

CertiStr: A Certified String Solver

Shuanglong Kan

Department of Computer Science
Technische Universität Kaiserslautern
Kaiserslautern, Germany
shuanglong@cs.uni-kl.de

Philipp Rümmer

Department of Information Technology
Uppsala University
Uppsala, Sweden
philipp.ruemmer@it.uu.se

Anthony Widjaja Lin

Department of Computer Science
Technische Universität Kaiserslautern & MPI-SWS
Kaiserslautern, Germany
lin@cs.uni-kl.de

Micha Schrader

Department of Computer Science
Technische Universität Kaiserslautern
Kaiserslautern, Germany
schrader@rhrk.uni-kl.de

Abstract

Theories over strings are among the most heavily researched logical theories in the SMT community in the past decade, owing to the error-prone nature of string manipulations, which often leads to security vulnerabilities (e.g. cross-site scripting and code injection). The majority of the existing decision procedures and solvers for these theories are themselves intricate; they are complicated algorithmically, and also have to deal with a very rich vocabulary of operations. This has led to a plethora of bugs in implementation, which have for instance been discovered through fuzzing.

In this paper, we present CertiStr, a certified implementation of a string constraint solver for the theory of strings with concatenation and regular constraints. CertiStr aims to solve string constraints using a forward-propagation algorithm based on symbolic representations of regular constraints as symbolic automata, which returns three results: *sat*, *unsat*, and *unknown*, and is guaranteed to terminate for the string constraints whose concatenation dependencies are acyclic. The implementation has been developed and proven correct in Isabelle/HOL, through which an effective solver in OCaml was generated. We demonstrate the effectiveness and efficiency of CertiStr against the standard Kaluza benchmark, in which 80.4% tests are in the string constraint fragment of CertiStr. Of these 80.4% tests, CertiStr can solve 83.5% (i.e. CertiStr returns *sat* or *unsat*) within 60s.

CCS Concepts: • Theory of computation → Automated reasoning; Formal languages and automata theory.



This work is licensed under a Creative Commons Attribution 4.0 International License.

CPP '22, January 17–18, 2022, Philadelphia, PA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9182-5/22/01.

<https://doi.org/10.1145/3497775.3503691>

Keywords: string theory, Isabelle, symbolic automata, SMT solvers

ACM Reference Format:

Shuanglong Kan, Anthony Widjaja Lin, Philipp Rümmer, and Micha Schrader. 2022. CertiStr: A Certified String Solver. In *Proceedings of the 11th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP '22)*, January 17–18, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3497775.3503691>

1 Introduction

Strings are among the most fundamental and commonly used data types in virtually all modern programming languages, especially with the rapidly growing popularity of dynamic languages, including JavaScript and Python. Programs written in such languages often implement security-critical infrastructure, for instance web applications, and they tend to process data and code in string representation by applying built-in string-manipulating functions; for instance, to split, concatenate, encode/decode, match, or replace parts of a string. Functions of this kind are complex to reason about and can easily lead to programming mistakes. In some cases, such mistakes can have serious consequences, e.g., in the case of web applications, cross-site scripting (XSS) attacks can be used by a malicious user to attack both the web server or the browsers of other users.

One promising research direction, which has been intensively pursued in the SMT community in the past ten years, is the development of SMT solvers for theories of strings (dubbed *string solvers*) including Kaluza [35], CVC4 [28], Z3 [17], Z3-str3 [4], Z3-Trau [9], S3P [39], OSTRICH [13], SLOTH [22], and Norn [1], to name only a few. Such solvers are highly optimized and find application, among others, in bounded model checkers and symbolic execution tools [14, 34], but also in tools tailored to verification of the security properties [3].

It has long been observed that constraint solvers are extremely complicated procedures, and that implementations

are prone to bugs. Defects that affect soundness or completeness are routinely found even in well-maintained state-of-the-art tools, both in real-world applications and through techniques like fuzzing [6, 8, 31]. String solvers are particularly troublesome in this context, since, unlike other SMT theories, the theory of strings requires a multitude of operations including concatenation, regular constraints, string replacement, length constraints, and many others. This inherent intricacy is also reflected in the recently developed SMT-LIB standard for strings (<http://smtlib.cs.uiowa.edu/theories-UnicodeStrings.shtml>). Many of the string solvers that arose in the past decade relied on inevitably intricate decision procedures, which resulted in subtle bugs in implementations themselves even among the mainstream string solvers (e.g. see [8, 31]). This implies the unfortunate fact that one *cannot* blindly trust the answer provided by string solvers.

