

Static Analysis of String Values

Giulia Costantini¹, Pietro Ferrara², and Agostino Cortesi¹

¹ University Ca' Foscari of Venice, Italy

{costantini,cortesi}@dsi.unive.it

² ETH Zurich, Switzerland

pietro.ferrara@inf.ethz.ch

Abstract. In this paper we propose a unifying approach for the static analysis of string values based on abstract interpretation, and we present several abstract domains that track different types of information. In this way, the analysis can be tuned at different levels of precision and efficiency, and it can address specific properties.

1 Introduction

Strings are widely used in modern programming languages. Their applications vary from providing an output to a user to the construction of programs executed through reflection. For instance, in PHP strings can be a way of communicating programs, while in Java they are widely used as SQL queries, or to access information about the classes through reflection. The execution of `str.substring(str.indexOf('a'))` raises an exception if `str` does not contain an 'a' character: in this case, it would be useful being able to track the characters surely contained on the variable `str`. As another example, when dealing with SQL queries, what happens if we execute the query “DELETE FROM `Table` WHERE `ID =` + `id` when `id` is equal to “10 OR TRUE”? The content of `Table` would be permanently erased! It's clear that a wrong manipulation of strings could lead not only to subtle exceptions, but to dramatic and permanent effects as well [20].

For all these reasons, the interest on approaches that automatically analyse and discover bugs on strings is constantly raising. On the other hand, the state-of-the-art in this field is still limited: approaches that rely on automata and use regular expressions are precise but slow, and they do not scale up [14,24,21,13], while many other approaches are focused on particular properties or class of programs [10,18,12]. Genericity and scalability are the main advantages of the abstract interpretation approach [4,5], though its instantiation to textual values has been quite limited up to now.

The main contribution of this paper is the formalisation of a *unifying* abstract interpretation based framework for string analysis, and its instantiations with *four different domains* that track distinct types of information. In this way, we can tune the analysis at diversified levels of *accuracy*, yielding to faster and rougher, or slower but more precise string analyses.

<pre> 1 var query = "SELECT '\$\\$\\$' 2 (RETAIL/100) FROM INVENTORY WHERE "; 3 if (l != null) 4 query = query+"WHOLESALE > "+l+" AND "; 5 6 var per = "SELECT TYPECODE, TYPEDESC FROM 7 TYPES WHERE NAME = 'fish' OR NAME = 'meat'"; 8 query = query+"TYPE IN (" + per + ");"; 9 return query; </pre> <p style="text-align: center;">(a) The first running example</p>	<pre> 1 string x = "a"; 2 while(cond) 3 x = "0" + x + "1"; 4 return x; </pre> <p style="text-align: center;">(b) The second running example</p>
---	---

Fig. 1. The running examples

We inspired our work looking at the approach adopted for numerical domains for static analysis of software [7,11,19]. The interface of a numerical domain is nowadays standard: each domain has to define the semantics of arithmetic expressions (like $i + 5$) and boolean conditions (like $i < 5$). Similarly, we consider a limited list of basic string operators that can be easily extended to the various programming languages. The concrete semantics of these operators is approximated in the four different abstract domains. In addition, after 30 years of practice with numerical domains, it is clear that a monolithic domain precise on any program and property (e.g., Polyhedra [7]) gives up in terms of efficiency, while to achieve scalability we need specific approximations on a given property (e.g., Pentagons [17]) or class of programs (e.g., ASTRÉE [6]). With this scenario in mind, we develop several domains inside the same framework to tune the analysis at different levels of precision and efficiency w.r.t. the analysed program and property. Other abstractions are possible and welcomed, and we expect our framework to be generic enough to support them.

The paper is structured as follows. In the rest of this Section we introduce two running examples, and we recall some basics of abstract interpretation. Section 2 defines the syntax of the string operators we will consider. Section 3 introduces the concrete semantics, while in Section 4 the abstract domains are formalised. Finally, Section 5 discusses the related work, and Section 6 concludes.

1.1 Running Examples

Along the paper, we will always refer to the two examples reported in Tables 1(a) and 1(b). The first Java program is taken from [10], and it dynamically builds an SQL query by concatenating some strings. One of these concatenations applies only if a certain value (unknown at compile time) is not null. We are interested in checking if the SQL query resulting by the execution of such code is well formed. For the sake of readability, we will use some shortcuts to identify string constants of this program, as reported in Table 1. The second program modifies a string inside a `while` loop whose condition cannot be statically evaluated. Therefore, we will need to apply a widening operator [2] to force the convergence of the analysis. Intuitively, this program produces strings in the form $0^n a 1^n$.

Table 1. Shortcuts of string constants in the first running example

Name	String constant
s ₁	“SELECT ‘\$’ (RETAIL/100) FROM INVENTORY WHERE ”
s ₂	“WHOLESALE > ”
s ₃	“ AND ”
s ₄	“SELECT TYPECODE, TYPEDESC FROM TYPES WHERE NAME = ‘fish’ OR NAME = ‘meat’”
s ₅	“TYPE IN (“
s ₆	“);”

1.2 Abstract Interpretation

Abstract interpretation is a theory to define and soundly approximate the semantics of a program [4,5], focusing on some runtime properties of interest. Usually, each concrete state is composed by a set of elements (e.g., all the possible computational states), that is approximated by an unique element in the abstract domain. Formally, the concrete domain $\wp(D)$ forms a complete lattice $\langle \wp(D), \subseteq, \emptyset, D, \cup, \cap \rangle$. On this domain, a semantics \mathbb{S} is defined. In the same way, an abstract semantics is defined, and it is aimed to approximate the concrete one in a computable way. Formally, the abstract domain \bar{A} has to form a complete lattice $\langle \bar{A}, \leq_{\bar{A}}, \perp_{\bar{A}}, \top_{\bar{A}}, \sqcup_{\bar{A}}, \sqcap_{\bar{A}} \rangle$. The concrete elements are related to the abstract domain by a concretization $\gamma_{\bar{A}}$ and an abstraction $\alpha_{\bar{A}}$ functions. In order to obtain a sound analysis, we require that the abstraction and concretization functions above form a Galois connection. An abstract semantics $\bar{\mathbb{S}}$ is defined as a sound approximation of the concrete one, i.e., $\forall \bar{a} \in \bar{A} : \alpha_{\bar{A}} \circ \mathbb{S}[\gamma_{\bar{A}}(\bar{a})] \leq_{\bar{A}} \bar{\mathbb{S}}[\bar{a}]$.

When abstract domains do not satisfy the ascending chain condition, a widening operator $\nabla_{\bar{A}}$ is required in order to guarantee the convergence of the fixed point computation. This is an upper bound operator such that for all increasing chains $\bar{a}_0 \leq_{\bar{A}} \dots \bar{a}_n \leq_{\bar{A}} \dots$ the increasing chain defined as $\bar{w}_0 = \bar{a}_0, \dots, \bar{w}_{i+1} = \bar{w}_i \nabla_{\bar{A}} \bar{a}_{i+1}$ is not strictly increasing.

2 Syntax

Different languages define different operators on strings, and usually each language supports a huge set of such operators: in Java 1.6 the `String` class contains 65 methods plus 15 constructors, `System.Text` in .Net contains about 12 classes that work with Unicode strings, and PHP provides 111 string functions. Considering all these operators would be quite verbose, and in addition the most part of them perform similar actions using slightly different data. We restrict our description on a small but representative set of common operators. We chose these operators looking at some case studies. Other operators can be easily added to our approach. For each operator, this would mean to define its concrete semantics, and its approximations on the different domains we will introduce.

