# High-Level Abstractions for Simplifying Extended String Constraints in SMT 

Andrew Reynolds ${ }^{1}$, Andres Nötzli ${ }^{2}$, Clark Barrett ${ }^{2}$, and Cesare Tinelli ${ }^{1}$<br>${ }^{1}$ Department of Computer Science, The University of Iowa<br>${ }^{2}$ Department of Computer Science, Stanford University


#### Abstract

Satisfiability Modulo Theories (SMT) solvers with support for the theory of strings have recently emerged as powerful tools for reasoning about string-manipulating programs. However, due to the complex semantics of extended string functions, it is challenging to develop scalable solvers for the string constraints produced by program analysis tools. We identify several classes of simplification techniques that are critical for the efficient processing of string constraints in SMT solvers. These techniques can reduce the size and complexity of input constraints by reasoning about arithmetic entailment, multisets, and string containment relationships over input terms. We provide experimental evidence that implementing them results in significant improvements over the performance of state-of-the-art SMT solvers for extended string constraints.


## 1 Introduction

Most programming languages support strings natively and a considerable number of programs perform some form of string manipulation. Automated reasoning about stringmanipulating programs for verification and test case generation purposes is then highly relevant for these languages and programs. Applications to security, such as finding SQL injection and XSS vulnerabilities in web applications [16, 14, 21] or proving their absence, are of critical importance. String constraints have also been used to generate relational database tables from SQL queries for unit testing purposes [19]. These applications require modeling all of the string operations that appear in real programs. This is challenging since some of those operations are complex and often realized by iterative applications of simpler operations. Additionally, since strings in many programming languages have variable length, reasoning accurately about them cannot be done by a reduction to bounded types such as bit-vectors, and requires instead the development of solvers for unbounded strings. To make this type of reasoning more scalable, the use of dedicated theory solvers natively supporting common string operations has been proposed [5, 8]. Some string solvers are fully integrated within Satisfiability Modulo Theories (SMT) solvers [11, 4]; some are built (externally) on top of such solvers [17, 8, 14]; and others are independent of SMT solvers [21].

A major challenge in developing solvers for unbounded string constraints is the complex semantics of extended string functions beyond the basic operations of string concatenation and equality. Extended functions include replace, which replaces a string
in another string, and indexof, which returns the position of a string in another string. Another challenge is that constraints using extended functions are often combined with constraints over other theories, e.g. integer constraints over string lengths or applications of indexof, which requires the involvement of solvers for those theories. Current string solvers address these challenges by reducing constraints with extended string functions to typically more verbose constraints over basic functions. As with every reduction, some of the higher level structure of the problem may be lost, with negative repercussions on the performance and scalability.

To address this issue, we have developed new techniques that reason about constraints with extended string operators before they are reduced to simpler ones. This analysis of complex terms can often eliminate the need for expensive reductions. The techniques are based on reasoning about relationships over strings with high-level abstractions, such as their arithmetic relationships (e.g., reasoning about their length), their string containment relationships, and their relationships as multisets of characters. We have implemented these techniques in CVC4, an SMT solver with native support for string reasoning. An experimental evaluation with benchmarks from various applications shows that our new techniques allows CVC4 to significantly outperform other state-of-the-art solvers that target extended string constraints.

Our main contributions are:

- A novel procedure for proving entailments over arithmetic predicates built from the theory of strings and linear integer arithmetic.
- Extensions of this technique for showing containment relationships between strings.
- A novel simplification technique based on abstracting strings as multisets.
- Experimental evidence that the simplification techniques provide significant performance improvements over current state-of-the-art solvers.

In the remainder of this section, we discuss related work. In Section 2, we provide some background on the theory of strings and how solvers reduce extended functions. In Sections 3, 4 and 5, we describe, respectively, our arithmetic-based, containment-based, and multiset-based simplification techniques. Section 6 describes our implementation of those techniques, and Section 7 presents our evaluation.

Related Work Various approaches to solving constraints over extended string functions have been proposed. Saxena et al. [14] showed that constraints from the symbolic execution of JavaScript code contain a significant number of extended string functions, which underlines their importance. Their approach translates string constraints to bit-vector constraints, similar to other approaches based on bounded strings such as HAMPI [8]. Bjørner et al. [5] proposed native support for extended string operators in string solvers for scaling symbolic execution of .NET code. They reduce extended string functions to basic ones after getting bounds for string lengths from an integer solver. They also showed that constraints involving unbounded strings and replace are undecidable. PASS [10] reduces string constraints over extended functions to arrays. Z3-str and its successors [23, 22, 4] reduce extended string functions to basic functions eagerly during preprocessing. S 3 [16] reduces recursive functions such as replace incrementally by splitting and unfolding. Its successor S3P [17] refines this reduction by pruning the resulting subproblems for better performance. CVC4 [3] reduces constraints with extended functions lazily and leverages context-dependent simplifications to simplify

$$
\begin{array}{lll}
n: \text { Int for all } n \in \mathbb{N} & +: \operatorname{Int} \times \operatorname{Int} \rightarrow \operatorname{Int} & -: \operatorname{Int} \rightarrow \operatorname{Int} \quad \geqslant: \operatorname{Int} \times \operatorname{Int} \rightarrow \text { Bool } \\
l: \text { Str for } \operatorname{lll} l \in \mathcal{A}^{*} & \text { con }: \operatorname{Str} \times \cdots \times \operatorname{Str} \rightarrow \operatorname{Str} & \text { len:Str } \rightarrow \operatorname{Int} \\
\text { substr: } & \operatorname{Str} \times \operatorname{Int} \times \operatorname{Int} \rightarrow \operatorname{Str} & \text { contains: Str } \times \operatorname{Str} \rightarrow \text { Bool } \\
\text { indexof: } & \operatorname{Str} \times \operatorname{Str} \times \operatorname{Int} \rightarrow \operatorname{Int} & \text { replace: } \operatorname{Str} \times \operatorname{Str} \times \operatorname{Str} \rightarrow \operatorname{Str} \\
\text { str.to.int: } \operatorname{Str} \rightarrow \operatorname{Int} & \text { int.to.str: } \operatorname{Int} \rightarrow \operatorname{Str}
\end{array}
$$

Fig. 1. Functions in signature $\Sigma_{\mathrm{s}}$. Str and Int denote strings and integers respectively.
the reductions [13]. TRAU [1] reduces certain extended functions, such as replace, to context-free membership constraints. OSTRICH [6] implements a decision procedure for a subset of constraints that include extended string functions. The simplification techniques presented in this paper are agnostic to the underlying solving procedure, so they can be combined with all of these approaches.

## 2 Preliminaries

We work in the context of many-sorted first-order logic with equality and assume the reader is familiar with the notions of signature, term, literal, formula, and formal interpretation of formulas. We review a few relevant definitions in the following. A theory is a pair $T=(\Sigma, \mathbf{I})$ where $\Sigma$ is a signature and $\mathbf{I}$ is a class of $\Sigma$-interpretations, the models of $T$. We assume $\Sigma$ contains the equality predicate $\approx$, interpreted as the identity relation, and the predicates $\top$ (for true) and $\perp$ (for false). A $\Sigma$-formula $\varphi$ is satisfiable (resp., unsatisfiable) in $T$ if it is satisfied by some (resp., no) interpretation in I. We write $\models_{T} \varphi$ to denote that the $\Sigma$-formula $\varphi$ is $T$-valid, i.e., is satisfied in every model of $T$. Two $\Sigma$-terms $t_{1}$ and $t_{2}$ are equivalent in $T$ if $\models_{T} t_{1} \approx t_{2}$.

We consider an extended theory $T_{\mathrm{S}}$ of strings and length equations, whose signature $\Sigma_{\mathrm{S}}$ is given in Figure 1 and whose models differ only on how they interpret variables. ${ }^{3}$ We assume a fixed finite alphabet $\mathcal{A}$ of characters which includes the digits $\{0, \ldots, 9\}$. The signature includes the sorts Bool, Int, and Str denoting the Booleans, the integers $(\mathbb{Z})$, and Kleene closure of $\mathcal{A}\left(\mathcal{A}^{*}\right)$, respectively. The top half of Figure 1 includes the usual symbols of linear integer arithmetic, interpreted as expected, a string literal $l$ for each word/string of $\mathcal{A}^{*}$, a variadic function symbol con, interpreted as word concatenation, and a function symbol len, interpreted as the word length function. We write $\epsilon$ for the empty word and abbreviate len $(s)$ as $|s|$. We use words over the characters $\mathrm{a}, \mathrm{b}$, and c , as in abca, as concrete examples of string literals.