Contributions. In this paper, we present CertiStr, the *first certified implementation of a string solver* for the standard theory of strings with concatenation and regular constraints (a.k.a. word equations with regular constraints [18, 20, 23, 30]). CertiStr aims to solve string constraints by means of the so-called *forward-propagation algorithm*, which relies on a simple idea of propagating regular constraints in a *forward direction* in order to derive contradiction, or prove the absence thereof (see Section 2 for an example). Similar ideas are already present in abstract interpretation of string-manipulating programs (e.g. see [32, 42, 43]), but not yet at the level of string solvers, i.e., which operate on string constraints. We have proven in Isabelle/HOL [33] some crucial properties of the forward-propagation algorithm and more specifically with respect to our implementation CertiStr: (i) [TERMINATION]: the algorithm terminates on string constraints without cyclic concatenation dependencies, (ii) [SOUNDNESS]: CertiStr is sound for *unsatisfiable* results (if forward-propagation detects inconsistencies in a constraint, then the constraint is indeed unsatisfiable), and (iii) [COMPLETENESS]: CertiStr is complete for constraints satisfying the *tree property* (i.e., in which on the right-hand side of each equation every variable appears at most once). Here completeness amounts to the fact that, if forward-propagation does not detect inconsistencies for a constraint satisfying the tree property, then the constraint is indeed *satisfiable*. For the constraints that do not satisfy the *tree property* and CertiStr cannot decide them as *unsatisfiable*, CertiStr returns *unknown*.

In order to facilitate the propagation of regular constraints, we also implement a certified library for Symbolic Non-deterministic Finite Automata (s-NFA), which contains various automata operations, such as the concatenation and product of two s-NFAs, the language emptiness checking for s-NFAs, among others. s-NFAs — as introduced by Veanes et al. [41] (see [16] for more details) — are different from classic

Non-deterministic Finite Automata (NFA) by allowing transition labels to be a set of characters in the (potentially infinite) alphabet, represented by an element of a boolean algebra (e.g. the interval algebra or the BDD algebra), instead of a single character. s-NFAs are especially crucial for an efficient implementation of automata-based string solving algorithms and many string processing algorithms [1, 13, 16, 24, 41], for which reason our certified implementation of an s-NFA library is of independent interests.

Last but not least, we have automatically generated a verified implementation CertiStr in Isabelle/HOL, which we have extensively evaluated against the standard string solving benchmark from Kaluza [35] with around 38000 string constraints. For the first time, we demonstrate that the simple forward-propagation algorithm in fact performs surprisingly well even compared to other highly optimized solvers (which are not verified implementations). In particular, we show that the majority of these constraints (83.5%) are solved by CertiStr; for the rest, the tool either returns *unknown*, or times out. Moreover, CertiStr terminates within 60 seconds on 98% of the constraints, witnessing its efficiency.

To summarize, our contributions are:

- we developed in Isabelle/HOL the tool CertiStr, the first certified implementation of a string solver.
- we implemented the first certified symbolic automata library, which is crucial for an efficient implementation of many string processing applications [16].
- CertiStr was evaluated over Kaluza benchmark [35], with around 38000 tests. In this benchmark, 83.5% of the tests can be solved with the results *sat* or *unsat*. Moreover, the solver terminates within 60 seconds on 98% of the tests, witnessing the surprising competitiveness of the simple forward-propagation algorithm against more complicated algorithms.

2 Motivating Example

We start by illustrating the decision procedure implemented by CertiStr. CertiStr uses the so-called *forward-propagation algorithm* for solving satisfiability of string theory with concatenation and regular membership constraints. To illustrate the algorithm, consider the formula:

$$\text{domain} \in /[\text{a-zA-Z.}] + / \quad (1)$$

$$\wedge \text{dir, file} \in /[\text{a-zA-Z0-9.}] + / \quad (2)$$

$$\wedge \text{path} = \text{dir} + "/" + \text{file} \quad (3)$$

$$\wedge \text{url} = \text{"http://"} + \text{domain} + "/" + \text{path} \quad (4)$$

$$[\wedge \text{url} \in /. * < \text{script} > . * /] \quad (5)$$

The formula contains string variables (*domain*, *dir*, etc.), and uses both regular membership constraints and word equations with concatenation to model the construction of a URL from individual components. Regular expressions are written using the standard PCRE syntax (<https://perldoc.perl>).

org/perltre). Each equation should be read as an assignment of the right-hand side term to the left-hand side variable, and is processed in this direction. We initially ignore the constraint (5).

The forward-propagation algorithm propagates regular membership constraints and derives new constraints for the variable on the left-hand side of equations. In this example, the algorithm starts with equation (3) and propagates constraint (2), resulting in the new constraint

$$path \in / [a-zA-Z0-9.]+\backslash/[a-zA-Z0-9.]+/ \quad (6)$$

We can propagate this new constraint further, using equation (4) and together with constraint (1), deriving:

$$url \in / http:\backslash\backslash/[a-zA-Z.]+\backslash/[a-zA-Z0-9.]+\backslash/[a-zA-Z0-9.]+/ \quad (7)$$

At this point, no more propagations are possible. Note that forward-propagation will terminate, in general, whenever no cyclic dependencies exist between the equations, which is the fragment of formulas we consider.

