of all coordinates selected by A on T when started from coordinate c .

DEFINITION B.4. Coordinate automaton A expresses coordinate expression φ if $\llbracket \varphi \rrbracket_T = \llbracket A \rrbracket_T$, for every tokenized table T . We say that φ is definable by means of a CA if there exists a CA that expresses φ . Similarly, if α is a navigational expression, then A expresses α if $\alpha(\{c\}) = \llbracket A \rrbracket_{T,c}$, for every tokenized table T and every $c \in \text{coords}(T)$.

It should be noted that the fact that a coordinate automaton can use coordinate expressions as oracles makes them very powerful. Indeed, they are almost trivially able to express any coordinate or navigational expression. In order to obtain Theorem 4.4, however, we will compile coordinate expressions into CA whose oracles are restricted. By restricted here we mean that the oracles are of a different level than coordinate or navigational expressions that

⁸Here, we tacitly make the convention that $\langle \text{new row} \rangle$ events do not corresponds to any cell, so σ_i is an event of the form $\langle \text{cell } \Gamma \rangle$.

are being expressed. The level intuitively corresponds to the maximum nesting level of coordinate and navigational expressions. The formal definition is as follows.

$$\begin{aligned} \text{lvl}(a) &= 0 \\ \text{lvl}(\text{root}) &= 0 \\ \text{lvl}(\text{true}) &= 0 \\ \text{lvl}(\varphi \wedge \psi) &= \max(\text{lvl}(\varphi), \text{lvl}(\psi)) \\ \text{lvl}(\varphi \vee \psi) &= \max(\text{lvl}(\varphi), \text{lvl}(\psi)) \\ \text{lvl}(\neg\varphi) &= \text{lvl}(\varphi) \\ \text{lvl}(\alpha(\varphi)) &= 1 + \max(\text{lvl}(\alpha), \text{lvl}(\varphi)) \end{aligned}$$

$$\begin{aligned} \text{lvl}(\varepsilon) &= 0 \\ \text{lvl}(\text{down}) &= 0 \\ \text{lvl}(\text{right}) &= 0 \\ \text{lvl}([\varphi]) &= 1 + \text{lvl}(\varphi) \\ \text{lvl}(\alpha \cdot \beta) &= \max(\text{lvl}(\alpha), \text{lvl}(\beta)) \\ \text{lvl}(\alpha + \beta) &= \max(\text{lvl}(\alpha), \text{lvl}(\beta)) \\ \text{lvl}(\alpha^*) &= \text{lvl}(\alpha) \end{aligned}$$

To illustrate, $\text{lvl}(\text{down} \cdot \text{right}^*) = 0$, $\text{lvl}([a] \cdot \text{right}) = 1$, and $\text{lvl}(\text{right}(a \wedge b)) = 1$.

Define the level of a CA to be the maximum level of any coordinate expression occurring in it. We now establish a number of technical lemmas that relate CA to coordinate and navigational expressions.

LEMMA B.5. Every forward navigational expression α is definable by means of a CA of level $\max(0, \text{lvl}(\alpha) - 1)$.

PROOF. We construct, by induction on α , a CA A that expresses α and is of level at most $\text{lvl}(\alpha) - 1$. The construction is essentially the same as the Thompson construction for transforming regular expressions into finite state automata.

- If $\alpha = \varepsilon$, then A is the CA with one state, which is both initial and final, and which does not have any transitions.
- If $\alpha = \text{down}$, then A is the CA with states q and q' , where q is initial and q' is final, with the single transition (q, down, q') .
- If $\alpha = \text{right}$, then A is the CA with states q and q' , where q is initial and q' is final, with the single transition (q, right, q') .
- If $\alpha = [\varphi]$, then A is the CA with states q and q' where q is initial and q' is final, with the single transition $(q, [\varphi], q')$.
- If $\alpha = \alpha_1 \cdot \alpha_2$, then let A_1 be the CA constructed by induction for α_1 , and A_2 the CA constructed by induction for α_2 . Take $A = A_1 \cdot A_2$, the CA obtained by the usual concatenation construction on A_1 with A_2 . (That is, we take the disjoint union of A_1 and A_2 , renaming states where necessary, and link the final states of A_1 to the initial state of A_2 by means of an ε -transition. The initial state is the initial state of A_1 ; the final states are the final states of A_2 .)
- If $\alpha = \alpha_1 + \alpha_2$, then let A_1 be the automaton constructed by induction for α_1 and A_2 the automaton