We refer to the function symbols in the bottom half of the figure as extended functions and refer to terms containing them as extended terms. A position in a string $l \in \mathcal{A}^{*}$ is a non-negative integer $n$ smaller than the length of $l$ that identifies the $(n+1)^{t h}$ character of $l$-with 0 identifying the first character, 1 the second, and so on. For all models $\mathcal{I}$ of $T_{\mathrm{S}}$, all $l, l_{1}, l_{2} \in \mathcal{A}^{*}$, and $n, m \in \mathbb{Z}$, $\operatorname{substr}^{\mathcal{I}}(l, n, m)$ (the interpretation of substr in $\mathcal{I}$ applied to $l, n, m)$ is the longest substring of $l$ starting at position $n$ with length at most

[^0]$m$, or $\epsilon$ if $n$ is an invalid position or $m$ is not positive; contains ${ }^{\mathcal{I}}\left(l_{1}, l_{2}\right)$ is true if and only if $l_{2}$ is a substring of $l_{1}$, with $\epsilon$ being a substring of every string; indexof ${ }^{\mathcal{I}}\left(l_{1}, l_{2}, n\right)$ is the position of the first occurrence of $l_{2}$ in $l_{1}$ at or after position $n, n$ if $l_{2}$ is empty and $0 \leqslant n \leqslant\left|l_{1}\right|$, and -1 if $n$ is an invalid position, or if no such occurrence exists; replace ${ }^{\mathcal{I}}\left(l, l_{1}, l_{2}\right)$ is the result of replacing the first occurrence of $l_{1}$ in $l$ by $l_{2}, l$ if $l$ does not contain $l_{1}$, or the result of prepending $l_{2}$ to $l$ if $l_{1}$ is empty; str.to.int ${ }^{\mathcal{I}}(l)$ is the non-negative integer represented by $l$ in decimal notation or -1 if the string contains non-digit characters; int.to.str ${ }^{\mathcal{I}}(n)$ is the result of converting $n$ to the corresponding string in decimal notation if $n$ is non-negative, or $\epsilon$ otherwise. We write substr$(t, u)$ as shorthand for the term $\operatorname{substr}(t, u,|t|)$, i.e. the suffix of $t$ starting at position $u$.

Note that the semantics for replace and indexof correspond to the semantics in the current draft of the SMT-LIB standard for the theory of strings [15]; they are slightly different from the ones described in previous work [4, 13, 18].

### 2.1 Solving Extended String Constraints (with Simplification)

Various efficient solvers have been designed for the satisfiability problem for quantifierfree $T_{\mathrm{S}}$-constraints, including CVC4 [3], S3\# [18] and Z3STR3 [4]. In this section, we give an overview of how these solvers process extended functions in practice.

Generally speaking, constraints involving extended functions are converted to basic ones through a series of reductions performed in an incremental fashion by the solver. Operators whose reduction requires universal quantification are dealt with by guessing upper bounds on the lengths of input strings or by lazily adding constraints that block models that do not satisfy extended string constraints.

Example 1. To determine the satisfiability of $\neg \operatorname{contains}(t, s)$, the application of contains is reduced to constraints that ensure that $s$ is not a substring of $t$ at any position. Assuming we have a fixed upper bound $n$ on the length of $t$, the above constraint is equivalent to the finite conjunction substr$(t, 0,|s|) \not \approx s \wedge \cdots \wedge \operatorname{substr}(t, n,|s|) \not \approx s$. Each application of substr is then eliminated by introducing an equality that constrains a fresh variable $x_{i}$ to have the semantics of that substring. Thus, reducing the formula above results in

$$
\bigwedge_{i=0}^{n}|t| \geqslant i+|s| \Rightarrow\left(x_{i} \not \approx s \wedge t \approx \operatorname{con}\left(x_{i}^{p r e}, x_{i}, x_{i}^{p o s t}\right) \wedge\left|x_{i}^{p r e}\right| \approx i \wedge\left|x_{i}\right| \approx|s|\right)
$$

where $x_{i}, x_{i}^{p r e}, x_{i}^{p o s t}$ are fresh string variables. ${ }^{4}$ The above conjunction involves only string concatenation, string length, and equality, and thus can be handled by a string solver with support for word equations with length constraints.

The reduction in Example 1 introduces $5 \cdot n$ theory literals over basic string functions and $3 \cdot n$ string variables. A full reduction accounting for all corner cases of substr is even more complex and thus more expensive to process, even for small values of $n$. These performance challenges can be addressed by aggressive simplifications that eliminate extended functions using high-level reasoning, as shown in the next example.

[^1]Example 2. Consider an instance of the previous example where $s=\operatorname{con}(\mathrm{a}, x)$ and $t=\operatorname{con}(\mathrm{b}, \operatorname{substr}(x, 0, n))$. A full reduction of $\neg \operatorname{contains}(t, s)$ that eliminates all applications of substr, including those in $t$, introduces $5 \cdot n+5$ new theory literals and $3 \cdot n+3$ string variables. However, based on the semantics of contains it is easy to see that $\neg$ contains $(t, s)$ is $T_{\mathrm{S}}$-valid: if $t$ were to contain $s$, then $s$ would have to occur in the portion of $t$ after its first character b , since the first character of $s$ is a. However, con $(\mathrm{a}, x)$ cannot be contained in substr$(x, 0, n)$, since the length of the former is at least $|x|+1$, while the length of the latter is at most $|x|$. A solver which recognizes that $\neg$ contains $(t, s)$ can be simplified to $T$ in this case can avoid the reduction altogether.

We advocate for aggressive simplification techniques to improve the performance of string solvers for extended functions. In the next sections, we describe several classes of such techniques that can be applied to inputs as a preprocessing step or during solving as part of a context-dependent solving strategy [13]. We present them as sets $R$ of rewrite rules of the form $t \rightarrow_{R} s$, where $s$ is a (simplified) term equivalent to $t$ in $T_{\mathrm{s}}$. We assume a deterministic application strategy for these rules, such that each term $t$ rewrites to a unique simplified form, denoted by $t \downarrow$, which is irreducible by the rules. We split our simplifications into four categories, presented in Figures 4, 6, 7 and 8. ${ }^{5}$

## 3 Arithmetic-Based String Simplification

To simplify string terms, it is useful to establish relationships between quantities such as the lengths of strings. For example, contains $(t, s)$ can be simplified to $\perp$ for a particular $s$ and $t$ if it can be inferred that $|s|$ is strictly greater than $|t|$. This section defines an inference system for such arithmetic relationships and the simplifications that it enables.

We are interested in proving the $T_{\mathrm{S}}$-validity of formulas of the form $u \geqslant 0$, where $u$ is a $\Sigma_{\mathrm{S}}$-term of integer type. We describe an inference system as a set of rules for deriving judgments of the form $\vdash u \geqslant 0$ and a specific rule application strategy we have implemented. The inference system is sound in the sense that $\models_{T_{\mathrm{S}}} u \geqslant 0$ whenever $\vdash u \geqslant 0$ is derivable in it. It is, however, incomplete as it may fail to derive $\vdash u \geqslant 0$ in some cases when $\models_{T_{\mathrm{S}}} u \geqslant 0$. This incompleteness is by design, since proving the $T_{\mathrm{S}}$-validity of inequalities is generally expensive due to the NP-hardness of linear integer arithmetic. Without loss of generality, we require that the term $u$ be in a simplified form, where terms of the form $|l|$ with $l$ a string literal of $n$ characters are rewritten to $n$, terms of the form $\left|\operatorname{con}\left(t_{1}, \ldots, t_{n}\right)\right|$ are rewritten to $\left|t_{1}\right|+\cdots+\left|t_{n}\right|$, and like monomials in arithmetic terms are combined in the usual way (e.g., $2 \cdot|x|+|x|$ is rewritten to $3 \cdot|x|$ ).

Definition 1 (Polynomial Form). An arithmetic term $u$ is in polynomial form if $u=$ $m_{1} \cdot u_{1}+\ldots m_{n} \cdot u_{n}+m$, where $m_{1}, \ldots, m_{n}$ are non-zero integer constants, $m$ is an integer constant, and each $u_{1}, \ldots, u_{n}$ is a unique term and one of the following:

1. an integer variable,
2. an application of length to a string variable, e.g. $|x|$,
3. an application of length to an extended function, e.g. $|\operatorname{substr}(t, v, w)|$, or
4. an application of an extended function of integer type, e.g. indexof $(t, s, v)$.