Still ignoring (5), we can observe that forward-propagation has not discovered any inconsistent constraints. In order to conclude the *satisfiability* of the formula from this, we need a meta-result about forward-propagation: we show that forward propagation is complete for acyclic formulas in which each variable occurs at most once on the right-hand side of equations. Since this criterion holds for the formula (1) $\wedge \dots \wedge$ (4), we have indeed proven that it is satisfiable.

Consider now also (5), modelling a simple form of injection attack. We can observe that this additional constraint is inconsistent with the derived constraint (7), since the intersection of the asserted regular languages is empty. Forward-propagation detects this conflict as soon as (7) has been computed: each time a new membership constraint is derived, the algorithm will check the consistency with constraints already assumed. Since we also show that forward-propagation is sound w.r.t. inconsistency, the detection of a conflict immediately implies that the formula (1) $\wedge \dots \wedge$ (5) is unsatisfiable.

In summary, the forward-propagation is defined using two main inference steps: the post-image computation for word equations, and a consistency check for regular constraints. Although the example uses regular expressions, both operations can be defined more easily through a translation to finite-state automata. To obtain a verified string solver, the implementation of both operations has to be shown correct, and in addition meta-results need to be derived about the soundness and completeness of the overall algorithm.

We can also observe that classical finite-state automata, over concrete alphabets, are cumbersome even for toy examples; regular expressions often talk about character ranges that would require a large number of individual transitions. In practice, string constraints are usually formulated over the

Unicode alphabet with currently 3×2^{16} characters, which necessitates symbolic character representation. In our verified implementation, we therefore use a simple form of symbolic automata [15] in which transitions are labelled with character intervals.

3 String Constraint Fragment

In this section, we present the string constraint fragment that CertiStr supports and its semantics of satisfiability.

Let Σ be an alphabet and Σ^* denote the set of words over Σ . Let $w_1, w_2 \in \Sigma^*$ be two words, $w_1 w_2$ denotes the concatenation of the two words, i.e., w_2 is appended to the end of w_1 to get a new string. We consider the following string constraint language over Σ :

$$c ::= x \in \mathcal{A} \mid x = x_1 + x_2 \mid c_1 \wedge c_2,$$

where x denotes a variable ranging over Σ^* and \mathcal{A} denotes an NFA representing a regular language over Σ . The language accepted by \mathcal{A} is denoted by $\mathcal{L}(\mathcal{A})$. The semantics of a constraint c is defined in an obvious way by interpreting $x \in \mathcal{A}$ as a membership of x in the language $\mathcal{L}(\mathcal{A})$, $x = x_1 + x_2$ as a string equation, and $c_1 \wedge c_2$ as a conjunction of c_1 and c_2 . More precisely, given a function μ mapping each variable in the constraint c to a string in Σ^* , we say that μ *satisfies* c if: (i) $\mu(x) \in \mathcal{L}(\mathcal{A})$ for each constraint $x \in \mathcal{A}$ in c , and (ii) $\mu(x) = \mu(x_1)\mu(x_2)$ for each constraint $x = x_1 + x_2$. The constraint c is satisfiable if one such *solution* μ for c exists.

Our string constraint language is as general as word equations with regular constraints [18, 20, 23, 30], which forms the basis of the recently published Unicode string constraint language in SMT-LIB 2.6, and already makes up the bulk of existing string constraint benchmarks. Notice that our restriction to string equation of the form $x = x_1 + x_2$ is not a real restriction since general string constraints can be obtained by means of desugaring. For instance, the concatenation of more than two variables, like $x = x_1 + x_2 + x_3$ can be desugared as two concatenation constraints: (1) $x' = x_1 + x_2$; $x = x' + x_3$. Moreover, a subset of string constraints with length functions (called monadic length) can also be translated into this fragment with regular constraints. For instance, $|x| \leq 6$, where $|x|$ denotes the length of the word of x , can be translated to a regular membership constraint as $x \in \mathcal{A}$, where \mathcal{A} is an NFA that accepts any word whose length is at most 6. Recent experimental evidence shows that a substantial portion of length constraints that appear in practice are essentially monadic [21]. There are also some string constraints using disjunction (\vee). For instance, $c_1 \vee c_2$, where c_1 and c_2 are string constraints in our fragment. This also can be solved with CertiStr by first solving c_1 and c_2 separately, and then checking whether one of them is satisfiable.