constructed for α_2 . Then take A to $A_1 + A_2$, the automaton obtained by performing the usual union construction on automata. (Take their disjoint union, renaming states where necessary, add a new initial state and add epsilon transition from this state to the initial states of B_1 and B_2 , respectively. The final states consists of the final states of B_1 and B_2 .)

- If $\alpha = \beta^*$, then let B be the automaton created for β . Let A be the automaton we obtain by adding an ε -loop from the final states of B to its initial state.

In all cases, it is now routine to check that A defines α and that $\text{lvl}(A) \leq \max(0, \text{lvl}(\alpha) - 1)$. \square

LEMMA B.6. *Every forward coordinate expression of the form $\alpha(\varphi)$ is definable by means of a CA of level at most $\text{lvl}(\alpha(\varphi)) - 1$.*

PROOF. Observe that $\alpha(\varphi)$ is equivalent to $\beta(\text{root})$ where β is $(\text{right} + \text{down})^* \cdot [\varphi] \cdot \alpha$. By Lemma B.5, there exist a CA A defining β of level $\text{lvl}(\beta) - 1 = \max(\text{lvl}(\varphi) + 1, \text{lvl}(\alpha)) - 1 = \max(\text{lvl}(\varphi), \text{lvl}(\alpha) - 1) \leq \text{lvl}(\alpha(\varphi)) - 1$. Now observe that, since $\alpha(\varphi) \equiv \beta(\text{root})$, it immediately follows that A defines $\alpha(\varphi)$. (Recall that a CA expresses a CA if it gives the same result when evaluation starts from root coordinate $(1,1)$.) \square

LEMMA B.7. *Every forward coordinate expression is definable by a CA of the same level.*

PROOF. Note that every forward coordinate expression φ is equivalent to $\beta(\text{root})$ where $\beta = (\text{right} + \text{down})^* \cdot [\varphi]$. By Lemma B.5, there exists a CA A defining β of level $\text{lvl}(\beta) - 1 = \text{lvl}(\varphi)$. Now observe that, since $\varphi \equiv \beta(\text{root})$, it immediately follows that A defines φ . (Recall that a CA expresses a CA if it gives the same result when evaluation starts from root coordinate $(1,1)$.) \square

Convention. In what follows, by Turing Machine we understand a Turing Machine with a read-only input tape, a read-write work tape, and a write-only output tape where the cursor on the input tape can only advance to the right (never go left). We say that Turing machine M implements CA A if, for every tokenized table T of dimension $n \times m$ and event stream $s = \sigma_1 \cdots \sigma_k$ of T it is the case that M outputs on its output tape the coordinates of the cells selected by A on σ .

PROPOSITION B.8. *For every CA A there exists a Turing Machine M that implements A and that uses at most $O(m \log(m) + \log(n))$ space on its work tape, where m is the number of columns and n is the number of rows in the input table event stream.*

PROOF. Let $A = (Q, q_0, F, \delta)$. Let Ω be the set of all oracles used in transitions of A , i.e., $\Omega = \{\varphi \mid (q, [\varphi], q') \in \Delta, q, q' \in Q\}$. We construct M by induction on the level λ of A .