[^2]\[

$$
\begin{aligned}
& \frac{u \rightarrow_{U}^{*} n \quad n \geqslant 0}{\vdash u \geqslant 0} \quad \text { where } \\
& |t| \rightarrow_{U} \quad 0 \\
& |\operatorname{substr}(t, v, w)| \rightarrow_{U} \begin{cases}w & \text { if } \vdash v \geqslant 0 \text { and } \vdash|t| \geqslant v+w \\
|t|-v & \text { if } \vdash v \geqslant 0 \text { and } \vdash v+w \geqslant|t|\end{cases} \\
& \mid \text { replace }(t, s, r) \left\lvert\, \rightarrow_{U}\left\{\begin{array}{l}
|t| \\
|t|-|s|
\end{array} \quad \text { if } \vdash|r| \geqslant|s| \text { or } \vdash|r| \geqslant|t|\right.\right. \\
& \mid \text { int.to. } \operatorname{str}(v) \mid \rightarrow_{U} \quad 1 \quad \text { if } \vdash v \geqslant 0 \\
& \text { indexof }(t, s, v) \rightarrow_{U} \quad-1 \\
& \text { str.to.int }(t) \rightarrow_{U} \quad-1 \\
& m \cdot v+u^{\prime} \rightarrow_{U} \quad m \cdot w+u^{\prime} \text { if } v \rightarrow_{U} w \text { and } m>0 \text { or } v \rightarrow O w \text { and } m<0 \\
& |\operatorname{substr}(t, v, w)| \rightarrow 0 \quad w \quad \text { if } \vdash w \geqslant 0 \\
& |\operatorname{substr}(t, v, w)| \rightarrow O\left\{\begin{array}{l}
|t|-v \quad \text { if } \vdash|t| \geqslant v \\
|t|
\end{array}\right. \\
& \mid \text { replace }(t, s, r) \left\lvert\, \rightarrow O\left\{\begin{array}{l}
|t| \\
|t|+|r|
\end{array} \quad \text { if } \vdash|s| \geqslant|r|\right.\right. \\
& \text { |int.to. } \operatorname{str}(v) \left\lvert\, \rightarrow_{O} \begin{cases}v & \text { if } \vdash v>0 \\
v+1 & \text { if } \vdash v \geqslant 0\end{cases} \right. \\
& \text { indexof }(t, s, v) \rightarrow O\left\{\begin{array}{l}
|t|-|s| \quad \text { if } \vdash|t| \geqslant|s| \\
|t|
\end{array}\right. \\
& m \cdot v+u^{\prime} \rightarrow_{O} \quad m \cdot w+u^{\prime} \text { if } v \rightarrow_{O} w \text { and } m>0 \text { or } v \rightarrow_{U} w \text { and } m<0
\end{aligned}
$$
\]

Fig. 2. Rules for arithmetic entailment based on under- and over-approximations computed for arithmetic terms containing extended string operators. We write $t, s, r$ to denote string terms, $u, u^{\prime}, v, w$ to denote integer terms and $m, n$ to denote integer constants.

Given $u$ in polynomial form, our inference system uses a set of over- and under-approximations for showing that $u \geqslant 0$ holds in all models of $T_{\mathrm{s}}$. We define two auxiliary rewrite systems, denoted $\rightarrow_{O}$ and $\rightarrow_{U}$. If $u$ rewrites to $v$ (in zero or more steps) in $\rightarrow_{O}$, written $u \rightarrow{ }_{O}^{*} v$, we say that $v$ is an over-approximation of $u$. We can prove in that case that $\models_{T_{\mathrm{s}}} v \geqslant u$. Dually, if $u$ rewrites to $v$ in $\rightarrow_{U}$, written $u \rightarrow_{U}^{*} v$, we say that $v$ is an under-approximation of $u$ and can prove that $\models_{T_{\mathrm{S}}} u \geqslant v$. Based on these definitions, the core of our inference system can be summarized by the single inference rule schema provided in Figure 2 together with the conditional rewrite systems $\rightarrow_{O}$ and $\rightarrow_{U}$ which are defined inductively in terms of the inference system and each other.

A majority of the rewrite rules have side conditions requiring the derivability of certain judgments in the same inference system. To improve their readability we take some liberties with the notation and write $\vdash u_{1} \geqslant u_{2}$, say, instead of $\vdash u_{1}-u_{2} \geqslant 0$. For example, $|\operatorname{substr}(t, v, w)|$ is under-approximated by $w$ if it can be inferred that the interval from $v$ to $v+w$ is a valid range of positions in string $t$, which is expressed by the side conditions $\vdash v \geqslant 0$ and $\vdash|t| \geqslant v+w$. Note that some arithmetic terms, such as $|\operatorname{substr}(t, v, w)|$, can be approximated in multiple ways-hence the need for a strategy for choosing the best approximation for arithmetic string terms, described later. The rules for polynomials are written modulo associativity of + and state that a monomial $m \cdot v$ in them can be over- or under-approximated based on the sign of the coefficient $m$.

```
Str-Arith-Approx \((u)\), where \(u=u_{x}+u_{\ell}+u_{s}+m\) and:
    \(-u_{x}=m_{1}^{y} \cdot y_{1}+\ldots+m_{n}^{y} \cdot y_{n}\),
    - \(u_{\ell}=m_{1}^{\ell} \cdot\left|x_{1}\right|+\ldots+m_{p}^{\ell} \cdot\left|x_{p}\right|\),
    - \(u_{s}=m_{1}^{v} \cdot v_{1}+\ldots+m_{q}^{v} \cdot v_{q}\).
for variables \(x_{1}, \ldots, x_{p}, y_{1}, \ldots, y_{n}\) and extended terms \(v_{1}, \ldots v_{q}\) :
    1. If \(q>0\), choose a \(v_{i}\) and \(v_{i}^{a}\) that maximize the following criteria (in descending order),
        where \(u^{\prime}=\left(u\left[m_{i}^{v} \cdot v_{i} \mapsto m_{i}^{v} \cdot v_{i}^{a}\right]\right) \downarrow\) :
            (a) (Soundness) \(v_{i} \rightarrow_{U} v_{i}^{a}\) if \(m_{i}^{v}>0\) and \(v_{i} \rightarrow_{O} v_{i}^{a}\) if \(m_{i}^{v}<0\);
            (b) (Avoids new terms) Minimizes the size of negcoeff \(\left(u^{\prime}\right) \backslash\) negcoeff \((u)\);
            (c) (Cancels existing terms) Maximizes the size of negcoeff \((u) \backslash \operatorname{negcoeff}\left(u^{\prime}\right)\).
        Return \(u \rightarrow_{U} u^{\prime}\).
    2. If \(p>0\) and \(m_{j}^{\ell}>0\) for some \(j\), return \(u \rightarrow_{U}\left(u\left[m_{j}^{\ell} \cdot\left|x_{j}\right| \mapsto 0\right]\right) \downarrow\).
```

Fig. 3. A greedy strategy for showing arithmetic entailments in the theory $T_{\mathrm{s}}$. We write negcoeff $(u)$ to denote the set of terms whose coefficient is negative in $u$.

For simplicity, we silently assume in the figure that basic arithmetic simplifications are applied after each rewrite step to put the right-hand side in polynomial form.

Example 3. Let $u$ be $|\operatorname{replace}(x, \mathrm{aa}, \mathrm{b})|$. Because $\vdash|\mathrm{aa}| \geqslant|\mathrm{b}|$, the first case of the over-approximation rule for replace applies, and we get that $u \rightarrow_{O}|x|$. This reflects that the result of replacing the first occurrence, if any, of aa in $x$ with b is no longer than $x$.

Example 4. Let $u$ be the same as in the previous example and let $v$ be $-1 \cdot u+2 \cdot|x|$. Since $u \rightarrow_{O}|x|$ and the coefficient of $u$ in $v$ is negative, we have that $v \rightarrow_{U}-1 \cdot|x|+2 \cdot|x|$, which simplifies to $|x|$; moreover, $|x| \rightarrow_{U} 0$. Thus, $v \rightarrow_{U}^{*} 0$ and so $\vdash v \geqslant 0$. In other words, we can use the approximations to show that $u$ is at most $2 \cdot|x|$.

### 3.1 A Strategy for Approximation

The rewrite systems $\rightarrow_{O}$ and $\rightarrow_{U}$ allow for many possible derivations. Thus, it is important to devise a strategy that is efficient and succeeds often in practice. We use a greedy rule application strategy that favors rule applications leading to the cancellation of monomials. For example, consider the term $|x|-|\operatorname{substr}(y, 0,|x|)|$, and observe that the subtrahend can be over-approximated either by $|y|$ or by $|x|$. However, proving the $T_{\mathrm{S}}$-validity of $|x|-|\operatorname{substr}(y, 0,|x|)| \geqslant 0$ with the former over-approximation is impossible since $|x|-|y| \geqslant 0$ does not hold in all models of $T_{\mathrm{s}}$. In contrast, the latter approximation produces $|x|-|x| \geqslant 0$ which is trivially $T_{\mathrm{S}}$-valid.

Recall that, given an arithmetic inequality $u \geqslant 0$, our goal is to find a reduction $u \rightarrow{ }_{U}^{*} n$ where $n$ is a non-negative constant. Our strategy for choosing which rule of $\rightarrow_{U}$ to apply to $u$ is given in Figure 3. We decompose $u$ into three parts: the portion $u_{x}$ consisting of a sum of integer variables, the portion $u_{\ell}$ consisting of a sum of lengths of string variables, and the remaining portion $u_{s}$ which is a sum of monomials involving extended terms $v_{1}, \ldots, v_{q}$ as defined in Definition 1.

Since there are multiple choices for how terms in $u_{s}$ are approximated, the strategy focuses primarily on this portion. In particular, we apply an approximation for one of
the terms $v_{i}$, under-approximating or over-approximating depending on the sign of its coefficient, and replace the monomial in $t$ by its corresponding approximation. The choice of $v_{i}$ and $v_{i}^{a}$ is based on maximizing the likelihood that the overall derivation will produce a non-negative constant.