Figure 1 shows the framework of CertiStr. The input of the solver is an SMT file with string constraints. The tool has two parts: (1) the front-end (non-certified) and (2) the back-end (certified). The front-end parses the SMT file, which

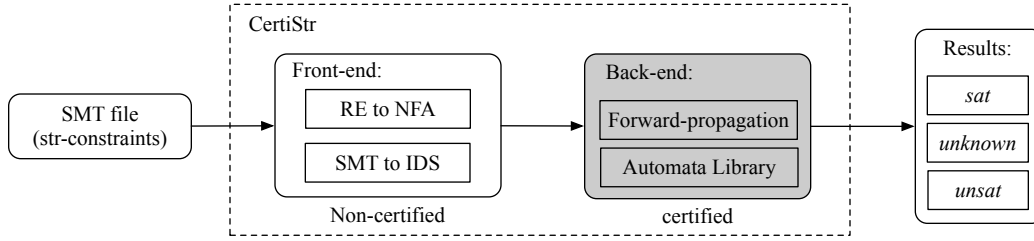


Figure 1. The framework of the certified string solver

will first translate all Regular Expressions (RE) to NFAs and desugar some string constraints that are not in the fragment, and then translate the string constraints to the Intermediate Data Structures (IDS), which are the inputs of the back-end. The IDS contains three parts: (1) a set S of variables that are used in the string constraints, (2) a concatenation constraint map $Concat$, which maps variables in S to the set of their concatenation constraints, and (3) a regular constraint map Reg , which maps each variable in S to its regular membership constraints.

The back-end contains two parts: (1) the Automata Library, which contains a collection of automata operations, such as the product and concatenation of two NFAs, and (2) a forward-propagation to check whether the string constraints are satisfiable. The results of the certified string solver can be (1) *sat* the string constraint is satisfiable, (2) *unsat* the string constraint is unsatisfiable, and (3) *unknown* the solver cannot decide whether the string constraint is satisfiable or not. Note that, *unknown* does not mean that CertiStr is non-terminating. It means the results after executing the forward-propagation for string constraints are not sufficient to decide their satisfiability.

4 Forward-propagation of String Constraints

In this section, we present the algorithm of the forward-propagation and the properties proven for it.

4.1 The Algorithm of Forward-propagation

We use the following example of string constraints to illustrate our forward-propagation analysis.

Example 4.1. $x_1 \in \mathcal{A}_1 \wedge x_2 \in \mathcal{A}_2 \wedge x_3 \in \mathcal{A}_3 \wedge x_5 = x_3 + x_4 \wedge x_3 = x_1 + x_2$.

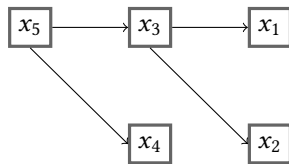


Figure 2. Dependence graph of variables in Example 4.1

Example 4.1 contains five variables “ x_1, x_2, x_3, x_4, x_5 ” and two concatenation constraints “ $x_5 = x_3 + x_4 \wedge x_3 = x_1 + x_2$ ”. We can view these concatenation constraints as a dependence graph shown in Figure 2. For a concatenation constraint $x_k = x_i + x_j$, it yields two dependence relations: $x_k \mapsto x_i$ and $x_k \mapsto x_j$ ($a \mapsto b$ denotes a depends on b). This means that before propagating the concatenation of the regular constraints of x_i and x_j to x_k , we need to first compute the propagation to x_i and x_j . Let the initial regular constraint of x_k be \mathcal{A}_k . Let the regular constraints of x_i and x_j be \mathcal{A}_i and \mathcal{A}_j , respectively. The new regular constraint of x_k after the propagation can be refined to an automaton, whose language is $\mathcal{L}(\mathcal{A}_k) \cap (\mathcal{L}(\mathcal{A}_i) + \mathcal{L}(\mathcal{A}_j))$, where $\mathcal{L}_1 + \mathcal{L}_2 \triangleq \{w_1w_2 \mid w_1 \in \mathcal{L}_1 \wedge w_2 \in \mathcal{L}_2\}$. Here, the term “refine” means narrowing the regular constraint of a variable by propagating the concatenation constraints.

The forward-propagation repeats the following computation until all variables are refined:

- detect a set of variables such that for each variable in the set, all its dependence variables have already been refined.
- refine the regular constraints of the variables detected in the last step with respect to their concatenation constraints.

We illustrate the idea of the forward-propagation with Example 4.1. The first iteration of the forward-propagation will detect the set of variables $\{x_1, x_2, x_4\}$ since they have no dependence variables and we do not need to refine their regular constraints. The second iteration of the forward-propagation will detect the set of variables: $\{x_3\}$, as x_3 only depends on the variables in $\{x_1, x_2, x_4\}$. Its regular constraint will be refined to an automaton, whose language is $\mathcal{L}(\mathcal{A}_3) \cap (\mathcal{L}(\mathcal{A}_1) + \mathcal{L}(\mathcal{A}_2))$ after the propagation for the constraint $x_3 = x_1 + x_2$. Assume this new regular constraint as \mathcal{A}'_3 . The third iteration detects the variable x_5 , whose regular constraint will be refined to an automaton, whose language is $\Sigma^* \cap (\mathcal{L}(\mathcal{A}'_3) + \Sigma^*)$. Note that we do not specify the initial regular constraints of x_4 and x_5 in the string constraint and therefore their initial regular constraints are Σ^* by default.