Table 2. The concrete semantics, where \top_B represents that the condition could be evaluated to true or false depending on the string in S_1 we are considering

$$\begin{aligned}
\mathbb{S}[\text{new String}(\text{str})]() &= \{\text{str}\} \\
\mathbb{S}[\text{concat}](S_1, S_2) &= \{s_1s_2 : s_1 \in S_1 \wedge s_2 \in S_2\} \\
\mathbb{S}[\text{readLine}]() &= S \\
\mathbb{S}[\text{substring}_b^e](S_1) &= \{c_b..c_e : c_1..c_n \in S_1 \wedge n \geq e \wedge b \leq e\} \\
\mathbb{B}[\text{contains}_c](S_1) &= \begin{cases} \text{true} & \text{if } \forall s \in S_1 : c \in \text{char}(s) \\ \text{false} & \text{if } \forall s \in S_1 : c \notin \text{char}(s) \\ \top_B & \text{otherwise} \end{cases}
\end{aligned}$$

A common operation is the creation of a new constant string (`new String(str)` where `str` is a sequence of characters). Usually programs concatenate strings (`concat(s1,s2)` where `s1` and `s2` are strings), read inputs from the user (`readLine()`), and take a substring of a given string (`substring_b^e(s)`, where `s` is a string, and `b` and `e` are integer values) as well. A common test is to check if a string contains a character (`contains_c(s)`, where `s` is a string and `c` is a character).

3 Concrete Domain and Semantics

3.1 Concrete Domain

Our concrete domain is simply made of strings. Given an alphabet K , that is a finite set of characters, we define strings as (possibly infinite) sequences of characters. Formally, $S = K^*$, where A^* is an ordered sequence of elements in A , that is, $A^* = \{a_1 \cdots a_n : \forall i \in [1..n] : a_i \in A\}$. A string variable in our program could have different values in different executions, and our goal is to approximate all these values (potentially infinite, e.g., when dealing with user input) in a finite, computable, and hopefully efficient manner. Our lattice will be made of sets of strings. As usual in abstract interpretation, the partial order is the set inclusion. Formally, our concrete domain is defined by $\langle \wp(S), \subseteq, \emptyset, S, \cup, \cap \rangle$.

3.2 Semantics

Table 2 formalises the concrete semantics. For each statement of the language we introduced in Section 2, we define its semantics. For the first four statements, we define a semantics \mathbb{S} that, given the statement and eventually some sets of concrete string values in S , returns a set of strings resulting from that operation. The semantics of `new String(str)` returns a singleton containing `str`, while the semantics of `readLine` returns a set containing all the possible strings, since we may read any string from the standard input. The semantics of `concat` returns all the possible concatenations of a string taken from the first set and a string taken from the second set (we denote by s_1s_2 the concatenation of strings s_1 and s_2), while the semantics of `substring_b^e` returns all the substrings from the

b -th to e -th character of the given strings (note that if one of the strings is too short, there is not any substrings for it in the resulting set, since this would cause a runtime error without producing any value). For `containsc` we define a particular semantics $\mathbb{B} : [\wp(S) \rightarrow \{\text{true}, \text{false}, \top_{\mathbb{B}}\}]$ that, given a set of strings, returns `true` if all the strings contains the character c , `false` if none contains this character, and $\top_{\mathbb{B}}$ otherwise. This special boolean value represents a situation in which the boolean condition may be evaluated to `true` some times, and to `false` other times. We denoted by `char` a function that returns the set of characters contained in the string in input.

4 Abstract Domains and Semantics

What is the *relevant* information contained in a string? How can we approximate it in an *efficient* way? Tracking both sound and precise information at compile time on strings in an efficient way is infeasible. Then we need to introduce *approximation*. We want to track information precise enough to efficiently analyse the behaviours of interest, considering the string operators we defined in the previous section. Our purpose is to approximate strings as much as we can, preserving the information we deem relevant.

4.1 Character Inclusion

For the first abstract domain we aim at approximating a string through the characters we know it surely contains or it could contain. This information could be particularly useful to track if the indexes extrapolated from a string with operators like `indexOf(c)` could be used to cut the string (because c is surely contained in the string), or they could be invalid (e.g., -1). A string will be represented by a pair of sets: the set of *certainly* contained characters $\overline{\mathcal{C}}$ and the set of *maybe* contained characters $\overline{\mathcal{MC}}$ ($\overline{\mathcal{CI}} = \{(\overline{\mathcal{C}}, \overline{\mathcal{MC}}) : \overline{\mathcal{C}}, \overline{\mathcal{MC}} \in \wp(K) \wedge \overline{\mathcal{C}} \subseteq \overline{\mathcal{MC}}\} \cup \perp_{\overline{\mathcal{CI}}}$). The partial order $\leq_{\overline{\mathcal{CI}}}$ on $\overline{\mathcal{CI}}$ is the following one:

$$(\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1) \leq_{\overline{\mathcal{CI}}} (\overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_2) \Leftrightarrow (\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1) = \perp_{\overline{\mathcal{CI}}} \vee (\overline{\mathcal{C}}_1 \supseteq \overline{\mathcal{C}}_2 \wedge \overline{\mathcal{MC}}_1 \subseteq \overline{\mathcal{MC}}_2)$$

This is because the more information we have on the string (that is, the more characters are certainly contained and the less characters are maybe contained), the less number of strings we are representing. For example the abstract element represented by the pair $(\{a\}, \{a\})$ is more precise than the one represented by $(\emptyset, \{a, b\})$. In fact, the first pair represents the concrete set of strings $\{a, aa, aaa, \dots\}$ while the second pair corresponds to $\{\epsilon, a, b, aa, bb, ba, ab, \dots\}$.

For these reasons, the least upper bound is defined by $\sqcup_{\overline{\mathcal{CI}}}((\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1), (\overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_2)) = (\overline{\mathcal{C}}_1 \cap \overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_1 \cup \overline{\mathcal{MC}}_2)$, and the greatest lower bound is defined by $\sqcap_{\overline{\mathcal{CI}}}((\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1), (\overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_2)) = (\overline{\mathcal{C}}_1 \cup \overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_1 \cap \overline{\mathcal{MC}}_2)$. The widening operator corresponds to the $\sqcup_{\overline{\mathcal{CI}}}$ operator, and it ensures the convergence of the analysis since we supposed that the alphabet is finite. The top element of the lattice is $\top_{\overline{\mathcal{CI}}} = (\emptyset, K)$, while the bottom element $\perp_{\overline{\mathcal{CI}}}$ corresponds to a “failure” state.

The function which abstracts a single string s is: $\alpha'_{\overline{\mathcal{CI}}}(s) = (\text{char}(s), \text{char}(s))$. The abstraction function takes us from a set of strings to an element in $\overline{\mathcal{CI}}$, and

Table 3. The abstract semantics of $\overline{\mathcal{CI}}$

$$\begin{aligned} \overline{\mathbb{S}}_{\mathcal{CI}}[\text{new String}(\text{str})]() &= (\text{char}(\text{str}), \text{char}(\text{str})) \\ \overline{\mathbb{S}}_{\mathcal{CI}}[\text{concat}]((\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1), (\overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_2)) &= (\overline{\mathcal{C}}_1 \cup \overline{\mathcal{C}}_2, \overline{\mathcal{MC}}_1 \cup \overline{\mathcal{MC}}_2) \\ \overline{\mathbb{S}}_{\mathcal{CI}}[\text{readLine}]() &= (\emptyset, \mathbb{K}) \\ \overline{\mathbb{S}}_{\mathcal{CI}}[\text{substring}_b^s]((\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1)) &= (\emptyset, \overline{\mathcal{MC}}_1) \\ \overline{\mathbb{B}}_{\mathcal{CI}}[\text{contains}_c]((\overline{\mathcal{C}}_1, \overline{\mathcal{MC}}_1)) &= \begin{cases} \text{true} & \text{if } c \in \overline{\mathcal{C}}_1 \\ \text{false} & \text{if } c \notin \overline{\mathcal{MC}}_1 \\ \top_B & \text{otherwise} \end{cases} \end{aligned}$$