Base case. If the level λ of A is 0, then M operates as follows. It processes the event stream on its input tape from left to right, one event at a time, starting at the first event. During the processing it maintains on its work tape a tuple $(k, \ell, Act, Susp, \Phi)$, where (k, ℓ) is the coordinate corresponding to the current event being processed, Act is the set of all active states that A can be in any run of A after having processed the events so far, $Susp$ is the set of all

suspended states that A can be in in any run after having processed the events so far, and Φ is the set of all oracles in Ω that select the current coordinate (k, ℓ) . As such, M simulates all possible runs of A on the input, similarly to how one normally simulates a non-deterministic finite state automaton by tracking all possible runs at once.

In particular, M starts with the tuple $(1, 0, \{q_0\}, \emptyset, \emptyset)$ on its work tape before processing any event. When processing the next event, M checks whether this is of the form $\langle \text{cell } \Gamma \rangle$ or $\langle \text{new row} \rangle$. If it is of the form $\langle \text{cell } \Gamma \rangle$ then it:

- 1) Increments ℓ ;
- 2) Computes Φ for the new coordinate (k, ℓ) . Note that this can be done using only the information in Γ (the tokens of the current event) and the coordinate (k, ℓ) . Indeed: since all oracles $\varphi \in \Omega$ are forward and level 0, they cannot contain subexpressions of the form $\langle \alpha \rangle$ or $\alpha(\psi)$. Therefore, each such φ is a boolean combinations of literals, where each literal is either (1) a token a , (2) the constant **true**, or (3) the root coordinate **root**. Checking (2) is trivial, whereas checking (1) amounts to checking whether $a \in \Gamma$ and (3) amounts to comparing the coordinate (k, ℓ) with $(1,1)$.
- 3) It replaces Act by $\{q' \mid q \in Act, (q, \text{right}, q') \in \delta\} \cup \{q \mid q: \ell \in Susp\}$. It then adds to Act all states p that can be reached from a state in this new Act by traversing only ε -transitions or $[\varphi]$ -transitions, with $\varphi \in \Phi$.
- 4) Finally, it updates $Susp$ to

$$\{q' : \ell \in Q \mid q \in Act, (q, \text{down}, q') \in \Delta\}.$$

If the current event is $\langle \text{new row} \rangle$ then it :

- 1) increments k and resets ℓ to 1;
- 2) moves to the next event on the input tape, which must be of the form $\langle \text{cell } \Gamma \rangle$;⁹
- 3) Computes Φ for the new coordinate (k, ℓ) .
- 4) Replaces Act by $\{q' \mid q \in Act, (q, \langle \text{new row} \rangle, q') \in \delta\} \cup \{q' \mid q': 1 \in Susp\}$. It then adds to Act all states p that can be reached from a state in this new Act by traversing only ε -transitions or $[\varphi]$ -transitions, with $\varphi \in \Phi$.
- 5) Finally, it updates $Susp$ to

$$\{q' : \ell \in Q \mid q \in Act, (q, \text{down}, q') \in \Delta\}.$$

After updating $(k, \ell, Act, susp, \Phi)$, M checks if $Act \cap F \neq \emptyset$. If so, it outputs (k, ℓ) .

Observe that the space required by M is:

- $\log(n) + \log(m)$ bits to store the coordinate (k, ℓ) ;
- at most $|Q|$ bits for storing Act ;
- at most $|Q|m \log(m)$ bits for storing $Susp$;
- at most $|\Omega|$ bits for storing Φ

Hence, since $|Q|$ and $|\Omega|$ are constant, M runs in space $O(m \log(m) + \log(n))$, as desired.

Induction step. If the level λ of A is > 0 , then let N be the set of all coordinate expressions of the form $\alpha(\varphi)$ that occur as a subexpression of some oracle in Ω . By definition, each of these $\alpha(\varphi)$ is of level at most λ . By Lemma B.6, there hence exists for each such $\alpha(\varphi)$ a CA $A_{\alpha(\varphi)}$ of level at most $\lambda - 1$ that expresses it. By induction hypothesis, there hence exists, for each $\alpha(\varphi) \in N$, a Turing Machine $M_{\alpha(\varphi)}$ that implements $\alpha(\varphi)$ in space $O(m \log(m) + \log(n))$.