Algorithm 1 shows the forward-propagation algorithm. It contains three procedures: **Forward_Prop**, **Ready_Set**, and **Var_Lang**. The procedure **Forward_Prop** is the entry point

of the algorithm and has three parameters: (1) S is the set of variables used in the string constraints, for instance, in Example 4.1, there are five variables: $S = \{x_1, x_2, x_3, x_4, x_5\}$; (2) $Concat : S \rightarrow 2^{(S \times S)}$ is a partial map which maps a variable x to a set of pairs of variables. For each pair (x_1, x_2) in the set, there is a concatenation constraint: $x = x_1 + x_2$; (3) $Reg : S \rightarrow \mathcal{N}$ is a map from variables to automata. These automata denote the variables' initial regular constraints. Here \mathcal{N} denotes the set of NFAs.

Algorithm 1 The algorithm of forward-propagation

```

1: procedure Forward_Prop( $S, Concat, Reg$ )
2:    $R \leftarrow \emptyset$ 
3:   while  $S \neq \emptyset$  do
4:      $C \leftarrow \text{Ready\_Set}(S, Concat, R)$ 
5:      $Reg \leftarrow \text{Var\_Lang}(C, Concat, Reg)$ 
6:      $S \leftarrow S - C; R \leftarrow R \cup C$ 
7:   end while
8:   return  $Reg$ 
9: end procedure
10: procedure Var_Lang( $C, Concat, Reg$ )
11:   for each  $v \in C$  do
12:      $\mathcal{A} \leftarrow (Reg\ v)$ 
13:     for each  $(v_1, v_2) \in (Concat\ v)$  do
14:        $\mathcal{A}' \leftarrow \text{NFA\_concat}(Reg\ v_1)\ (Reg\ v_2)$ 
15:        $\mathcal{A} \leftarrow \text{NFA\_product}\ \mathcal{A}\ \mathcal{A}'$ 
16:     end for
17:      $Reg \leftarrow Reg\ [v \mapsto \mathcal{A}]$ 
18:   end for
19:   return  $Reg$ 
20: end procedure
22: procedure Ready_Set( $S, Concat, R$ )
23:    $C \leftarrow \emptyset$ 
24:   for each  $v \in S$  do
25:      $D \leftarrow \emptyset$ 
26:     for each  $(v_1, v_2) \in (Concat\ v)$  do
27:        $D \leftarrow D \cup \{v_1, v_2\}$ 
28:     end for
29:     if  $D \subseteq R$  then
30:        $C \leftarrow C \cup \{v\}$ 
31:     end if
32:   end for
33:   return  $C$ 
34: end procedure

```

In the while loop of the procedure **Forward_Prop**, the computation iteratively refines the regular constraints of variables by propagating the regular constraints of variables on the right hand side of the concatenation constraints. The first step in the loop computes the set C of variables whose dependence variables are all in the set R (computed by the procedure **Ready_Set**). R is initialized empty and it denotes

the set of variables that have already been refined by the previous loops. The second step in the loop updates the regular constraints of the variables in C stored in Reg by propagating the concatenation constraints in $Concat$ (computed by **Var_Lang**).

At the end of the loop body, the set S is updated by removing the variables in C as they have already been refined and these removed variables are added to the set R , i.e., they have already been refined.

The procedure **Ready_Set** computes the set of variables that are only dependent on the variables in R . For each variable v , it stores all its dependence variables in D and then checks whether D is a subset of R . If D is a subset of R then v is moved to C as all its dependence variables have been refined.

The procedure **Var_Lang** refines the regular constraints of variables in C . In the inner loop of **Var_Lang**, it traverses all concatenation constraints $v = v_1 + v_2$ and refines the regular constraint of v by propagating the concatenation of the regular constraints of v_1 and v_2 . $Reg\ [v \mapsto \mathcal{A}]$ means updating the regular constraint of v in Reg to \mathcal{A} . The functions **NFA_concat** and **NFA_product** denote constructing the concatenation and product of two NFAs, respectively. These two functions will be introduced in Section 5.