For a term $u$ in polynomial form, let negcoeff $(u)$ be a set of integer terms whose coefficient is negative in $u$, e.g. negcoeff $\left(y_{1}+-1 \cdot y_{2}\right)=\left\{y_{2}\right\}$. Terms in this set can be seen as obligations for proving entailments in our derivations since if $y_{2} \in \operatorname{negcoeff}(u)$, it must be the case that our derivation applies a rule that introduces a term with a positive coefficient for $y_{2}$. In Figure 3, we say that our choice of $v_{i} \rightarrow_{U} v_{i}^{a}$ avoids new terms if it does not have the effect of adding any new terms to negcoeff $(u)$, and cancels existing terms if it has the effect of removing terms from this set. If the portion $u_{s}$ is empty, we apply the rule $\left|x_{j}\right| \rightarrow_{U} 0$ if there exists a monomial $m_{j}^{\ell} \cdot\left|x_{j}\right|$ where $m_{j}^{\ell}$ is positive. This rule is applied with lowest priority because these monomials may help to cancel negative terms introduced by the other steps.

Step 1 depends on knowing the set of possible one-step approximations $v_{i} \rightarrow{ }_{U} v_{i}^{a}$ and $v_{i} \rightarrow_{O} v_{i}^{a}$ for terms from $u$. These are determined using the rules of Figure 2. Whenever applicable, we break ties between rewrites in Step 1 by considering a fixed arbitrary ordering over extended terms.

Example 5. Let $u$ be $1+\left|t_{1}\right|+\left|t_{2}\right|-\left|x_{1}\right|$, where $t_{1}$ is $\operatorname{substr}\left(x_{2}, 1,\left|x_{2}\right|+\left|x_{4}\right|\right)$ and $t_{2}$ is replace $\left(x_{1}, x_{2}, x_{3}\right)$. Step 1 of STR-ARITH-APPROX considers the possible approximations $\left|t_{1}\right| \rightarrow_{U}\left|x_{2}\right|-1$ and $\left|t_{2}\right| \rightarrow_{U}\left|x_{1}\right|-\left|x_{2}\right|$. Note that under-approximations are needed because the coefficients of $\left|t_{1}\right|$ and $\left|t_{2}\right|$ are positive. The first approximation is an instance of the third rule in Figure 2, noting that both $\vdash 1 \geqslant 0$ and $\vdash 1+\left|x_{2}\right|+\left|x_{4}\right| \geqslant\left|x_{2}\right|$ are derivable by a basic strategy that, wherever applicable, under-approximates string length terms as zero. Our strategy chooses the first approximation since it introduces no new negative coefficient terms, thus obtaining: $u \rightarrow_{U}\left|x_{2}\right|+\left|t_{2}\right|-\left|x_{1}\right|$. We now choose the approximation $\left|t_{2}\right| \rightarrow_{U}\left|x_{1}\right|-\left|x_{2}\right|$, noting that it introduces no new negative coefficient terms and cancels an existing one, $\left|x_{1}\right|$. After arithmetic simplification, we have derived $u \rightarrow_{U}^{*} 0$, and hence $\vdash u \geqslant 0$.

One can show that our strategy is sound, terminating, and deterministic. This means that applying STR-ARITH-APPROX to completion produces a unique rewrite chain of the form $t \rightarrow_{U} u_{1} \rightarrow_{U} \ldots \rightarrow_{U} u_{n}$ for a finite $n$, where each step is an application of one of the rewrite rules from Figure 2.

### 3.2 Simplification Rules with Arithmetic Side Conditions

We use the inference system from the previous section for simplifications of string terms with arithmetic side conditions. Figure 4 summarizes those simplifications.

The first rule rewrites a string equality to $\perp$ if one of the two sides can be inferred to be strictly longer than the other. In the second rule, if one side of an equality, con $(s, r, q)$, is such that the sum of lengths of $s$ and $q$ alone can be shown to be greater than or equal to the length of the other side, then $r$ must be empty. The third rule recognizes that string containment reduces to string equality when it can be inferred that string $s$ is at least as long as the string $t$ that must contain it. The next rule captures the fact

$$
\begin{aligned}
& t \approx s \rightarrow \quad \perp \\
& \text { if } \vdash|t| \geqslant|s|+1 \\
& t \approx \operatorname{con}(s, r, q) \rightarrow \quad t \approx \operatorname{con}(s, q) \wedge r \approx \epsilon \\
& \text { contains }(t, s) \rightarrow \quad t \approx s \\
& \operatorname{substr}(t, v, w) \rightarrow \quad \epsilon \\
& \text { if } \vdash|s|+|q| \geqslant|t| \\
& \text { if } \vdash|s| \geqslant|t| \\
& \text { if } \vdash 0>v \vee v \geqslant|t| \vee 0 \geqslant w \\
& \operatorname{substr}(\operatorname{con}(t, s), v, w) \rightarrow \quad \operatorname{substr}(s, v-|t|, w) \\
& \text { if } \vdash v \geqslant|t| \\
& \operatorname{substr}(\operatorname{con}(s, t), v, w) \rightarrow \quad \operatorname{substr}(s, v, w) \quad \text { if } \vdash|s| \geqslant v+w \\
& \operatorname{substr}(\operatorname{con}(t, s), 0, w) \rightarrow \quad \operatorname{con}(t, \operatorname{substr}(s, 0, w-|t|)) \text { if } \vdash w \geqslant|t| \\
& \text { indexof }(t, s, v) \rightarrow \quad \text { ite }(\operatorname{substr}(t, v) \approx s, v,-1) \text { if } \vdash v+|s| \geqslant|t|
\end{aligned}
$$

Fig. 4. String simplification rules. Letters $t, s, r, q$ denote string terms; $v, w$ denote integer terms.
that substring simplifies to the empty string if it can be inferred that its position $v$ is not within bounds, or its length $w$ is not positive. In the figure, we write that rule with a disjunctive side condition; this is a shorthand to denote that we can pick any disjunct and show that it holds assuming the negation of the other disjuncts. We can use those assumptions to perform substitutions to simplify the derivation. Concretely, to show $\vdash u_{1} \geqslant u_{2} \vee \ldots \vee u \not \approx u^{\prime}$ it is sufficient to infer $\vdash\left(u_{1} \geqslant u_{2}\right)\left[u \mapsto u^{\prime}\right]$. We demonstrate this with an example.

Example 6. Consider the term substr $(t,|t|+w, w)$. Our rules may simplify this term to $\epsilon$ by inferring that its start position $(|t|+w)$ is not within the bounds of $t$ if we assume that its size $(w)$ is positive. In detail, assume that $w>0$ (the negation of the last disjunct in the side condition of the fourth rule), which is equivalent to $w \approx|x|+1$ where $x$ is a fresh string variable and $|x|$ denotes an unknown non-negative quantity. It is sufficient to derive the formula obtained by replacing all occurrences of $w$ by $|x|+1$ in the disjunct $|t|+w \geqslant|t|$ to show that the start position of our term is out of bounds. After simplification, we obtain $|x|+1 \geqslant 0$, which is trivial to derive.

The next two rules in Figure 4 apply if we can infer respectively that the start position of the substring comes strictly after a prefix $t$ or that the end position of the substring comes strictly before a suffix $t$ of the first argument string. In either case, $t$ can be dropped.

Example 7. Let $t$ be $\operatorname{substr}\left(\operatorname{con}\left(x_{1}\right.\right.$, replace $\left.\left.\left(x_{2}, x_{3}, x_{4}\right)\right), 0, w\right)$, where $w$ is $\left|x_{1}\right|-\left|x_{2}\right|$. We have that $t \rightarrow \operatorname{substr}\left(x_{1}, 0, w\right)$, noting that $\vdash\left|x_{1}\right| \geqslant 0+\left|x_{1}\right|-\left|x_{2}\right|$. In other words, only the first component $x_{1}$ of the string concatenation is relevant to the substring since its end point must occur before the end of $x_{1}$.

The final rule for substr shows that a prefix of a substring can be pulled upwards if the start position is zero and we can infer that the substring is guaranteed to include at least a prefix string $t$. Finally, if we can infer that the last position of $s$ in $t$ starting from position $v$ is at or beyond the end of $t$, then the indexof term can be rewritten as an if-then-else (ite) term that checks whether $s$ is a suffix of $t$.