#I	Var	$\overline{\mathcal{CI}}$
1	query	$\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_1)$
3	1	(\emptyset, \mathbb{K})
3	query	$(\pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_1)) \cup \pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_2)) \cup \pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_3)), \mathbb{K})$
4	query	$(\pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_1)), \mathbb{K})$
5	per	$\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_1)$
7	query	$(\pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_1)) \cup \pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_4)) \cup \pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_5)) \cup \pi_1(\alpha'_{\overline{\mathcal{CI}}}(\mathbf{s}_6)), \mathbb{K})$

(a) First running example

#I	Var	$\overline{\mathcal{CI}}$
1	x	$(\{a\}, \{a\})$
3	x	$(\{0, a, 1\}, \{0, a, 1\})$
4	x	$(\{a\}, \{0, a, 1\})$

(b) Second running example

Fig. 2. The results of $\overline{\mathcal{CI}}$

it returns the upper bound of the abstraction of all the concrete strings. Let π_i be the projection on the i-th component of a tuple.

$$\alpha_{\overline{\mathcal{CI}}}(\mathbf{S}_1) = \bigsqcup_{\overline{\mathcal{CI}}, s \in \mathbf{S}_1} \alpha'_{\overline{\mathcal{CI}}}(s) = (\bigcap_{s \in \mathbf{S}_1} \pi_1(\alpha'_{\overline{\mathcal{CI}}}(s)), \bigcup_{s \in \mathbf{S}_1} \pi_2(\alpha'_{\overline{\mathcal{CI}}}(s)))$$

Semantics. Table 3 defines the abstract semantics of the operators introduced in Section 2 on the abstract domain $\overline{\mathcal{CI}}$. We denote by $\overline{\mathbb{S}}_{\mathcal{CI}}$ and $\overline{\mathbb{B}}_{\mathcal{CI}}$ the abstract counterparts of \mathbb{S} and \mathbb{B} respectively.

When we evaluate a string, we know that the characters that are surely or maybe included are exactly the ones that appear in the string. The concatenation of two strings will contain all the characters that are surely or maybe contained in the two strings. `readLine` returns a top value, while if we take a substring of a given string, the result will possibly contain all the characters that are possibly contained in the initial string, while we know nothing about the surely contained characters. Finally, the semantics of `containsc` is quite precise, as it checks if a character is surely contained or not contained respectively through $\overline{\mathcal{C}}$ and $\overline{\mathcal{MC}}$.

Running Example. Consider the examples introduced in Section 1.1. The results of the analysis of the first program using $\overline{\mathcal{CI}}$ are depicted in Figure 2(a). At the beginning, variable `query` is related to a state that contains the abstraction of \mathbf{c}_1 , that is, both $\overline{\mathcal{C}}$ and $\overline{\mathcal{MC}}$ contain all the characters of \mathbf{s}_1 . Since we do not know the value of `1`, we compute the least upper bound between the abstract values of `query` after instructions 1 and 3. In this way, we obtain that after the `if` statement the abstract value of `query` contains the abstraction of \mathbf{s}_1 in the $\overline{\mathcal{C}}$ component (since it surely contains all the characters of that constant string), and the top value in the

\overline{MC} component (since we may have concatenated a string that may contain any character). At the end of the given code, `query` surely contains the characters of s_1, s_4, s_5 , and s_6 , and it may contain any character, since we possibly concatenated in `query` an input string (the `l` variable).

As for the second program, in Figure 2(b) we see that after instruction 1 `x` surely contains ‘a’. Inside the loop (line 3), `x` surely contains ‘a’, ‘0’ and ‘1’. In line 4 we report the least upper bound between the value of `x` *before* entering the loop (line 1) and the value *after* the loop (line 4): variable `x` surely contains the character ‘a’, and it also may contain the characters ‘0’ and ‘1’.

4.2 Prefix and Suffix

The next abstract domain we consider approximates strings by their *prefix*. A string will be a sequence of characters which *begins* with a certain sequence of characters and ends with any string (we use $*$ to represent any string, ϵ included). For example, $abc*$ represents all the strings which begin with “abc”, including “abc” itself. Since the asterisk $*$ at the end of the representation is always present, we do not include it in the domain and consider abstract elements made only of sequence of characters: $\overline{\mathcal{PR}} = K^* \cup \perp_{\overline{\mathcal{PR}}}$. The partial order on this domain is:

$$\overline{S} \leq_{\overline{\mathcal{PR}}} \overline{T} \Leftrightarrow \overline{S} = \perp_{\overline{\mathcal{PR}}} \vee (\forall i \in [0, \text{len}(\overline{T}) - 1] : \text{len}(\overline{T}) \leq \text{len}(\overline{S}) \wedge \overline{T}[i] = \overline{S}[i])$$

An abstract string \overline{S} is smaller than \overline{T} if \overline{T} is a prefix of \overline{S} or if \overline{S} is the bottom $\perp_{\overline{\mathcal{PR}}}$ of the domain. The least upper bound operator is defined as the longest common prefix of two strings. The greater lower bound is defined by:

$$\sqcap_{\overline{\mathcal{PR}}}(\overline{S}_1, \overline{S}_2) = \begin{cases} \overline{S}_1 & \text{if } \overline{S}_1 \leq_{\overline{\mathcal{PR}}} \overline{S}_2 \\ \overline{S}_2 & \text{if } \overline{S}_2 \leq_{\overline{\mathcal{PR}}} \overline{S}_1 \\ \perp_{\overline{\mathcal{PR}}} & \text{otherwise} \end{cases}$$

The widening operator is simply the upper bound operator above, as the latter converges in finite time. Top and bottom elements are, respectively, ϵ (the empty prefix) and $\perp_{\overline{\mathcal{PR}}}$. The function which abstracts a single string is $\alpha'_{\overline{\mathcal{PR}}}(\mathfrak{s}) = \mathfrak{s}$. The abstraction function is $\alpha_{\overline{\mathcal{PR}}}(\mathcal{S}_1) = \sqcup_{\overline{\mathcal{PR}}.s \in \mathcal{S}_1} \alpha'_{\overline{\mathcal{PR}}}(s)$. This means that we consider the longest common prefix amongst all strings in \mathcal{S}_1 .

We can track information about the *suffix* of a string as well. We define another abstract domain, $\overline{\mathcal{SU}}$, where a string will be something which *ends* with a certain sequence of characters. The notation and all the operators of this domain are dual to those of the previous domain. The definition of the domain is: $\overline{\mathcal{SU}} = K^* \cup \perp_{\overline{\mathcal{SU}}}$. The partial order is:

$$\overline{S} \leq_{\overline{\mathcal{SU}}} \overline{T} \Leftrightarrow \overline{S} = \perp_{\overline{\mathcal{SU}}} \vee (\forall i \in [0, \text{len}(\overline{T}) - 1] : \text{len}(\overline{T}) \leq \text{len}(\overline{S}) \wedge \overline{T}[i] = \overline{S}[i + \text{len}(\overline{S}) - \text{len}(\overline{T})])$$

The least upper bound $\sqcup_{\overline{\mathcal{SU}}}$ is the longest common suffix, while the greatest lower bound $\sqcap_{\overline{\mathcal{SU}}}$ is the smallest suffix (if they are comparable) or $\perp_{\overline{\mathcal{SU}}}$ (if they are not comparable). The widening operator is the least upper bound operator above. The top element is ϵ . The function which abstracts a single string is: $\alpha'_{\overline{\mathcal{SU}}}(\mathfrak{s}) = \mathfrak{s}$, and the abstraction function is $\alpha_{\overline{\mathcal{SU}}}(\mathcal{S}_1) = \sqcup_{\overline{\mathcal{SU}}.s \in \mathcal{S}_1} \alpha'_{\overline{\mathcal{SU}}}(s)$.

These abstract domains could be particularly useful to check if some simple syntactic properties (e.g., a string that is used as an SQL command always begins with “SELECT” and ends with “;”) are respected by all possible executions.