Assume a string constraint is stored in S , $Concat$ and Reg , then after calling "**Forward_Prop**($S, Concat, Reg$)", we will get a new Reg' such that for each variable v , $(Reg'\ v)$ stores the refined regular constraint of v , which is an NFA. We have two cases to decide whether the string constraint is satisfiable or not:

- If there exists a variable v such that the language of $(Reg'\ v)$ is empty then the string constraint is unsatisfiable, because the refined regular constraint of v is the empty language.
- If for all variables v such that the language of $(Reg'\ v)$ is not empty then unfortunately we cannot decide whether the string constraint is satisfiable or not. Consider the string constraint: "(1) $y = x + x \wedge$ (2) $y \in \{ab\} \wedge$ (3) $x \in \{a, b\}$ ", where ab , a , and b are words, the refined regular constraints of x and y are both not empty. But this string constraint is not satisfiable as constraints (1) and (2) require x to have two different values. We will introduce a property, called *tree property* later, such that if the string constraint satisfies the *tree property* then it is satisfiable, otherwise CertiStr returns *unknown*.

4.2 Theorems for Forward-propagation

Now we introduce the key theorem proven for the algorithm. Firstly we present the definitions of well-formed inputs (**wf** in Figure 3) and the acyclic property (**acyclic** in Figure 3) for concatenation constraints, which are the premises of the correctness theorem.

definition wf where

$$\begin{aligned} \text{wf } S \text{ Concat } Reg &= \\ &(\text{dom } Concat \subseteq S) \wedge (\text{dom } Reg = S) \wedge \\ &(\forall v \ v_1 \ v_2. \ v \in (\text{dom } Concat) \wedge (v_1, v_2) \in (Concat \ v) \\ &\quad \longrightarrow v_1 \in S \wedge v_2 \in S) \wedge \text{finite } S \end{aligned}$$

definition acyclic where

$$\begin{aligned} \text{acyclic } S \text{ Concat } l &= S = \bigcup(\text{set } l) \wedge ((l = []) \\ &\vee (l = s\#l' \longrightarrow (s \cap (\bigcup(\text{set } l')) = \emptyset) \wedge \\ &(\forall v_1 \ v_2 \ v. \ v \in s \wedge (v_1, v_2) \in (Concat \ v) \longrightarrow \\ &\{v_1, v_2\} \subseteq \bigcup(\text{set } l')) \wedge (\text{acyclic } (S - s) \text{ Concat } l'))) \end{aligned}$$

Figure 3. The definitions of **wf** and **acyclic**

The **wf** predicate requires that all variables appearing in the domain and codomain of *Concat* should be in the set *S*, and the domain of *Reg* should be the set *S*, i.e., all variables should have their initial regular constraints. *S* should be finite. Now we consider the definition of **acyclic**, which checks whether the dependence graphs of concatenation constraints are acyclic. **Forward_Prop** cannot solve string constraints that are not acyclic. For instance, the string constraint “ $x = y_1 + y_2 \wedge y_1 = z_1 + x$ ” does not satisfy the acyclic property, as $x \mapsto y_1 \mapsto x$.

In order to specify the acyclic property, we arrange the set *S* of variables to a list in which elements are disjoint sets of variables in *S* and the list characterizes the dependence levels among variables. For instance, in Example 4.1, there are 5 variables and they can be organized in a list: $[\{v_5\}, \{v_3\}, \{v_1, v_2, v_4\}]$. Indeed, the list characterizes the dependence levels among variables. For instance, the variable v_5 only depends on the variables in $\{v_3\}$ and $\{v_1, v_2, v_4\}$ and v_3 only depends on the variables in $\{v_1, v_2, v_4\}$ but not $\{v_5\}$. The variables in $\{v_1, v_2, v_4\}$ do not depend on any variables. If *Concat* does not satisfies the acyclic property then we cannot arrange all variables in such a list, as all variables in a cycle always depend on another variable in the cycle.

The predicate **acyclic** contains two parts. (1) $S = \bigcup(\text{set } l)$, which means the elements in *S* and *l* should be exactly the same. “**set** *l*” translates the list *l* to a set of elements in *l* and “ $\bigcup(\text{set } l) = \{x \mid \exists y. y \in \text{set } l \wedge x \in y\}$ ”. (2) Part 2 further has two cases. The first case is “ $l = []$ ”, i.e., empty list. In this case, the predicate is obviously true. The second case is “ $l = s\#l'$ ”, i.e., it has at least one element. *s* is the first element and *l'* is the tail. In this case, firstly, it requires the elements in *s* to be disjoint with the elements in the sets of the tail *l'* (checked by $s \cap (\bigcup(\text{set } l')) = \emptyset$). Secondly, let *v* be a variable in *s*, for any two variables v_1 and v_2 that *v* depends on, v_1 and v_2 can only be in the sets of the tail *l'*. Thirdly, **acyclic** $(S - s)$ *Concat* *l'* checks whether the remaining elements in $S - s$ are still acyclic w.r.t. *Concat* and *l'*. In order to check whether the

variables in *S* are acyclic w.r.t. *Concat*, we can specify it as “ $\exists l. \text{acyclic } S \text{ Concat } l$ ”.