We then construct the Turing Machine M for A as follows. For ease of exposition, M will have multiple (but a fixed

⁹since we assume that all tables have at least one column.

number of) work tapes. Since each of these will use only $O(m \log(m) + \log(n))$ cells, it is standard to transform M in a single-tape Turing Machine that runs in space $O(m \log(m) + \log(n))$.

In particular, M has a $|N|+1$ work tapes: a principal work tape and an auxiliary work tape for each $\alpha(\varphi) \in N$. During processing, M simulates A on its principal work tape, and, in parallel, the Turing Machine for $M_{\alpha(\varphi)}$ on the auxiliary tape for $\alpha(\varphi)$. During the simulation of $M_{\alpha(\varphi)}$ we take care to never construct any output, but merely checks whether the current cell on the input tape should be output according to $M_{\alpha(\varphi)}$.

The simulation of A on its principal work tape happens in the exact same way as for the case where $\lambda = 0$. That is, we maintains a tuple $(k, \ell, Act, Susp, \Phi)$ for A that simulates all possible runs of A . The only difference is that when we update this tuple in response to reading a new event from the input, we first update all the auxiliary work tapes, and then compute the set $\Phi \subseteq \Omega$ of all of A 's oracles that select the current cell (k, ℓ) as follows.

Since each $\psi \in \Omega$ is forward, each ψ is a boolean combination of *literals*, where each literal is either (1) a with $a \in \Delta$; (2) **true**; (3) **root**; or (4) $\alpha(\varphi)$. Cases (1)–(3) can be checked as before. We can check whether current coordinate (k, ℓ) is selected by $\alpha(\varphi)$ simply by looking at the simulation of $M_{\alpha(\varphi)}$ on the work tape for $\alpha(\varphi)$ and verify whether $M_{\alpha(\varphi)}$ would output the current coordinate. As such, we can easily compute at any given instant whether the current coordinate is selected by ψ .

As before, after the update of $(k, \ell, Act, Susp, \phi)$, M checks whether $Act \cap F \neq \emptyset$ and, if so, writes (k, ℓ) on its output tape.

Now note that, as before, M uses $O(m \log(m) + \log(n))$ space on its principal work tape, and (by induction hypothesis) $O(m \log(m) + \log(n))$ on each of its auxiliary work tapes. It hence uses $O(m \log(m) + \log(n))$ space in total, as desired. \square

From this, we derive Theorem 4.4 as follows.

PROOF OF THEOREM 4.4. Let R be a tabular schema, i.e. a set of rules of the form $\varphi \rightarrow e$ with φ a coordinate expression and e a content expression.

By Lemma B.7 every forward coordinate expression can be expressed by means of a CA which, by Proposition B.8, can be evaluated by a Turing Machine in space $O(m \log(m) + \log(n))$, where m is the number of columns in the input, and n the number of rows.

From this, we construct a Turing Machine M that validates its input event stream w.r.t. R as follows. M has a fixed number of work tapes. In particular, for each coordinate expression φ serving as the left-hand side of a rule in R , M has one tape on which it simulates the Turing Machine that evaluates φ . Here, M prevents any output that may be generated by the Turing Machine for φ ; but records when this machine would be doing so.

All of these left-hand-sides are simulated in parallel upon reading the event stream. In addition, for each right-hand side c , M has an auxiliary tape on which it simulates a finite state automaton for c . For each rule $\varphi \rightarrow c$, and in each position in the event stream of the form $\langle \text{cell } \Gamma \rangle$, whenever it finds that φ would select the current cell, M simulates reading Γ in the NFA for c : it proceeds from the current state set for c according to all $\tau \in \Gamma$.

Under the row-based semantics, whenever we encounter a **(new row)**, each c must be in a finite state, otherwise the input is invalid w.r.t. R . Whenever we see **(new row)**, we move each of the NFAs for c back to their initial state.