Table 4. The abstract semantics of $\overline{\mathcal{PR}}$

$$\begin{aligned} \overline{\mathcal{SR}}[\text{new String}(\text{str})]() &= \text{str} \\ \overline{\mathcal{SR}}[\text{concat}](\overline{p}_1, \overline{p}_2) &= \overline{p}_1 \\ \overline{\mathcal{SR}}[\text{readLine}]() &= \epsilon \\ \overline{\mathcal{SR}}[\text{substring}_b^e](\overline{p}) &= \begin{cases} \overline{p}[b \dots e - 1] & \text{if } e \leq \text{len}(\overline{p}) \\ \overline{p}[b \dots \text{len}(\overline{p}) - 1] & \text{if } e > \text{len}(\overline{p}) \wedge b < \text{len}(\overline{p}) \\ \epsilon & \text{otherwise} \end{cases} \\ \overline{\mathcal{BR}}[\text{contains}_c](\overline{p}) &= \begin{cases} \text{true} & \text{if } c \in \text{char}(\overline{p}) \\ \top_B & \text{otherwise} \end{cases} \end{aligned}$$

$\begin{aligned} \overline{\mathcal{SU}}[\text{new String}(\text{str})]() &= \text{str} \\ \overline{\mathcal{SU}}[\text{concat}](\overline{s}_1, \overline{s}_2) &= \overline{s}_2 \\ \overline{\mathcal{SU}}[\text{readLine}]() &= \epsilon \\ \overline{\mathcal{SU}}[\text{substring}_b^e](\overline{s}) &= \epsilon \\ \overline{\mathcal{BU}}[\text{contains}_c](\overline{s}) &= \\ &= \begin{cases} \text{true} & \text{if } c \in \text{char}(\overline{s}) \\ \top_B & \text{otherwise} \end{cases} \end{aligned}$	<table style="border-collapse: collapse; text-align: center;"> <thead> <tr> <th>#I</th> <th>Var</th> <th>$\overline{\mathcal{PR}}$</th> <th>$\overline{\mathcal{SU}}$</th> </tr> </thead> <tbody> <tr><td>1</td><td>query</td><td>\overline{s}_1</td><td>\overline{s}_1</td></tr> <tr><td>3</td><td>l</td><td>ϵ</td><td>ϵ</td></tr> <tr><td>3</td><td>query</td><td>\overline{s}_1</td><td>\overline{s}_3</td></tr> <tr><td>4</td><td>query</td><td>\overline{s}_1</td><td>" "</td></tr> <tr><td>5</td><td>per</td><td>\overline{s}_4</td><td>\overline{s}_4</td></tr> <tr><td>7</td><td>query</td><td>\overline{s}_1</td><td>\overline{s}_6</td></tr> </tbody> </table>	#I	Var	$\overline{\mathcal{PR}}$	$\overline{\mathcal{SU}}$	1	query	\overline{s}_1	\overline{s}_1	3	l	ϵ	ϵ	3	query	\overline{s}_1	\overline{s}_3	4	query	\overline{s}_1	" "	5	per	\overline{s}_4	\overline{s}_4	7	query	\overline{s}_1	\overline{s}_6	<table style="border-collapse: collapse; text-align: center;"> <thead> <tr> <th>#I</th> <th>Var</th> <th>$\overline{\mathcal{PR}}$</th> <th>$\overline{\mathcal{SU}}$</th> </tr> </thead> <tbody> <tr><td>1</td><td>x</td><td>a</td><td>a</td></tr> <tr><td>3</td><td>x</td><td>0</td><td>1</td></tr> <tr><td>4</td><td>x</td><td>\top</td><td>\top</td></tr> </tbody> </table>	#I	Var	$\overline{\mathcal{PR}}$	$\overline{\mathcal{SU}}$	1	x	a	a	3	x	0	1	4	x	\top	\top
#I	Var	$\overline{\mathcal{PR}}$	$\overline{\mathcal{SU}}$																																											
1	query	\overline{s}_1	\overline{s}_1																																											
3	l	ϵ	ϵ																																											
3	query	\overline{s}_1	\overline{s}_3																																											
4	query	\overline{s}_1	" "																																											
5	per	\overline{s}_4	\overline{s}_4																																											
7	query	\overline{s}_1	\overline{s}_6																																											
#I	Var	$\overline{\mathcal{PR}}$	$\overline{\mathcal{SU}}$																																											
1	x	a	a																																											
3	x	0	1																																											
4	x	\top	\top																																											

(a) The abstract semantics of $\overline{\mathcal{SU}}$ (b) First running example (c) Second running example

Fig. 3. The abstract semantics of $\overline{\mathcal{SU}}$ and the running examples

Semantics. Table 4 and 3(a) define the abstract semantics on $\overline{\mathcal{PR}}$ and $\overline{\mathcal{SU}}$ respectively. The most precise suffix and prefix of a constant string are the string itself. When we concatenate two strings, we consider as prefix and suffix of the resulting string the abstract value of the left and right operand respectively. As usual, the semantics of `readLine` returns the top value. The same happens for `substringbe` in $\overline{\mathcal{SU}}$, since we do not know how many characters there are before the suffix. Instead, $\overline{\mathcal{PR}}$ can be more precise if *b* (and eventually *e*) are smaller than the length of the prefix we have. Finally, the semantics of `containsc` returns `true` iff *c* is in the prefix or suffix, and \top_B otherwise, since we have no information at all about which characters are after the prefix or before the suffix.

Running Example. The results of the analyses using the prefix and suffix domains on our running examples are reported in Figures 3(b) and 3(c).

For the first program, at line 1, `query` contains the whole string *s*₁ as both prefix and suffix. As already pointed out, `l` is an input of the user, so we do not know what its prefix and suffix are. On the other hand, when we concatenate it at line 3, we still have some information on the prefix and suffix of the resulting string. Thus, at the end of the analyses, we get that the prefix of `query` is string *s*₁, its suffix is *s*₆, although we lose information about what there is in the middle.

For the second program, before entering the loop we know the prefix and suffix of `x`. Inside the loop after line 3, the convergence for `x` is ‘0’ as prefix and ‘1’ as suffix. This state, combined through the lub operator with the state before the loop, unfortunately goes to \top (the longest common prefixes and suffixes are empty), making us lose all the information.

4.3 Bricks

The next abstract domain, $\overline{\mathcal{BR}}$, captures both *inclusion and order* amongst characters, using a simplification of regular expressions. Therefore, the information tracked by this domain could be adopted to prove more sophisticated properties than the previous domains (e.g., the well-formedness of SQL queries). A string is approximated by a combination of *bricks*. A brick is defined as an element of: $\overline{\mathcal{B}} = [\varphi(\mathcal{S})]^{\min, \max}$, where \min and \max are two integer positive values. A brick represents all the strings which can be built through the given strings, taken between \min and \max times altogether. For example, $\{[“mo”, “de”]\}^{1,2} = \{mo, de, momo, dede, mode, demo\}$. We represent strings as ordered lists of bricks. For example we have that $\{[“straw”]\}^{0,1} \{[“berry”]\}^{1,1} = \{berry, strawberry\}$ since $\{[“straw”]\}^{0,1}$ concretizes to $\{\epsilon, “straw”\}$ and $\{[“berry”]\}^{1,1}$ to $\{“berry”\}$. Since a particular set of strings could be represented by more than one combination of bricks, we adopted a normalised form in which the lists are made of bricks like $[\mathbb{T}]^{1,1}$ or $[\mathbb{T}]^{0, \max > 0}$, where \mathbb{T} is a set of strings. We defined a function $\overline{normBricks}(\overline{\mathbb{L}})$ which, given a list of bricks $\overline{\mathbb{L}}$, returns its normalized version.