## 4 Containment-Based String Simplification

This section provides an overview of simplifications that are based on reasoning about the containment relationship between strings. We describe an inference system for

$$
\begin{array}{ccccc}
\frac{l_{1} \text { contains } l_{2}}{\vdash \operatorname{con}\left(l_{1}, t\right) \ni l_{2}} & \frac{\vdash s \ni r}{\vdash \operatorname{con}(t, s) \ni r} & \frac{\vdash t \ni^{s} r}{\vdash \operatorname{con}(t, s) \ni \operatorname{con}(r, q)} & \stackrel{\vdash \ni^{p} q}{\vdash t \ni \operatorname{substr}(t, v, w)} \\
\frac{l_{1} \text { does not contain } l_{2}}{\vdash} & \frac{\vdash r \nexists t}{\vdash l_{1} \nexists \operatorname{con}\left(l_{2}, t\right)} & \frac{\vdash r \nexists \operatorname{con}(s, t)}{\vdash} & \frac{\vdash l_{1} \backslash l_{2} \nexists t}{\vdash l_{1} \nexists \operatorname{con}\left(l_{2}, t\right)} \\
\frac{l_{2} \text { is a prefix of } l_{1}}{\vdash \operatorname{con}\left(l_{1}, t\right) \ni^{p} l_{2}} & \frac{\vdash s \ni^{p} r}{\vdash \operatorname{con}(t, s) \ni^{p} \operatorname{con}(t, r)} & \overline{\vdash t \ni^{p} t} & \frac{\vdash v \leqslant 0}{\vdash t \ni^{p} \operatorname{substr}(t, v, w)} \\
\frac{l_{2} \text { is a suffix of } l_{1}}{\vdash \operatorname{con}\left(t, l_{1}\right) \ni^{s} l_{2}} & \frac{\vdash s \ni^{s} r}{\vdash \operatorname{con}(s, t) \ni^{s} \operatorname{con}(r, t)} & \overline{\vdash t \ni^{s} t} & \frac{\vdash v+w \geqslant|t|}{\vdash t \ni^{s} \operatorname{substr}(t, v, w)}
\end{array}
$$

Fig. 5. Inferences for string containment $\ni$, is-prefix $\ni^{p}$ and is-suffix $\ni^{s}$. We write $l_{1}, l_{2}$ to denote string literals, $v, w$ to denote integer terms and $t, s, r, q$ to denote string terms. Although not shown here, we also assume rules that reflect the transitivity of $\ni, \ni^{p}$ and $\ni^{s}$.
deriving when one string is definitely contained or not contained in another. Following the notation from the last section, we write $\vdash t \ni s$ to denote the judgment of our inference system, denoting that string $t$ contains string $s$ in all models of $T_{\mathrm{S}}$. Conversely, we write $\vdash t \nexists s$ to denote string $t$ does not contain string $s$. We write $\vdash t \ni^{p} s$ (resp., $\vdash t \ni^{s} s$ ) to denote the judgment indicating that $s$ must be a prefix (resp., suffix) of $t$.

Rules for inferring judgments of these forms are given in Figure 5. Like our rules for arithmetic, these rules are solely based on the syntactic structure of terms, so inferences in this system can be computed statically. Both the assumptions and conclusions of the rules assume associativity of string concatenation with identity element $\epsilon$, that is, $\operatorname{con}(t, s)$ may refer to a term of the form $\operatorname{con}\left(\operatorname{con}\left(t_{1}, t_{2}\right), s\right)=\operatorname{con}\left(t_{1}, t_{2}, s\right)$ or alternatively to $\operatorname{con}(\epsilon, s)=s$. Most of the rules are straightforward. The inference system has special rules for substring terms $\operatorname{substr}(t, v, w)$, using arithmetic entailments from Section 3 to show prefix and suffix relationships with the base string $t$. For negative containment, the rules of the inference system together can show a (possibly non-constant) string cannot occur in a constant string by reasoning that its characters cannot appear in order in that string. We write $l_{1} \backslash l_{2}$ to denote the empty string if $l_{1}$ does not contain $l_{2}$, or the result of removing the smallest prefix of $l_{1}$ that contains $l_{2}$ from $l_{1}$ otherwise.

Example 8. Let $t$ be abcab and let $s$ be con(b, $x, \mathrm{a}, y, \mathrm{c})$. String $s$ is not contained in $t$ for any value of $x, y$. We derive $\vdash t \nexists s$ using two applications of the rightmost rule for negative containment in Figure 5, noting $a b c a b \backslash b=c a b$, $c a b \backslash a=b$, and $b$ does not contain $c$. In other words, the containment does not hold since the characters $b, a$ and $c$ cannot be found in order in the constant abcad.

### 4.1 Simplification Rules based on String Containment

Figure 6 gives rules for simplifying extended function terms based on the aforementioned judgments pertaining to string containment. First, equalities can be rewritten to false and applications of contains can be rewritten to a constant based on the appropriate judgment of our inference system. Applications of indexof can be simplified to -1 if it can be shown that the second argument does not appear in the suffix of the first

$$
\begin{aligned}
& t \approx s \rightarrow \quad \perp \quad \text { if } \vdash t \nexists s \\
& \text { contains }(t, s) \rightarrow \quad \perp \quad \text { if } \vdash t \nexists s \\
& \text { contains }(t, s) \rightarrow \quad \top \quad \text { if } \vdash t \ni s \\
& \begin{array}{lll}
\text { indexof }(t, s, v) \rightarrow & -1 & \text { if } \vdash \operatorname{substr}(t, v) \nexists s
\end{array} \\
& \text { indexof }(\operatorname{con}(t, r), s, v) \rightarrow \quad \text { indexof }(t, s, v) \quad \text { if } \vdash \operatorname{substr}(t, v) \ni s \\
& \text { indexof }(\operatorname{con}(t, r), s, v) \rightarrow \quad \operatorname{indexof}(r, s, v-|t|)+|t| \text { if } \vdash \operatorname{substr}(\operatorname{con}(t, r), v) \ni s \text { and } \\
& \vdash v \geqslant|t| \\
& \text { indexof }(t, s, v) \rightarrow \quad v \quad \text { if } \vdash \operatorname{substr}(t, v) \ni^{p} s \text { and } \vdash v<|t| \\
& \text { replace }(t, s, r) \rightarrow \quad t \quad \text { if } \vdash t \nexists s \\
& \text { replace }(\operatorname{con}(t, q), s, r) \rightarrow \quad \operatorname{con}(\text { replace }(t, s, r), q) \quad \text { if } \vdash t \ni s \\
& \text { replace }(t, s, r) \rightarrow \quad \operatorname{con}(r, \operatorname{substr}(t,|s|)) \quad \text { if } \vdash t \ni^{p} s
\end{aligned}
$$

Fig. 6. Simplification rules based on string containment.
argument starting at the position given by the third argument. The next two rules reason about cases where the second argument $s$ definitely occurs in the first argument starting from position $v$. In this case, if we additionally know that $s$ occurs within (beyond) a prefix $t$ of the first argument, then the suffix $r$ (prefix $t$ ) can be dropped, where the start position and the return value of the result are modified accordingly. If we know $s$ is a prefix of the first argument at position $v$, then the result is $v$ if indeed $v$ is in the bounds of $t$. Notice that the latter condition is necessary to handle the case where $s$ is the empty string. The three rules for replace are analogous. First, the replace rewrites to the first argument if we know it does not contain the second argument $s$. If we know $s$ is definitely contained in a prefix of the first argument, then we can pull the remainder of that string upwards. Finally, if we know $s$ is a prefix of the first argument, then we can replace that prefix with $r$ while concatenating the remainder. We use the term $\operatorname{substr}(t,|s|)$ to denote the remainder after the replacement for the sake of brevity, although this term typically does not involve extended functions after simplification, e.g. replace $(\operatorname{con}(x, y), x, z) \rightarrow \operatorname{con}(z, y)$ noting that $(\operatorname{substr}(\operatorname{con}(x, y),|x|)) \downarrow=y$, or replace $(\mathrm{ab}, \mathrm{a}, x) \rightarrow \operatorname{con}(x, \mathrm{~b})$ noting that $(\operatorname{substr}(\mathrm{ab},|\mathrm{a}|)) \downarrow=\mathrm{b}$.

### 4.2 Simplifications based on Equivalence of String Containment

We further refine our approach based on inferring when one containment is equivalent to another one. For example, con $(\mathrm{a}, x)$ is contained in con $(\mathrm{b}, y)$ if and only if con $(\mathrm{a}, x)$ is contained in $y$ alone. We introduce simplifications for such equivalences by reasoning about the maximal overlap between two strings.