Now we are ready to introduce our correctness theorem (Theorem 4.2). The keyword **fixes** in Isabelle is used to define universally quantified variables. The keyword **assumes** is used to specify assumptions of the theorem and the keyword **shows** is used to specify the conclusion of the theorem. The **shows** part is of the form “ $M \leq \text{SPEC } (\lambda \text{Reg}'. \Phi)$ ”, which means the algorithm *M* is correct with respect to the specification Φ . More precisely, the result of applying the predicate $(\lambda \text{Reg}'. \Phi)$ to the output of *M* should be true. “**NFA_accept** *w* \mathcal{A} ” checks whether *w* is accepted by \mathcal{A} , which will be introduced in Section 5. The symbol “@” is the list concatenation operation in Isabelle. We use a list of elements in Σ to represent a word.

Theorem 4.2 (Correctness of **Forward_Prop**).

$$\begin{aligned} &\text{fixes } S \text{ Concat } Reg \\ &\text{assumes } 1. \text{wf } S \text{ Concat } Reg \\ &\quad 2. (\exists l. \text{acyclic } S \text{ Concat } l) \\ &\text{shows } \text{Forward_Prop}(S, \text{Concat}, \text{Reg}) \leq \text{SPEC}(\lambda \text{Reg}'. \\ &\quad (\forall w \ v. \text{NFA_accept } w \ (\text{Reg}' \ v) \longleftrightarrow \\ &\quad (\text{NFA_accept } w \ (\text{Reg } v) \wedge \\ &\quad (\forall v_1 \ v_2. (v_1, v_2) \in (\text{Concat } v) \longrightarrow \\ &\quad (\exists w_1 \ w_2. w = w_1 @ w_2 \wedge \\ &\quad \text{NFA_accept } w_1 \ (\text{Reg}' \ v_1) \wedge \\ &\quad \text{NFA_accept } w_2 \ (\text{Reg}' \ v_2)))))) \end{aligned}$$

Theorem 4.2 shows that if *S*, *Concat*, *Reg* are well-formed and acyclic then after calling **Forward_Prop**(*S*, *Concat*, *Reg*) we can get *Reg'*, which satisfies: any *w* is accepted by the automaton $(\text{Reg}' \ v)$ if and only if

1. *w* is accepted by the original regular constraint of *v*, i.e., $(\text{Reg } v)$, and
2. for all variables v_1 and v_2 , $(v_1, v_2) \in \text{Concat } v$ implies there exist w_1, w_2 such that $w = w_1 @ w_2$ and w_i , for $i = 1, 2$, is accepted by the regular constraint $(\text{Reg}' \ v_i)$.

That is, in the new *Reg'*, the concatenation constraints are used to refine the corresponding variables.

Moreover, in Isabelle, Theorem 4.2 also requires us to prove the termination of Algorithm 1. The termination is proven by showing that in the procedure **Forward_Prop**, the number of variables in *S* decreases progressively for each iteration. Since *S* is finite, the algorithm must terminate.

Even though Theorem 4.2 is the key of the correctness of the forward-propagation, it does not give us an intuition with respect to the satisfiability semantics of string constraints introduced in Section 3. In the following, we will show that our forward-propagation is sound w.r.t. *unsat* results.

Firstly we need to formalize the semantics of *sat* and *unsat* for string constraints. Let μ be an assignment, which is a map $S \rightarrow \Sigma^*$ from variables in *S* to words. We define the predicate **sat_str** in Figure 4 for the semantics of the satisfiability of string constraints.

definition sat_str where

sat_str S $Concat$ Reg $\mu =$
 $(\forall v \in S. (\mu v) \in \mathcal{L}(Reg v)) \wedge$
 $(\forall v v_1 v_2. v \in S \wedge (v_1, v_2) \in (Concat v) \longrightarrow$
 $(\mu v) = (\mu v_1) @ (\mu v_2))$

definition tree where

tree $Concat =$
 $(\forall v v_1 v_2. (v_1, v_2) \in Concat v \longrightarrow v_1 \neq v_2) \wedge$
 $(\forall v v_1 v_2 v_3 v_4. (v_1, v_2) \neq (v_3, v_4) \wedge$
 $(v_1, v_2) \in Concat v \wedge (v_3, v_4) \in Concat v$
 $\longrightarrow \{v_1, v_2\} \cap \{v_3, v_4\} = \emptyset) \wedge$
 $(\forall v v' v_1 v_2 v_3 v_4. v \neq v' \wedge (v_1, v_2) \in Concat v$
 $\wedge (v_3, v_4) \in Concat v' \longrightarrow \{v_1, v_2\} \cap \{v_3, v_4\} = \emptyset)$

Figure 4. The definitions of `sat_str` and `tree`

The predicate `sat_str` contains two parts: (1) the assignment μ should respect all regular constraints and (2) the assignment μ should respect all concatenation constraints.