The abstract domain of bricks is defined as: $\overline{\mathcal{BR}} = \overline{\mathcal{B}}^*$, that is, the set of all finite sequences composed of bricks. The top element $\top_{\overline{\mathcal{BR}}}$ is a list containing only $\top_{\overline{\mathcal{B}}}$. The bottom element is $\perp_{\overline{\mathcal{BR}}}$, an empty list or any list which contains at least one invalid element ($\perp_{\overline{\mathcal{B}}}$). The partial order between single bricks is: $[\overline{\mathcal{C}}_1]^{\min_1, \max_1} \leq_{\overline{\mathcal{B}}} [\overline{\mathcal{C}}_2]^{\min_2, \max_2} \Leftrightarrow (\overline{\mathcal{C}}_1 \subseteq \overline{\mathcal{C}}_2 \wedge \min_1 \geq \min_2 \wedge \max_1 \leq \max_2) \vee [\overline{\mathcal{C}}_2]^{\min_2, \max_2} = \top_{\overline{\mathcal{B}}} \vee [\overline{\mathcal{C}}_1]^{\min_1, \max_1} = \perp_{\overline{\mathcal{B}}}$ where $\top_{\overline{\mathcal{B}}}$ and $\perp_{\overline{\mathcal{B}}}$ are special bricks, respectively greater and smaller than any other brick. The partial order between lists of bricks $\overline{\mathbb{L}}_1$ and $\overline{\mathbb{L}}_2$ is as follows:

$$\overline{\mathbb{L}}_1 \leq_{\overline{\mathcal{BR}}} \overline{\mathbb{L}}_2 \Leftrightarrow (\overline{\mathbb{L}}_2 = \top_{\overline{\mathcal{BR}}}) \vee (\overline{\mathbb{L}}_1 = \perp_{\overline{\mathcal{BR}}}) \vee (\forall i \in [1, n] : \overline{\mathbb{L}}_1[i] \leq_{\overline{\mathcal{B}}} \overline{\mathbb{L}}_2[i])$$

where we make $\overline{\mathbb{L}}_1$ and $\overline{\mathbb{L}}_2$ have the same size n by adding empty bricks ($\{\emptyset\}^{0,0}$) at the end of the shorter list. The upper bound operator on a single brick is:

$$\sqcup_{\overline{\mathcal{B}}}([\overline{\mathcal{S}}_1]^{m_1, M_1}, [\overline{\mathcal{S}}_2]^{m_2, M_2}) = [\overline{\mathcal{S}}_1 \cup \overline{\mathcal{S}}_2]^{\min(m_1, m_2), \max(M_1, M_2)}$$

The upper bound operator on lists of bricks (elements of our domain) is as follows: given two lists $\overline{\mathbb{L}}_1$ and $\overline{\mathbb{L}}_2$, we make them to have the same size n adding empty bricks to the shorter one. Then: $\sqcup_{\overline{\mathcal{BR}}}(\overline{\mathbb{L}}_1, \overline{\mathbb{L}}_2) = \overline{\mathbb{L}}_R[1] \overline{\mathbb{L}}_R[2] \dots \overline{\mathbb{L}}_R[n]$ where $\forall i \in [1, n] : \overline{\mathbb{L}}_R[i] = \sqcup_{\overline{\mathcal{B}}}(\overline{\mathbb{L}}_1[i], \overline{\mathbb{L}}_2[i])$.

Let k_L , k_I and k_S be three constant integer values. The widening operator $\nabla_{\overline{\mathcal{BR}}} : (\overline{\mathcal{BR}} \times \overline{\mathcal{BR}}) \rightarrow \overline{\mathcal{BR}}$ is defined as follows:

$$\nabla_{\overline{\mathcal{BR}}}(\overline{\mathbb{L}}_1, \overline{\mathbb{L}}_2) = \begin{cases} \top_{\overline{\mathcal{BR}}} & \text{if } (\overline{\mathbb{L}}_1 \not\leq_{\overline{\mathcal{BR}}} \overline{\mathbb{L}}_2 \wedge \overline{\mathbb{L}}_2 \not\leq_{\overline{\mathcal{BR}}} \overline{\mathbb{L}}_1) \vee \\ & (\exists i \in [1, 2] : \text{len}(\overline{\mathbb{L}}_i) > k_L) \\ w(\overline{\mathbb{L}}_1, \overline{\mathbb{L}}_2) & \text{otherwise} \end{cases}$$

where $w(\overline{\mathbb{L}}_1, \overline{\mathbb{L}}_2) = [\overline{\mathcal{B}}_1^{\text{new}}(\overline{\mathbb{L}}_1[1], \overline{\mathbb{L}}_2[1]); \overline{\mathcal{B}}_2^{\text{new}}(\overline{\mathbb{L}}_1[2], \overline{\mathbb{L}}_2[2]); \dots; \overline{\mathcal{B}}_n^{\text{new}}(\overline{\mathbb{L}}_1[n], \overline{\mathbb{L}}_2[n])]$, with n being the size of the bigger list (we make them to have the same size n adding empty bricks to the shorter one), and $\overline{\mathcal{B}}_i^{\text{new}}(\overline{\mathbb{L}}_1[i], \overline{\mathbb{L}}_2[i])$ is defined by:

Table 5. The abstract semantics of $\overline{\mathcal{BR}}$

$$\begin{aligned}
\overline{\mathcal{S}}_{\mathcal{BR}}[\text{new String}(\text{str})](\cdot) &= [\{\text{str}\}]^{1,1} \\
\overline{\mathcal{S}}_{\mathcal{BR}}[\text{concat}](\overline{\mathbf{b}}_1, \overline{\mathbf{b}}_2) &= \text{normBricks}(\text{concatList}(\overline{\mathbf{b}}_1, \overline{\mathbf{b}}_2)) \\
\overline{\mathcal{S}}_{\mathcal{BR}}[\text{readLine}](\cdot) &= \top_{\overline{\mathcal{BR}}} \\
\overline{\mathcal{S}}_{\mathcal{BR}}[\text{substring}_e^e](\overline{\mathbf{b}}) &= \begin{cases} [\overline{\mathbf{T}}']^{1,1} & \text{if } \overline{\mathbf{b}}[0] = [\overline{\mathbf{T}}]^{1,1} \wedge \forall \bar{\mathbf{t}} \in \overline{\mathbf{T}} : \text{len}(\bar{\mathbf{t}}) \geq e \\ \top_{\overline{\mathcal{BR}}} & \text{otherwise} \end{cases} \\
\overline{\mathcal{B}}_{\mathcal{BR}}[\text{contains}_c](\overline{\mathbf{b}}) &= \begin{cases} \text{true} & \text{if } \exists \overline{\mathbf{B}} \in \overline{\mathbf{B}} : \overline{\mathbf{B}} = [\overline{\mathbf{T}}]^{m,M} \wedge 1 \leq m \leq M \wedge (\forall \bar{\mathbf{t}} \in \overline{\mathbf{T}} : c \in \text{char}(\bar{\mathbf{t}})) \\ \text{false} & \text{if } \forall [\overline{\mathbf{T}}]^{m,M} \in \overline{\mathbf{b}}, \forall \bar{\mathbf{t}} \in \overline{\mathbf{T}} : c \notin \text{char}(\bar{\mathbf{t}}) \\ \top_{\mathbf{B}} & \text{otherwise} \end{cases} \\
\overline{\mathcal{B}}_i^{\text{new}}([\overline{\mathcal{S}}_{1i}]^{m_{1i}, M_{1i}}, [\overline{\mathcal{S}}_{2i}]^{m_{2i}, M_{2i}}) &= \begin{cases} \top_{\overline{\mathcal{B}}} & \text{if } |\overline{\mathcal{S}}_{1i} \cup \overline{\mathcal{S}}_{2i}| > k_S \\ & \quad \forall \overline{\mathcal{L}}_1[i] = \top_{\overline{\mathcal{B}}} \vee \overline{\mathcal{L}}_2[i] = \top_{\overline{\mathcal{B}}} \\ [\overline{\mathcal{S}}_{1i} \cup \overline{\mathcal{S}}_{2i}]^{(0, \infty)} & \text{if } (M - m) > k_I \\ [\overline{\mathcal{S}}_{1i} \cup \overline{\mathcal{S}}_{2i}]^{(m, M)} & \text{otherwise} \end{cases}
\end{aligned}$$