Under the region-base semantics, each c has to be in a final state at the end of the input (and we never need to move c back to its initial state upon **(new row)**). \square

C. PROOF OF THEOREM 4.6

THEOREM 4.6. *Guarded forward core-SCULPT schemas are strongly streamable.*

PROOF SKETCH. We can use the algorithm of Theorem 4.4 for weak streamability with the additional observation that, for guarded schemas, the columns that need to be remembered when going from one row to the next are independent of the width of the table T . More precisely, when going from one row to the next, we can store for each subexpression φ of a region selection expression a set of pairs P_φ of the form (c, e) where c is a column number and $e \in \{=, \geq\}$. The semantics is that, on the next row r , we need to continue the evaluation of this subexpression in the cells $\{(r, c) \mid (c, =) \in P_\varphi\} \cup \{(r, i) \mid (c, \geq) \in P_\varphi \text{ and } i \geq c\}$. The size of each such set P_φ can be bounded by $2^{O(|\varphi|)}$, which is formalized in Lemma C.1. \square

For a coordinate expression φ , table T , and a row number r we denote by $\llbracket \varphi \rrbracket_{T,r}$ the cells of $\llbracket \varphi \rrbracket_T$ in row r . That is, $\llbracket \varphi \rrbracket_{T,r} = \llbracket \varphi \rrbracket_T \cap \{(r, k) \mid k \in \mathbb{N}\}$.

In the next lemma, a *right-open interval* (on a row r) is a set of cells S for which there exists a $k \in \mathbb{N}$ such that $S = \{(r, i) \mid i \geq k\}$.

LEMMA C.1. (a) *If a coordinate expression φ is row-guarded then, for each table T and row coordinate r , the set $\llbracket \varphi \rrbracket_{T,r}$ consists of $2^{O(|\varphi|)}$ cells plus, optionally, a right-open interval on r .*

(b) *If a coordinate expression φ is guarded then, for each table T and row coordinate r , the set $\llbracket \text{down}^*(\varphi) \rrbracket_{T,r}$ consists of $2^{O(|\varphi|)}$ cells plus, optionally, a right-open interval on r .*

PROOF. The lemma is proved by a straightforward induction on forward coordinate and navigational expressions. We provide it in full detail for the sake of completeness.

For a coordinate expression φ and a row r , we denote by $|\llbracket \varphi \rrbracket_r|$ the number of coordinates of $\llbracket \varphi \rrbracket$ in row r , that is, the number of elements in $\{(r, k) \in \llbracket \varphi \rrbracket \mid k \in \mathbb{N}\}$ if it is finite, and ∞ otherwise.

(a) We prove this case by induction on the definition of guardedness in forward coordinate expressions. The induction base cases for coordinate expressions φ are $\varphi = a$, $\varphi = \text{root}$, $\varphi = \text{true}$, and $\varphi = \text{right}^*(\psi)$.

In the first case, if $\varphi = a$ is row-guarded then it only appears at most once per row by definition of the predicate **unique-per-row**. Therefore, for each row r , $|\llbracket \varphi \rrbracket_r| \leq 1$ and (a) is fulfilled.

In the second case, if $\varphi = \text{root}$, then $\llbracket \varphi \rrbracket$ only contains a single cell. Again, for each row r , $|\llbracket \varphi \rrbracket_r| \leq 1$ and (a) is fulfilled.

In the third case, if $\varphi = \text{true}$, then, for each row r , $|\llbracket \varphi \rrbracket_r| = \infty$. Moreover, $\llbracket \varphi \rrbracket_r = \{(r, k) \mid k \geq 1\}$, which is a right-open interval. Again, (a) is fulfilled.