We adapt and extend the notation given in previous work [13]. Given string literals $l_{1}$ and $l_{2}$, the sufficient left overlap of $l_{1}$ and $l_{2}$, written $l_{1} \sqcup_{l} l_{2}$, is the largest suffix of $l_{1}$ that is a prefix of $l_{2}$ or has $l_{2}$ as a prefix. For example, we have abc $\sqcup_{l} \mathrm{~cd}=\mathrm{c}$, $\mathrm{abc} \sqcup_{l} \mathrm{~b}=\mathrm{bc}$, and $\mathrm{abc} \sqcup_{l} \mathrm{ba}=\epsilon$. We extend this definition to arbitrary strings $s$ such that $l_{1} \sqcup_{l} s$ is equivalent to $l_{1} \sqcup_{l} l_{2}$ for the largest constant prefix $l_{2}$ of $s$, where notice that $l_{2}$ is the empty string if $s$ does not have a constant prefix. For example, we have $\mathrm{abc} \sqcup_{l} \operatorname{con}(\mathrm{cde}, y)=\mathrm{c}$, $\mathrm{abc} \sqcup_{l} \operatorname{con}(\mathrm{~b}, y)=\mathrm{bc}$, and $\mathrm{abc} \sqcup_{l} \operatorname{con}(\mathrm{a}, y)=\mathrm{abc}$. We define the dual operator sufficient right overlap, written $l_{1} \sqcup_{r} l_{2}$, which is the largest prefix of $l_{1}$ that is a suffix of $l_{2}$ or has $l_{2}$ as a suffix, e.g. abc $\sqcup_{r} \mathrm{~b}=\mathrm{ab}$, and extend this to

```
contains \((\operatorname{con}(t, l), s) \rightarrow \quad\) contains \(\left(\operatorname{con}\left(t, l \sqcup_{r} s\right), s\right)\)
contains \((\operatorname{con}(l, t), s) \rightarrow \quad\) contains \(\left(\operatorname{con}\left(l \sqcup_{l} s, t\right), s\right)\)
\(\operatorname{indexof}(\operatorname{con}(t, l), s, v) \rightarrow \quad \operatorname{indexof}\left(\operatorname{con}\left(t, l \sqcup_{r} s\right), s, v\right)\)
indexof \((\operatorname{con}(l, t), s, v) \rightarrow \quad\) indexof \(\left(\operatorname{con}\left(l_{2}, t\right), s, v-\left|l_{1}\right|\right) \quad\) if \(l=l_{1} \cdot l_{2}\) and \(l_{2}=l \sqcup_{l} s\)
    \(+\left|l_{1}\right| \quad \vdash \operatorname{substr}(\operatorname{con}(l, t), v) \ni s\)
replace \((\operatorname{con}(t, l), s, r) \rightarrow \quad \operatorname{con}\left(\right.\) replace \(\left.\left(\operatorname{con}\left(t, l_{1}\right), s, r\right), l_{2}\right)\) if \(l=l_{1} \cdot l_{2}\) and \(l_{1}=l \sqcup_{r} s\)
replace \((\operatorname{con}(l, t), s, r) \rightarrow \quad \operatorname{con}\left(l_{1}\right.\), replace \(\left.\left(\operatorname{con}\left(l_{2}, t\right), s, r\right)\right)\) if \(l=l_{1} \cdot l_{2}\) and \(l_{2}=l \sqcup_{l} s\)
```

Fig. 7. Simplification rules based on equivalence of string containment. We write $l, l_{1}, l_{2}$ to denote string literals, $v, w$ to denote integer terms and $t, s$ to denote string terms.
arbitrary strings in an analogous way. The sufficient left (resp., right) overlap operator can be used to determine how much of a constant string prefix $l_{1}$ (resp., suffix) can be safely removed from a string without impacting whether it contains another string.

The rules in Figure 7 simplify extended terms by considering string overlaps. The first two rules drop parts of string literals from the suffix or prefix of their first arguments. The two rules for indexof are similar: a suffix of the first argument can be dropped if it does not contribute to whether it contains the second argument. A prefix of an indexof term can be dropped if it does not contribute to containment, but only in the case where we know the second argument is definitely contained in the first argument. This is to guard against the case where the entire indexof term returns -1 . The rules for replace are similar to those for contains, except that the suffix (resp., prefix) of the first argument is pulled upwards instead of being dropped.

## 5 Multiset-Based String Simplification

Next, we introduce simplifications based on reasoning about strings as multisets, i.e. collections of unordered characters. Such reasoning is sufficient for showing that equalities like $\operatorname{con}(\mathrm{a}, x) \approx \operatorname{con}(x, \mathrm{~b})$ are equivalent to $\perp$, since the left side of the equality contains exactly one more occurrence of character a than the right-hand side. Similar to arithmetic reasoning from Section 3, we use approximations when reasoning about strings as multisets. We define the multiset abstraction of $t$, written $\mathcal{M}_{t}$, as the multiset $\left\{t_{1}, \ldots, t_{n}\right\}$ where $t$ is equivalent to con $\left(t_{1}, \ldots, t_{n}\right)$ and all constants in this set are characters. For example, $\mathcal{M}_{\mathrm{con}(\mathrm{aba}, x)}=\{\mathrm{a}, \mathrm{a}, \mathrm{b}, x\}$. We define a rewrite system $\rightarrow{ }_{O}^{\mathcal{M}}$ over strings where a rewritten string over-approximates the original string in the following sense: if $t \rightarrow{ }_{O}^{\mathcal{M}} s$, then for all models of $T_{\mathrm{S}}$ and any character $c$, the number of occurrences of $c$ in the strings in $\mathcal{M}_{s}$ is greater than or equal to the number of occurrences in the strings in $\mathcal{M}_{t}$.

Figure 8 lists the rules for the rewrite system $\rightarrow{ }_{O}^{\mathcal{M}}$ and the simplifications based on multiset reasoning. Given a predicate contains $(t, s)$, if over-approximating $t$ with respect to the rules of $\rightarrow{ }_{O}^{\mathcal{M}}$ results in a string $r$, and it can be determined that $s$ contains strictly more occurrences of some character $c$ than $r$, then it cannot be the case that $s$ is contained in $t$. To establish this, we check whether the multiset difference of $\mathcal{M}_{s}$ and $\mathcal{M}_{r}$ contains $c$, and conversely the difference of $\mathcal{M}_{r}$ and $\mathcal{M}_{s}$ contains only character constants which are distinct from $c$. In the second rule, if one side of an equality can be determined to contain only a character $c$, then one occurrence of that character can be dropped from both sides of the equality, since the relative position of that character does not

$$
\left.\begin{array}{rl}
\text { contains }(t, s) \rightarrow \perp \quad \text { if } t \rightarrow_{O}^{\mathcal{M}} * r \\
& \mathcal{M}_{s} \backslash \mathcal{M}_{r}=\left\{c, s_{1}, \ldots, s_{n}\right\} \text { and } \\
\mathcal{M}_{r} \backslash \mathcal{M}_{s}=\left\{c_{1}, \ldots, c_{m}\right\}
\end{array}\right] \begin{aligned}
& \operatorname{con}(t, s, s) \approx \operatorname{con}(q, c, r) \rightarrow \operatorname{con}(t, s) \approx \operatorname{con}(q, r) \text { if } \mathcal{M}_{\operatorname{con}(t, c, s)} \rightarrow \mathcal{M}_{O} *_{p} \text { and } \\
& \mathcal{M}_{p}=\{c, \ldots, c\}
\end{aligned}
$$

Fig. 8. Simplification rules based on multiset reasoning. We write $c, c_{1}, \ldots$ to denote characters, $v, w$ to denote integer terms, and $t, s, r, q, p$ to denote string terms.
matter. The three rules for $\rightarrow{ }_{O}^{\mathcal{M}}$ state that the multiset abstraction of a term of the form $\operatorname{substr}(t, v, w)$ can be over-approximated as the entire string $t$; a term replace $(t, s, r)$ can be over-approximated as a string having both $t$ and $r$; and over-approximation can be applied to the children of con terms.

Example 9. We have that con(aaa, $\left.\operatorname{substr}\left(x, y_{1}, y_{2}\right)\right) \approx \operatorname{con}(x, \mathrm{~b}) \rightarrow \perp$ by noting that con $\left(\right.$ aaa, $\left.\operatorname{substr}\left(x, y_{1}, y_{2}\right)\right) \rightarrow_{O}^{\mathcal{M}} * \operatorname{con}(\mathrm{aaa}, x), \mathcal{M}_{\mathrm{con}(\mathrm{aaa}, x)}=\{\mathrm{a}, \mathrm{a}, \mathrm{a}, x\}$ and $\mathcal{M}_{\mathrm{con}(x, \mathrm{~b})}=\{\mathrm{b}, x\}$. The difference of the latter with the former is $\{\mathrm{b}\}$, and the former with the latter is $\{a, a, a\}$. Thus, the right side of the equality contains at least one more occurrence of $b$ than the left side; hence, the equality is equivalent to false.