where $m = \min(m_{1i}, m_{2i})$ and $M = \max(M_{1i}, M_{2i})$. $\nabla_{\overline{\mathcal{BR}}}$ is an upper bound operator because it returns either $\top_{\overline{\mathcal{BR}}}$ or $w(\overline{\mathcal{L}}_1, \overline{\mathcal{L}}_2)$, which builds a new list of bricks which is bigger (with respect to $\leq_{\overline{\mathcal{BR}}}$) than both $\overline{\mathcal{L}}_1$ and $\overline{\mathcal{L}}_2$. The resulting list is greater or equal because each brick is greater than or equal to the two corresponding bricks in $\overline{\mathcal{L}}_1$ and $\overline{\mathcal{L}}_2$, since we always take the union of the two strings sets and an index range bigger than the initial two. Moreover, this operator converges because a value of an ascending chain can increase along three axes: (i) the length of the brick list, (ii) the indices range of a certain brick, and (iii) the strings contained in a certain brick. The growth of an abstract value is bounded along each axis with the help of the three constants. After the list has reached k_L elements, the entire abstract value is approximated to $\top_{\overline{\mathcal{BR}}}$. If the range of a certain brick becomes larger than k_I , the range is approximated to $(0, +\infty)$. Finally, if the strings set of a certain brick reaches k_S elements, the brick is approximated to $\top_{\overline{\mathcal{B}}}$. The lower bound operator is dual with respect to the upper bound operator above. Formally, $\sqcap_{\overline{\mathcal{B}}}([\overline{\mathcal{S}}_1]^{m_1, M_1}, [\overline{\mathcal{S}}_2]^{m_2, M_2}) = [\overline{\mathcal{S}}_1 \cap \overline{\mathcal{S}}_2]^{\max(m_1, m_2), \min(M_1, M_2)}$. The abstraction function is defined by: $\alpha'_{\overline{\mathcal{BR}}}(\mathbf{s}) = [\{\mathbf{s}\}]^{(1,1)}$ and

$$\alpha_{\overline{\mathcal{BR}}}(S_1) = \bigsqcup_{\overline{\mathcal{BR}}, s \in S_1} \alpha'_{\overline{\mathcal{BR}}}(s) = [S_1]^{(1,1)}$$

Semantics. Table 5 defines the abstract semantics on $\overline{\mathcal{BR}}$. When a constant string is evaluated, the semantics returns a single brick containing exactly that string with $[1, 1]$ as index. For the concatenation of two strings, we rely on the *concatList* function that concatenates two lists of bricks, and then we normalise its result. *readLine* returns the top value, while *substring_e^e* returns the substring iff the first brick of the list has index $[1, 1]$ and the length of all the strings contained in it is greater than e . Notice that $\overline{\mathbf{T}}' = \{\bar{\mathbf{t}}.\text{substring}(\mathbf{b}, e) \mid \bar{\mathbf{t}} \in \overline{\mathbf{T}}\}$. Finally, the semantics of *contains_c* returns *true* iff there is surely at least one brick that contains c and whose minimal index is at least 1. It returns *false* iff all the bricks do not contain c , and $\top_{\mathbf{B}}$ otherwise.

#I	Var	$\overline{\mathcal{BR}}$
1	query	$[\{s_1\}]^{1,1}$
3	l	$\top_{\overline{\mathcal{B}}}$
3	query	$[\{s_1 + s_2\}]^{1,1} \top_{\overline{\mathcal{B}}} [\{s_3\}]^{1,1}$
4	query	$[\{s_1, s_1 + s_2\}]^{1,1} \top_{\overline{\mathcal{B}}} [\{s_3\}]^{0,1}$
5	per	$[\{s_4\}]^{1,1}$
7	query	$[\{s_1, s_1 + s_2\}]^{1,1} \top_{\overline{\mathcal{B}}} [\{s_3\}]^{0,1}$ $[\{s_5 + s_4 + s_6\}]^{1,1}$

(a) First running example

#I	Var	$\overline{\mathcal{BR}}$
1	x	$[\{“a”\}]^{1,1}$
3	x	\top
4	x	\top

(b) Second running example

Fig. 4. The results of $\overline{\mathcal{BR}}$

Running Example. The results of the analysis of the running examples using $\overline{\mathcal{BR}}$ are depicted in Figures 4(a) and 4(b). For the first program, the bricks of the final result on query are four: (i) the first brick represents a string between s_1 and $s_1 + s_2$, (ii) the second brick corresponds to the input l, (iii) the third brick could be the empty string ϵ or s_3 , and (iv) the fourth brick represents the concatenation of s_5 , s_4 , and s_6 . We can see that the precision is higher than in the previous domains, but still not the best we aim to get: amongst the concrete results we have, for example, $s_1 + s_3 + s_5 + s_4 + s_6$, which cannot be computed in any execution of the analysed code. For the second program, the result is unsatisfactory: the use of the widening operator makes us lose all information. At the end of the program, variable x has value \top .

4.4 String Graphs

The last abstract domain we introduce exploits type graphs, a data structure which represents tree automata [15], adapting them to represent sets of strings. A type graph \overline{T} is a triplet $(\overline{N}, \overline{A}_F, \overline{A}_B)$ where $(\overline{N}, \overline{A}_F)$ is a rooted tree whose arcs in \overline{A}_F are called forward arcs, and \overline{A}_B is a restricted class of arcs, backward arcs, superimposed on $(\overline{N}, \overline{A}_F)$. Each node $\overline{n} \in \overline{N}$ of a type graph has a label, denoted by $\overline{lb}(\overline{n})$, indicating the kind of term it describes, and the nodes are divided into three classes: simple, functor and OR nodes. We use the convention that \overline{n}/i denotes the i-th son of node \overline{n} , and the set of sons of a node \overline{n} is then denoted as $\{\overline{n}/1, \dots, \overline{n}/k\}$ with $\overline{k} = \overline{outdegree}(\overline{n})$ where $\overline{outdegree}$ is a function that given a node returns the number of its sons. We define a modified version of type graphs, called string graphs, which represent strings instead of types. String graphs have the same basic structure of type graphs. The following differences distinguish them: (i) simple nodes have labels from the set $\{\max, \perp, \epsilon\} \cup K$; (ii) the only functor we consider is concat (with its obvious meaning of string concatenation). Thus, functor nodes are labelled with concat/k. An example is depicted in Figure 5. The root of the string graph is an OR node with two sons: a simple node (b) and a concat node with two sons of its own. The second son of the concat node is the root (with the use of a backward arc). Such string graph represents the following set of strings: $\{b, ab, aab, aaab, \dots\} = a^*b$.

The abstract domain is: $\overline{\mathcal{SG}} = \overline{\text{NSG}}$, where $\overline{\text{NSG}}$ is the set of all Normal String Graphs. In fact, the type graphs are very suitable for representing a set of terms. However, several distinct type graphs can have the same denotation. The existence of superfluous nodes and arcs makes operations needed during abstract interpretation, such as the \leq -operation, quite complex and inefficient. In order to reduce this variety of type graphs, additional restrictions are imposed (for details see [15]), defining normal type graphs. We added a few other restrictions (specific for string graphs), thus obtaining the definition of normal string graphs. For example, we impose that `concat` nodes are not allowed to have only one son (they should be replaced by the son itself) or that a `concat` node cannot have two successive sons with both label `concat` (they should be merged together). An algorithm of normalisation ($\overline{\text{normStringGraph}}$), encapsulating all those rules, is defined as well.