Fourth, when $\varphi = \text{right}^*(\psi)$, we have that ψ is a boolean combination of **root**, **true**, and **tokens**. Therefore, it can be decided whether a cell c is in $\llbracket \psi \rrbracket$ by looking at the predicates of c and by testing whether it is cell $(1, 1)$ or not. Such tests can be made by inspecting c alone. For each row r , we either have that $\llbracket \psi \rrbracket_r$ is empty or not. If it is empty, then $\llbracket \varphi \rrbracket_r$ is also empty and (a) follows. If it is non-empty, then $\llbracket \varphi \rrbracket_r = \{(r, j) \mid j \geq i\}$, where i is minimal such that $(r, i) \in \llbracket \psi \rrbracket_r$. Since this is a right-open interval, (a) follows.

For the inductive step, we consider the cases $\varphi = \psi_1 \vee \psi_2$, $\varphi = \psi_1 \wedge \psi_2$, $\varphi = \varepsilon(\psi)$, $\varphi = \text{down}(\psi)$, $\varphi = \text{right}(\psi)$, $\varphi = \text{down}(\alpha(\psi))$, $\varphi = \text{right}(\alpha(\psi))$, and $\varphi = (\alpha + \beta)(\psi)$.

In the first case, for each row r , we have that $\llbracket \varphi \rrbracket_r = \llbracket \psi_1 \rrbracket_r \cup \llbracket \psi_2 \rrbracket_r$. Here (a) follows immediately from the inductive hypothesis and the observation that the union of two right-open intervals is again a right-open interval.

The second case is analogous to the first, but we observe that also the intersection of two right-open intervals is again a right-open interval.

Third, when $\varphi = \varepsilon(\psi)$ case (a) follows immediately by induction and the observation that $\llbracket \varepsilon(\psi) \rrbracket_r = \llbracket \psi \rrbracket_r$ for each r .

Cases four to seven are similar to case three. For example, when $\varphi = \text{down}(\psi)$, then, for $r = 1$ we have $\llbracket \varphi \rrbracket_r = \emptyset$ and for $r > 1$ we have $\llbracket \varphi \rrbracket_r = \{(i, r) \mid (i, r-1) \in \llbracket \psi \rrbracket_{r-1}\}$. Again (a) follows by induction.

In the last case we have $\varphi = (\alpha + \beta)(\psi)$ and we already know that $\alpha(\psi)$ and $\beta(\psi)$ are row-guarded. Here, we again have that $\llbracket \varphi \rrbracket_r = \llbracket \alpha(\psi) \rrbracket_r \cup \llbracket \beta(\psi) \rrbracket_r$. Here, (a) follows by induction in a similar way as the first inductive case. However, we need to take a little bit more care about the finite part of $\llbracket \varphi \rrbracket_r$. Assume w.l.o.g. that $|\alpha| \geq |\beta|$ (if not, the roles of α and β can be interchanged in the following). Then, the finite part of $\llbracket \varphi \rrbracket_r$ consists of $2^{O(|\alpha|+|\psi|)} + 2^{O(|\beta|+|\psi|)}$ cells, which is at most $2 \cdot 2^{O(|\alpha|+|\psi|)}$ cells and which, in turn, is bounded from above by $2^{O(|\varphi|)}$ cells, which also proves (a) in this case.

This concludes the proof of Lemma C.1(a).

(b) We again proceed by induction on the definition of guardedness in forward coordinate expressions. The induction base cases for coordinate expressions φ are $\varphi = a$, $\varphi = \text{root}$, $\varphi = \text{true}$, and $\varphi = \text{right}^*(\psi)$.

In the first case, if $\varphi = a$ is guarded then it only appears at most once in the table by definition of the predicate **unique**. If $\llbracket \varphi \rrbracket$ is empty, then $\text{down}^*(\varphi)$ is empty in which case (b) is fulfilled. If not, we have that $\llbracket \varphi \rrbracket = \{(i, j)\}$ for some i and j . Here, we have that $\llbracket \text{down}^*(\varphi) \rrbracket = \{(k, j) \mid k \geq i\}$, which fulfils the conditions of case (b).

In the second case, if $\varphi = \text{root}$, then $\llbracket \varphi \rrbracket$ only contains a single cell. Here, we have that $\llbracket \text{down}^*(\varphi) \rrbracket = \{(k, 1) \mid k \geq 1\}$, which fulfils (b).