## 6 Implementation

We implemented the above simplification rules and others in the DPLL-based SMT solver CVC4, which implements a theory solver for a basic fragment of word equations with length, several other theory solvers, and reduction techniques for extended string functions as described in Section 2.1. Our simplification rules are run in a preprocessing pass as well as an inprocessing pass during solving. For the latter, we use a contextdependent simplification strategy that infers when an extended string constraint, e.g., $\neg$ contains $(t, s)$, simplifies to $\perp$ based on other assertions, e.g., $s \approx \epsilon$. Our simplification techniques do not affect the core procedure for the theory of strings, nor the compatibility of the string solver with other theories. In total, our implementation is about 3,500 lines of C++ code. We cache the results of the simplifications and the approximation-based arithmetic entailments to amortize their costs.
Additional Simplification Rules The simplification rules in this paper are a subset of the rules in the implementation. We omit other uncategorized rules for lack of space. Many of these apply to specific term patterns, such as cases where two nested applications of substr can be combined; cases where an application of replace can be eliminated by case splitting; and other cases like con $(t, t) \approx a \rightarrow \perp$. An example of such rules is contains(replace $\left.\left(t, w_{1}, w_{2}\right), w_{3}\right) \rightarrow \operatorname{contains}\left(t, w_{3}\right)$ if $w_{3}$ does not overlap with either $w_{1}$ or $w_{2}$, because the replace does not change whether $t$ contains $w_{3}$ or not. Another class of rules only applies to strings of length one because they cannot span multiple components of a concatenations, e.g. contains $(\operatorname{con}(t, s), c) \rightarrow \operatorname{contains}(t, c) \vee \operatorname{contains}(s, c)$
where $c$ is a character. Finally, there are rewrites that benefit from multiple techniques presented in this paper. For example, we have a rewrite that splits string equations into multiple smaller equations if it can determine that prefixes must have the same length: $\operatorname{con}(\mathrm{a}, t, s) \approx \operatorname{con}(t, \mathrm{~b}, r) \rightarrow \operatorname{con}(\mathrm{a}, t) \approx \operatorname{con}(t, \mathrm{~b}) \wedge s \approx r \rightarrow \perp$.
Validating Simplification Rules The correctness of our simplification techniques is critical to the soundness of the overall solver. Due to the sophistication and breadth of those techniques, it is challenging to formally verify our implementation. As a pragmatic alternative, we periodically test our implementation using a testing infrastructure we developed for this purpose. We found this to be critical in our development process. Our testing infrastructure allows the developer to specify a context-free grammar in the syntax-guided synthesis format [2]. We generate all terms $t$ in this grammar up to a fixed size and test the equivalence of $t$ and its simplified form $t \downarrow$ on a set of randomly generated points. The most recent run of this system on two grammars (one for extended string terms and another for string predicates) up to a term size of three, validated 319,867 simplifications of string terms and 188,428 simplifications of string predicates on 1,000 sample points. This run took $924 s$ for string terms and $971 s$ for the string predicates using the same hardware as in Section 7.

## 7 Evaluation

We evaluate the impact of each simplification technique as implemented in CVC4 on three benchmark sets that use extended string operators: CMU, a dataset obtained from symbolic execution of Python code [13]; TERMEQ, a benchmark set consisting of the verification of term equivalences over strings [12]; and SLOG, a benchmark set extracted from vulnerability testing of web applications [20]. The SLOG set uses the replace function extensively but does not contain other extended functions. We also evaluate the impact on APLAS, a set of handcrafted benchmarks involving looping word equations [9] (string equalities whose left and right sides have variables in common).

We compare CVC 4 with Z 3 commit $9 \mathrm{cb} 1 \mathrm{aOf}[7],{ }^{6}$ a state-of-the-art string solver. Additionally, we compare against OSTRICH on the Slog benchmarks but not other sets because it does not support some functions such as contains and indexof. We omit a comparison with Z3STR3 4.8.4 because we found multiple issues in its latest release including wrong answers, which we have reported to the authors. We also omit a comparison with S3\# due to differing semantics. We compare four configurations of CVC4: all, which enables all optimizations; -arith, which disables arithmetic-based simplification techniques (discussed in Section 3); -contain, which disables containmentbased simplification techniques (discussed in Section 4); and -msets, which disables multiset-based simplification techniques (discussed in Section 5). Additionally, to test the applicability of our techniques to other solvers, we test the effect of our simplifications on Z3 by using CVC4 to generate simplified benchmarks and then running Z3 on those benchmarks. We generate a set of simplified benchmarks that are simplified with CVC4 with $\left(\mathrm{Z} 3_{f}\right)$ and without $\left(\mathrm{Z} 3_{b}\right)$ the simplification techniques presented in this paper.

We ran all benchmarks on a cluster equipped with Intel E5-2637 v4 CPUs running Ubuntu 16.04 and dedicated one core, 8GB RAM, and 600 seconds for each job. Table 1

[^3]| Set |  | all | -arith | -contain | -msets | z3 | z3 | z3 | OsTRICH | R\% |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
|  | sat | $\mathbf{5 7 0 3}$ | 5535 | $\mathbf{5 7 0 3}$ | $\mathbf{5 7 0 3}$ | 2343 | 3923 | 3943 |  |  |
| CMU | unsat | $\mathbf{6 5}$ | 29 | $\mathbf{6 5}$ | $\mathbf{6 5}$ | 50 | 58 | 61 |  | $32 \%$ |
|  | $\times$ | 154 | 358 | 154 | 154 | 3529 | 1941 | 1918 |  |  |
|  | sat | $\mathbf{1 0}$ | $\mathbf{1 0}$ | $\mathbf{1 0}$ | $\mathbf{1 0}$ | 4 | 5 | 5 |  |  |
| TERMEQ | unsat | 51 | 37 | 28 | 51 | 35 | 40 | $\mathbf{6 0}$ |  | $68 \%$ |
|  | $\times$ | 19 | 33 | 42 | 19 | 41 | 35 | 15 |  |  |
|  | sat | 1302 | 1302 | 1302 | 1302 | 1133 | 1225 | 1225 | $\mathbf{1 3 0 4}$ |  |
| SLOG | unsat | $\mathbf{2 0 8 2}$ | $\mathbf{2 0 8 2}$ | $\mathbf{2 0 8 2}$ | $\mathbf{2 0 8 2}$ | 2080 | 2080 | 2080 | $\mathbf{2 0 8 2}$ | $27 \%$ |
|  | $\times$ | 7 | 7 | 7 | 7 | 178 | 86 | 86 | 5 |  |
|  | sat | $\mathbf{1 3 5}$ | $\mathbf{1 3 5}$ | $\mathbf{1 3 5}$ | $\mathbf{1 3 5}$ | 9 | 51 | 46 |  |  |
| APLAS | unsat | $\mathbf{2 9 2}$ | $\mathbf{2 9 2}$ | 171 | 171 | 94 | 129 | $\mathbf{2 9 2}$ |  | n/a |
|  | $\times$ | 159 | 159 | 280 | 280 | 483 | 406 | 248 |  |  |
|  | sat | $\mathbf{7 1 5 0}$ | 6982 | $\mathbf{7 1 5 0}$ | $\mathbf{7 1 5 0}$ | 3489 | 5204 | 5219 | 1304 |  |
| Total | unsat | 2490 | 2440 | 2346 | 2369 | 2259 | 2307 | $\mathbf{2 4 9 3}$ | 2082 |  |
|  | $\times$ | 339 | 557 | 483 | 460 | 4231 | 2468 | 2267 | 5 |  |

Table 1. Number of solved problems per benchmark set. Best results are in bold. Gray cells indicate benchmark sets not supported by a solver. " $\mathrm{R} \%$ " indicates the reduction of extended string functions during preprocessing. All benchmarks ran with a timeout of 600 s .
summarizes the number of solved instances for each configuration and the baseline solvers grouped by benchmark sets. We remark that the average reduction of extended string functions (with all simplification techniques enabled) shown in column "R\%" is significant on all benchmark sets. The scatter plots in Figure 9 detail the effects of disabling each family of simplifications. They distinguish between satisfiable and unsatisfiable instances. To emphasize non-trivial benchmarks, we omit the benchmarks that are solved in less than a second by all solvers.

The arithmetic-based simplification techniques have the most significant performance impact on the symbolic execution benchmarks CMU. The number of solved benchmarks is significantly lower when disabling those techniques. The scatter plot shows that for longer running satisfiable queries there is a large portion of the benchmarks that are solved up to an order of magnitude faster with the simplifications. These improvements in runtime on the CMU set are particularly compelling because they come from a symbolic execution application, which involves a large number of queries with a short timeout. The improvements are more pronounced for unsatisfiable benchmarks, where our results show that simplifications often give the solver the ability to derive a refutation in a matter of seconds, something that is infeasible with configurations without these techniques. The Aplas set contains no extended string operators and hence our arithmetic-based simplification techniques have little impact on this set.

In contrast, both containment and multiset-based rewrites have a high impact on the APLAS set, as -contain and -msets both solve 121 fewer benchmarks. Additionally, -contain has a high impact on the TERMEQ set, where the simplifications enable the best configuration to solve 61 out of 80 benchmarks. Since these techniques apply most


Fig. 9. Scatter plots showing the impact of disabling simplification techniques in CVC 4 on both satisfiable and unsatisfiable benchmarks. All benchmarks ran with a timeout of 600 s .
frequently to looping word equations, they are less important for the CMU set, which does not have such equations. The containment-based and multiset-based techniques primarily help on unsatisfiable benchmarks, as shown in the scatter plots. On TERMEQ benchmarks, it tends to be easier to find counterexamples, i.e. to solve the satisfiable ones, so there is more to gain on unsatisfiable benchmarks.