Fig. 5. An example of string graph

The bottom element $\perp_{\overline{\mathcal{SG}}}$ is a string graph made by one bottom node. The top element is a string graph made by only one node, a `max`-node. To define the partial order of the domain we can exploit an algorithm defined in [15]: $\leq (\overline{n}, \overline{m}, \emptyset)$. The algorithm compares the two nodes in input $(\overline{n}, \overline{m})$. In some cases the procedure is recursively called, for example if \overline{n} and \overline{m} are both `concat` or `OR` nodes. Note that the recursive call adds a new edge $(\{\overline{n}, \overline{m}\})$ to the third input parameter (a set of edges). If, at the next execution of the procedure ($\leq (\overline{n'}, \overline{m'}, \overline{E})$), the edge $\{\overline{n'}, \overline{m}'\}$ is contained in \overline{E} then the procedure immediately returns true. The order is then:

$$\overline{T}_1 \leq_{\overline{\mathcal{SG}}} \overline{T}_2 \Leftrightarrow \overline{T}_1 = \perp_{\overline{\mathcal{SG}}} \vee (\leq (\overline{n}_0, \overline{m}_0, \emptyset) : \overline{n}_0 = \overline{root}(\overline{T}_1) \wedge \overline{m}_0 = \overline{root}(\overline{T}_2))$$

where $\overline{root}(\overline{T})$ is the root element of the tree defined in \overline{T} . The least upper bound between two string graphs \overline{T}_1 and \overline{T}_2 can be computed creating a new string graph \overline{T} whose root is an `OR`-node and whose sons are \overline{T}_1 and \overline{T}_2 . Then we apply the compaction algorithm that will transform \overline{T} in a normal string graph:

$$\bigsqcup_{\overline{\mathcal{SG}}}(\overline{T}_1, \overline{T}_2) = \overline{\text{normStringGraph}}(\text{OR}(\overline{T}_1, \overline{T}_2))$$

The greatest lower bound operator is described in the appendix of [15], while the widening operator is described in [23]. The abstraction of a string is: $\alpha'_{\overline{\mathcal{SG}}}(\mathbf{s}) = \text{concat}/k\{\mathbf{s}[i] : i \in [0, k - 1]\}$ where $k = \text{len}(\mathbf{s})$, and the abstraction function is:

$$\alpha_{\overline{\mathcal{SG}}}(\mathbf{S}_1) = \bigsqcup_{\overline{\mathcal{SG}}, \mathbf{s} \in \mathbf{S}_1} \alpha'_{\overline{\mathcal{SG}}}(\mathbf{s}) = \overline{\text{normStringGraph}}(\text{OR}\{\alpha'_{\overline{\mathcal{SG}}}(\mathbf{s}) : \mathbf{s} \in \mathbf{S}_1\})$$

Semantics. Table 6 defines the abstract semantics on $\overline{\mathcal{SG}}$. The evaluation of a string returns a `concat` containing the sequence of all the characters of the string. When we concatenate two strings, the semantics returns the normalisation of a `concat` node containing the two strings in sequence. As usual, the semantics of `readLine` returns the top value. The semantics of `substringbe` (where $\overline{\text{res}} = \text{concat}/(e - b)\{\overline{root}(\overline{T})/i : i \in [b, e - 1]\}$) returns a precise value only if the root is a `concat` node with at least `e` characters. Finally, `containsc` returns `true` iff there is a `concat` node containing `c` in the tree, and without any `OR` node in the path from the root to this node.

Table 6. The abstract semantics of \overline{SG}

$$\begin{aligned}
 \overline{SG}[\text{new String}(\text{str})]() &= \text{concat}/k\{\text{str}[i] : i \in [0, k - 1]\} \\
 \overline{SG}[\text{concat}](\bar{t}_1, \bar{t}_2) &= \text{normStringGraph}(\text{concat}/2\{\bar{t}_1, \bar{t}_2\}) \\
 \overline{SG}[\text{readLine}]() &= \top_{\overline{SG}} \\
 \overline{SG}[\text{substring}_e^s](\bar{t}) &= \begin{cases} \overline{\text{res}} & \text{if } \overline{\text{root}}(\bar{t}) = \text{concat}/k \wedge \forall i \in [0, e - 1] : \overline{\text{lb}}(\overline{\text{root}}(\bar{t})/i) \in K \\ \top_{\overline{SG}} & \text{otherwise} \end{cases} \\
 \overline{SG}[\text{contains}_c](\bar{t}) &= \begin{cases} \text{true} & \text{if } \exists \bar{m} \in \bar{t} : \bar{m} = \text{concat}/k \wedge \text{OR} \notin \overline{\text{path}}(\overline{\text{root}}, \bar{m}) \wedge \\ & \exists i : \overline{\text{lb}}(\bar{m}/i) = c \\ \text{false} & \text{if } \nexists \bar{n} \in \bar{t} : \overline{\text{lb}}(\bar{n}) = \max \vee \overline{\text{lb}}(\bar{n}) = c \\ \top_B & \text{otherwise} \end{cases}
 \end{aligned}$$

#I	Var	\overline{SG}
1	query	concat[s ₁]
3	l	max
3	query	concat[s ₁ + s ₂ ; max; s ₃]
4	query	SG ₁ = OR[concat[s ₁]; concat[s ₁ + s ₂ ; max; s ₃]]
5	per	concat[s ₄]
7	query	concat[SG ₁ ; concat[s ₅ + s ₄ + s ₆]]

(a) First running example

#I	Var	\overline{SG}
1	x	concat["a"]
3	x	OR ₁ ["a"; concat["0"; OR ₁ ; "1"]]
4	x	OR ₁ ["a"; concat["0"; OR ₁ ; "1"]]

(b) Second running example

Fig. 6. The results of \overline{SG}

Running Example. The results of the analysis of the running examples through string graphs are depicted in Figures 6(a) and 6(b). For sake of simplicity, we adopt the notation $\text{concat}[s]$ to indicate a string graph with a concat node whose sons are all the characters of string s . The symbol $+$ represents, as usual, string concatenation, while $;$ is used to separate different sons of a node.

For the first program, the resulting string graph for **query** represents exactly the two possible outcomes of the procedure. For the second program, the resulting string graph for **x** represents exactly all the concrete possible values of **x**. Note that the resulting string graph contains a backward arc to allow the repetition of the pattern $0^n \dots 1^n$. This abstract domain is the most precise domain for the analysis of both running examples: it tracks information similarly to \overline{BR} domain, but its lub and widening operators are definitely more accurate.

4.5 Discussion: Relations between the Four Domains

The abstract domains we introduced in the previous sections track different types of information. Let us discuss the relations between different domains. Intuitively, there are two axes on which the analyses of string values can work: the characters contained in a string, and their position inside the string. It is easy to see that the \overline{CI} , \overline{PR} and \overline{SU} are less precise than \overline{BR} and \overline{SG} . In fact, \overline{CI} domain considers only character inclusion and completely disregards the order. \overline{PR} and

\overline{SU} domains consider also the order, but limiting themselves to the initial/final segment of the string, and in the same way they collect only partial information about character inclusion. \overline{BR} and \overline{SG} , instead, track both inclusion and order along the string. In [3] we studied these relationships in details: we defined pairs of functions (abstraction and concretization) from domain to domain, and showed that \overline{CI} , \overline{PR} and \overline{SU} are more abstract (i.e., less precise) than both \overline{BR} and \overline{SG} . In the case of \overline{BR} versus \overline{SG} , the comparison is more complex, since they exploit very different data structures. For example, \overline{SG} has OR-nodes, while \overline{BR} can only trace alternatives inside bricks but not outside (like: “these three bricks *or* these other two”). From this perspective, \overline{SG} is more precise than \overline{BR} . Another important difference is that \overline{SG} has backward arcs which allow repetitions of patterns, but they can be traversed how many times we want (even infinite times). With \overline{BR} , instead, we can indicate exactly how many times a certain pattern should be repeated (through the range of bricks). This makes \overline{BR} more expressive than \overline{SG} in that respect. So, these domains are not directly comparable. We obtain the lattice depicted in Figure 7, where the upper domains are more approximated. We denote by \top the abstract domain that does not track any information about string values, and by $\wp(K^*)$ the (naïve and uncomputable) domain that tracks all the possible strings values we can have.