In the third case, if $\varphi = \text{true}$, then, $\llbracket \text{down}^*(\varphi) \rrbracket$ contains every cell in the table. Clearly, this is a right-open interval for each row and (b) is fulfilled.

In the fourth case, if $\varphi = \text{right}^*(\psi)$ we can decide for each individual cell c whether $c \in \llbracket \varphi \rrbracket$ by only inspecting c , analogously as in (a). Also analogously, For each row r , we either have that $\llbracket \varphi \rrbracket_r$ is empty or a right-open interval of the form $\llbracket \varphi \rrbracket_r = \{(r, j) \mid j \geq i\}$, where i is minimal such that $(r, i) \in \llbracket \psi \rrbracket_r$. Therefore, $\llbracket \text{down}^*(\varphi) \rrbracket_r$ is also empty or a right-open interval on each row r .

For the inductive step, we consider the cases $\varphi = \psi_1 \vee \psi_2$, $\varphi = \psi_1 \wedge \psi_2$, $\varphi = \varepsilon(\psi)$, $\varphi = \text{down}(\psi)$, $\varphi = \text{right}(\psi)$, $\varphi =$

$\text{down}(\alpha(\psi))$, $\varphi = \text{right}(\alpha(\psi))$, $\varphi = (\alpha + \beta)(\psi)$, $\varphi = \text{down}^* \cdot \alpha(\psi)$, and $\varphi = \text{right}^* \cdot \alpha(\psi)$.

In the first case, we have that $\text{down}^*(\psi_1)$ and $\text{down}^*(\psi_2)$ are guarded by induction, so they fulfil condition (b). Furthermore, we have that $\llbracket \text{down}^*(\psi_1 \vee \psi_2) \rrbracket = \llbracket \text{down}^*(\psi_1) \rrbracket \cup \llbracket \text{down}^*(\psi_2) \rrbracket$. This means that, for each row r , we also have that $\llbracket \varphi \rrbracket_r = \llbracket \psi_1 \rrbracket_r \cup \llbracket \psi_2 \rrbracket_r$. Here (b) follows immediately from the inductive hypothesis and the observation that the union of two right-open intervals is again a right-open interval.

The second case is analogous to the first, but takes intersections of right-open intervals instead of unions.

Case three, where $\varphi = \varepsilon(\psi)$, (b) follows immediately from the guardedness of ψ and the observation that $\llbracket \varepsilon(\psi) \rrbracket_r = \llbracket \psi \rrbracket_r$ for each r .

Cases four to eight are similar to case three. In case eight, where $\varphi = (\alpha + \beta)(\psi)$, we can bound the number of cells per row of $\llbracket \text{down}^*(\alpha + \beta)(\psi) \rrbracket$ that are not in the right-open interval in exactly the same way as for the analogous case in (a).

In case nine, we have that $\varphi = \text{down}^* \cdot \alpha(\psi)$ and $\alpha(\psi)$ is guarded. This case immediately follows from the induction hypothesis since $\llbracket \text{down}^* \varphi \rrbracket = \llbracket \text{down}^* \text{down}^* \alpha(\psi) \rrbracket = \llbracket \text{down}^* \alpha(\psi) \rrbracket$.

Finally, in the last case, we have $\varphi = \text{right}^* \cdot \alpha(\psi)$, where $\alpha(\psi)$ is row-guarded. Due to the row-guardedness of $\alpha(\psi)$, we know by (a) that $\llbracket \alpha(\psi) \rrbracket$ for each row r consists of $2^{O(|\varphi|)}$ cells in r plus, optionally, a right-open interval on r . Therefore, $\text{right}^* \cdot \alpha(\psi)$ on a row r is either empty, or a right-open interval. From this, we immediately have that also $\text{down}^* \text{right}^* \alpha(\psi)$, on each row r , is either empty or a right-open interval. This concludes the proof of Lemma C.1 (b). \square