On Slog, Ostrich solves two more instances than CVC4 but CVC4 is over 50 times faster on commonly solved instances while supporting a richer set of string operators. On all benchmark sets, CVC4 solves at least as many benchmarks as Z3 and CVC4 has $12 \times$ fewer timeouts than Z3. On the simplified benchmarks, Z3 performs significantly better. On the CMU and the APLAS benchmarks, $\mathrm{Z} 3_{b}$ outperforms Z 3 by a large margin. Additionally simplifying the benchmarks with the techniques presented in this paper improves performance further on most benchmark sets and allows $\mathrm{Z} 3_{f}$ to solve the most unsatisfiable benchmarks overall. These results indicate that Z3 could benefit from additional simplifications, and they underscore the importance of curating and publishing simplification techniques in order to improve the state-of-the-art.

## 8 Conclusion

We have presented a set of aggressive simplification techniques for reasoning about extended string constraints. Our results suggest that such techniques are key to advancing the state of the art in SMT string solving. Arithmetic-based simplifications lead to significant speedups in benchmarks from a symbolic execution application, while containment and multiset-based simplifications improve the performance on problems consisting of difficult term equivalences and looping word equations. Our approach is not limited to CVC4 and can be adapted to other solvers.

Given the encouraging results for each of the simplification techniques in our evaluation, we plan to extend them to other types of abstraction and make them context-aware. The latter extension involves taking into account other assertions when checking whether a side condition of a rule is fulfilled.

Acknowledgements This work was partially supported by the National Science Foundation under award 1656926, the Defense Advanced Research Projects Agency under award FA8650-18-2-7854, and Amazon Web Services.

## References

[1] P. A. Abdulla, M. F. Atig, Y. Chen, B. P. Diep, L. Holík, A. Rezine, and P. Rümmer. Trau: SMT solver for string constraints. In N. Bjørner and A. Gurfinkel, editors, 2018 Formal Methods in Computer Aided Design, FMCAD 2018, Austin, TX, USA, October 30 - November 2, 2018, pages 1-5. IEEE, 2018.
[2] R. Alur, R. Bodík, E. Dallal, D. Fisman, P. Garg, G. Juniwal, H. Kress-Gazit, P. Madhusudan, M. M. K. Martin, M. Raghothaman, S. Saha, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In M. Irlbeck, D. A. Peled, and A. Pretschner, editors, Dependable Software Systems Engineering, volume 40 of NATO Science for Peace and Security Series, D: Information and Communication Security, pages 1-25. IOS Press, 2015.
[3] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In G. Gopalakrishnan and S. Qadeer, editors, Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings, volume 6806 of Lecture Notes in Computer Science, pages 171-177. Springer, 2011.
[4] M. Berzish, V. Ganesh, and Y. Zheng. Z3str3: A string solver with theory-aware heuristics. In D. Stewart and G. Weissenbacher, editors, 2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, October 2-6, 2017, pages 55-59. IEEE, 2017.
[5] N. Bjørner, N. Tillmann, and A. Voronkov. Path feasibility analysis for string-manipulating programs. In S. Kowalewski and A. Philippou, editors, Tools and Algorithms for the Construction and Analysis of Systems, 15th International Conference, TACAS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings, volume 5505 of Lecture Notes in Computer Science, pages 307-321. Springer, 2009.
[6] T. Chen, M. Hague, A. W. Lin, P. Rümmer, and Z. Wu. Decision procedures for path feasibility of string-manipulating programs with complex operations. PACMPL, 3(POPL):49:149:30, 2019.
[7] L. M. de Moura and N. Bjørner. Z3: an efficient SMT solver. In C. R. Ramakrishnan and J. Rehof, editors, Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings, volume 4963 of Lecture Notes in Computer Science, pages 337-340. Springer, 2008.
[8] A. Kiezun, V. Ganesh, S. Artzi, P. J. Guo, P. Hooimeijer, and M. D. Ernst. HAMPI: A solver for word equations over strings, regular expressions, and context-free grammars. ACM Trans. Softw. Eng. Methodol., 21(4):25:1-25:28, 2012.
[9] Q. L. Le and M. He. A decision procedure for string logic with quadratic equations, regular expressions and length constraints. In S. Ryu, editor, Programming Languages and Systems - 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2-6, 2018, Proceedings, volume 11275 of Lecture Notes in Computer Science, pages 350-372. Springer, 2018.
[10] G. Li and I. Ghosh. PASS: string solving with parameterized array and interval automaton. In V. Bertacco and A. Legay, editors, Hardware and Software: Verification and Testing - 9th International Haifa Verification Conference, HVC 2013, Haifa, Israel, November 5-7, 2013, Proceedings, volume 8244 of Lecture Notes in Computer Science, pages 15-31. Springer, 2013.
[11] T. Liang, A. Reynolds, C. Tinelli, C. Barrett, and M. Deters. A DPLL(T) theory solver for a theory of strings and regular expressions. In A. Biere and R. Bloem, editors, Computer

Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings, volume 8559 of Lecture Notes in Computer Science, pages 646-662. Springer, 2014.
[12] A. Reynolds, H. Barbosa, A. Niemetz, A. Nötzli, M. Preiner, C. Barrett, and C. Tinelli. Rewrites for SMT solvers using syntax-guided enumeration. SMT, 2018.
[13] A. Reynolds, M. Woo, C. Barrett, D. Brumley, T. Liang, and C. Tinelli. Scaling up DPLL(T) string solvers using context-dependent simplification. In R. Majumdar and V. Kuncak, editors, Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part II, volume 10427 of Lecture Notes in Computer Science, pages 453-474. Springer, 2017.
[14] P. Saxena, D. Akhawe, S. Hanna, F. Mao, S. McCamant, and D. Song. A symbolic execution framework for javascript. In 31st IEEE Symposium on Security and Privacy, S\&P 2010, 1619 May 2010, Berleley/Oakland, California, USA, pages 513-528. IEEE Computer Society, 2010.
[15] C. Tinelli, C. Barrett, and P. Fontaine. Unicode Strings (Draft 1.0). http: / /smt lib.cs . uiowa.edu/theories-UnicodeStrings.shtml, 2018.
[16] M. Trinh, D. Chu, and J. Jaffar. S3: A symbolic string solver for vulnerability detection in web applications. In G. Ahn, M. Yung, and N. Li, editors, Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014, pages 1232-1243. ACM, 2014.
[17] M. Trinh, D. Chu, and J. Jaffar. Progressive reasoning over recursively-defined strings. In S. Chaudhuri and A. Farzan, editors, Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I, volume 9779 of Lecture Notes in Computer Science, pages 218-240. Springer, 2016.
[18] M. Trinh, D. Chu, and J. Jaffar. Model counting for recursively-defined strings. In R. Majumdar and V. Kuncak, editors, Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part II, volume 10427 of Lecture Notes in Computer Science, pages 399-418. Springer, 2017.
[19] M. Veanes, N. Tillmann, and J. de Halleux. Qex: Symbolic SQL query explorer. In E. M. Clarke and A. Voronkov, editors, Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16, Dakar, Senegal, April 25-May 1, 2010, Revised Selected Papers, volume 6355 of Lecture Notes in Computer Science, pages 425-446. Springer, 2010.
[20] H. Wang, T. Tsai, C. Lin, F. Yu, and J. R. Jiang. String analysis via automata manipulation with logic circuit representation. In S. Chaudhuri and A. Farzan, editors, Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I, volume 9779 of Lecture Notes in Computer Science, pages 241-260. Springer, 2016.
[21] F. Yu, M. Alkhalaf, and T. Bultan. Stranger: An automata-based string analysis tool for PHP. In J. Esparza and R. Majumdar, editors, Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings, volume 6015 of Lecture Notes in Computer Science, pages 154-157. Springer, 2010.
[22] Y. Zheng, V. Ganesh, S. Subramanian, O. Tripp, M. Berzish, J. Dolby, and X. Zhang. Z3str2: an efficient solver for strings, regular expressions, and length constraints. Formal Methods in System Design, 50(2-3):249-288, 2017.
[23] Y. Zheng, X. Zhang, and V. Ganesh. Z3-str: a z3-based string solver for web application analysis. In B. Meyer, L. Baresi, and M. Mezini, editors, Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of

Software Engineering, ESEC/FSE'13, Saint Petersburg, Russian Federation, August 18-26, 2013, pages $114-124$. ACM, 2013


[^0]:    ${ }^{3}$ Our implementation supports a larger set of symbols, but for brevity, we only show the subset of the symbols used throughout this paper.

[^1]:    ${ }^{4}$ This formula is a simplified form of the general reduction. The general reduction also expresses that $i$ is a valid position in $t$ and that the third argument of substr is non-negative [13].

[^2]:    ${ }^{5}$ Some specialized rules have been omitted for space reasons.

[^3]:    ${ }^{6} 9 \mathrm{cbla} 0 \mathrm{f}$ is newer than the current release 4.8.4 and includes several fixes for critical issues.