In conclusion, the first three domains (\overline{CI} , \overline{PR} , \overline{SU}) are not so precise but the complexity is kept linear, whereas the other domains (\overline{BR} and \overline{SG}) are more demanding (though in the practice complexity is still kept polynomial) but also more precise.




Fig. 7. The hierarchy of abstract domains

5 Related Work

The static analysis of strings was addressed in various directions.

Kim and Choe [16] introduced recently an approach based on abstract interpretation. They abstract strings with pushdown automata (PDA). The result of the analysis is compared with a grammar to determine if all the strings generated by the PDA belong to the grammar. This approach has a fixed precision, and in the worst case (not often encountered in practice) it has exponential complexity.

Hosoya and Pierce [14] used tree automata to verify dynamically generated XML documents. The regular expression types of this approach recall our \overline{BR} domain, while the tree automata recall our \overline{SG} domain. However, they are focused on building XML documents, while our focus is on collecting possible values of generic string variables. In addition, they require to manually annotate the code through types while our approach is completely automatic.

A more recent work was developed by Yu *et al.* [24]. It presented an automata-based approach for the verification of string operations in PHP programs. The information tracked by this analysis is fixed, and it is specific for PHP programs.

Tabuchi *et al.* [21] presented a type system based on regular expressions. It is focused on a λ -calculus supporting concatenation, and pattern matching. Some type annotation is required when dealing with recursive function.

Thiemann [22] introduced a type system for string analysis based on context-free grammars. Their analysis is more precise than those based on regular expressions, but the only supported string operator is concatenation, and the analysis is tuned at a fixed level of precision.

Context-free grammars are also the basis of the analysis of Christensen *et al.* [1]. This analysis is tuned at a fixed level of abstraction. In the second running example of this paper, \overline{SG} domain reaches a better precision than theirs.

Minamide [18] presented an analysis to statically check some properties of Web pages generated dynamically by a server-side program. This work is specific for HTML pages, while we do not need to know the reference grammar *a priori*. Also in this case, \overline{SG} obtain a better precision on the loop example.

Doh *et al.* [8] proposed a technique called “abstract parsing”: it combines LR(k)-parsing technology and data-flow analysis to analyse dynamically generated documents. Their technique is quite precise, but the level of abstraction is fixed, and it cannot be tuned at different levels of precision and efficiency.

Given this context, our work is the first one that (i) is a generic, flexible, and extensible approach to the analysis of string values, and (ii) can be tuned at different levels of precision and efficiency.

6 Conclusion and Future Work

In this paper we introduced a new framework for the static analysis of string values, and four different abstract domains. We chose some string operators on which we focused our approach defining the concrete and the abstract semantics.

Future Work. We are working on the implementation of our approach in **Sample** (Static Analyzer of Multiple Programming Languages) [9]. We plan to apply our analysis to some case studies to study the precision of our analysis. In order to check the scalability and performances of our approach, we plan to apply our analysis to some Scala standard libraries. Some preliminary experimental results point out that \overline{CI} and $\overline{PR} \times \overline{SU}$ are quite efficient, \overline{BR} is slower but still fast, while \overline{SG} 's performances seem to be still critical.

Acknowledgments. Work partially supported by RAS project “TESLA - Tecniche di enforcement per la sicurezza dei linguaggi e delle applicazioni”, and by SNF project “Verification-Driven Inference of Contracts”.

References

1. Christensen, A., Moller, A., Schwartzbach, M.: Precise analysis of string expressions. In: Cousot, R. (ed.) SAS 2003. LNCS, vol. 2694, pp. 1–18. Springer, Heidelberg (2003)
2. Cortesi, A., Zanioli, M.: Widening and narrowing operators for abstract interpretation. *Computer Languages, Systems and Structures* 37(1), 24–42 (2011)

3. Costantini, G.: Abstract domains for static analysis of strings. Master's thesis, Ca' Foscari University of Venice (2010)
4. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL 1977. ACM, New York (1977)
5. Cousot, P., Cousot, R.: Systematic design of program analysis frameworks. In: POPL 1979. ACM, New York (1979)
6. Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: The ASTREÉ analyzer. In: Sagiv, M. (ed.) ESOP 2005. LNCS, vol. 3444, pp. 21–30. Springer, Heidelberg (2005)
7. Cousot, P., Halbwachs, N.: Automatic discovery of linear restraints among variables of a program. In: Proceedings of POPL 1978. ACM Press, New York (1978)
8. Doh, K., Kim, H., Schmidt, D.: Abstract parsing: Static analysis of dynamically generated string output using LR-parsing technology. In: Palsberg, J., Su, Z. (eds.) SAS 2009. LNCS, vol. 5673, pp. 256–272. Springer, Heidelberg (2009)
9. Ferrara, P.: Static type analysis of pattern matching by abstract interpretation. In: Hatcliff, J., Zucca, E. (eds.) FMOODS 2010. LNCS, vol. 6117, pp. 186–200. Springer, Heidelberg (2010)
10. Gould, C., Su, Z., Devanbu, P.: Static checking of dynamically generated queries in database applications. In: Proceedings of ICSE 2004, pp. 645–654. IEEE Computer Society, Los Alamitos (2004)
11. Granger, P.: Static analysis of linear congruence equalities among variables of a program. In: Abramsky, S. (ed.) CAAP 1991 and TAPSOFT 1991. LNCS, vol. 493, pp. 169–192. Springer, Heidelberg (1991)
12. Gulwani, S.: Automating string processing in spreadsheets using input-output examples. In: Proceedings of POPL 2011. ACM, New York (2011)
13. Hooimeijer, P., Veanes, M.: An evaluation of automata algorithms for string analysis. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 248–262. Springer, Heidelberg (2011)
14. Hosoya, H., Pierce, B.: Xduce: A statically typed xml processing language. ACM Trans. Internet Technol. 3(2), 117–148 (2003)
15. Janssens, G., Bruynooghe, M.: Deriving description of possible values of program variables by means of abstract interpretation. Journal of Logic Programming 13(2-3), 205–258 (1992)
16. Kim, S.-W., Choe, K.-M.: String analysis as an abstract interpretation. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 294–308. Springer, Heidelberg (2011)
17. Logozzo, F., Fähndrich, M.: Pentagons: A weakly relational domain for the efficient validation of array accesses. In: Proceedings of SAC 2008. ACM Press, New York (2008)
18. Minamide, Y.: Static approximation of dynamically generated web pages. In: Proceedings of WWW 2005, pp. 432–441. ACM, New York (2005)
19. Miné, A.: The octagon abstract domain. Higher-Order and Symbolic Computation (2006)
20. Halder, R., Cortesi, A.: Obfuscation-based analysis of sql injection attacks. In: IEEE (ed.) Proceedings of ISCC 2010 (2010)
21. Tabuchi, N., Sumii, E., Yonezawa, A.: Regular expression types for strings in a text processing language. Electr. Notes Theor. Comput. Sci. 75 (2002)

22. Thiemann, P.: Grammar-based analysis of string expressions. In: Proceedings of TLDI 2005, pp. 59–70. ACM, New York (2005)
23. van Hentenryck, P., Cortesi, A., Le Charlier, B.: Type analysis of prolog using type graphs. *Journal of Logic Programming* 22(3), 179–208 (1995)
24. Yu, F., Bultan, T., Cova, M., Ibarra, O.: Symbolic string verification: An automata-based approach. In: Havelund, K., Majumdar, R. (eds.) SPIN 2008. LNCS, vol. 5156, pp. 306–324. Springer, Heidelberg (2008)