Theorem 4.3 shows that after executing `Forward_Prop`, we get a new map Reg' . If there exists a variable v , such that $\mathcal{L}(Reg' v) = \emptyset$, i.e., the language of the refined regular constraint of v is the empty set, then there is no assignment μ that makes the predicate “`sat_str S Concat Reg μ` ” true.

Theorem 4.3 (Soundness for *unsat* results).

fixes S $Concat$ Reg
assumes 1. **wf** S $Concat$ Reg
 2. $(\exists l. \text{acyclic } S \text{ Concat } l)$
shows $\text{Forward_Prop}(S, Concat, Reg) \leq \text{SPEC}(\lambda Reg')$
 $((\exists v. \mathcal{L}(Reg' v) = \emptyset) \longrightarrow$
 $(\forall \mu. \neg \text{sat_str } S \text{ Concat } Reg \mu))$

Unfortunately, the theorem holds only for *unsat* results. As we already analyzed before, the case $\forall v. \mathcal{L}(Reg' v) \neq \emptyset$ cannot make us conclude that the string constraint is satisfiable.

In this paper, we characterize a subset of the string constraints, using a property such that the forward-propagation’s *satisfiable* results are also sound. We call it the *tree property*. The intuition behind the *tree property* is that the dependence graph of a string constraint should constitute a tree or forest. The example $y = x + x$ does not satisfy the *tree property*, because y has two edges pointing to a same node x in the dependence graph. In other words, the node x has indegree 2. But a tree requires that each node has at most indegree 1. As the map `Concat` stores the dependence graphs of string constraints, we only need to check `Concat` for the tree property. The predicate `tree` is defined in Figure 4.

This definition has three cases separated by the conjunction operator: (1) the first case requires that for any concatenation constraint $v = v_1 + v_2$, we have $v_1 \neq v_2$. (2) The second

case requires that for any variable v with two concatenations $v = v_1 + v_2$ and $v = v_3 + v_4$, $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are disjoint. (3) The third case requires that for any two different variables v and v' with the concatenation constraints $v = v_1 + v_2$ and $v' = v_3 + v_4$, $\{v_1, v_2\}$ and $\{v_3, v_4\}$ are disjoint.

With this definition we have the following theorem.

Theorem 4.4 (Completeness for the tree property).

fixes S $Concat$ Reg
assumes 1. **wf** S $Concat$ Reg
 2. $(\exists l. \text{acyclic } S \text{ Concat } l)$
 3. **tree** $Concat$
shows $\text{Forward_Prop}(S, Concat, Reg) \leq \text{SPEC}(\lambda Reg')$
 $((\forall v. \mathcal{L}(Reg' v) \neq \emptyset) \longleftrightarrow$
 $(\exists \mu. \text{sat_str } S \text{ Concat } Reg \mu))$

This theorem shows that if `Concat` satisfies the tree property then we have: Reg' satisfies that for all v , the language of $Reg' v$ is not empty iff there exists an assignment, which makes the string constraint satisfiable.

In order to check whether a string constraint satisfies the tree property, we propose Algorithm 2. It first stores all the variables, which are on the right hand side of the concatenation constraints, into a list, and then checks whether the list is distinct. If it is not distinct then there must exist a variable with at least indegree 2 and thus `Concat` cannot satisfy the tree property.

Algorithm 2 Tree property checking

```

1: procedure Check_Tree( $S, Concat$ )
2:    $l \leftarrow []$ 
3:   for each  $v \in S$  do
4:     for each  $(v_1, v_2) \in (Concat v)$  do
5:        $l \leftarrow v_1 \# v_2 \# l$ 
6:     end for
7:   end for
8:   return distinct  $l$ 
9: end procedure

```

Now we show that if Algorithm 2 returns *True* then the tree property is satisfied.

Theorem 4.5 (Correctness of `Check_Tree`).

fixes S $Concat$
assumes 1. $\text{dom } Concat \subseteq S$
 2. $(\forall v v_1 v_2. (v_1, v_2) \in Concat v \longrightarrow \{v_1, v_2\} \subseteq S)$
shows $\text{Check_Tree}(S, Concat) \longrightarrow \text{tree } Concat$

With the function `Check_Tree`, CertiStr can return three results: *unsat*, *sat*, and *unknown*, in which *unknown* means the string constraint does not satisfy the tree property and the forward-propagation does not refine the regular constraint of any variable to empty.

We now finish introducing the forward-propagation procedure. As it uses many NFA operations, in the following

section, we will introduce our implementation for these NFA operations.

5 Symbolic Automata Formalization

Automata operations are the key to our certified string solver. For instance, we need to check the language emptiness of an NFA as well as construct the concatenation and product of two NFAs. In order to make the string solver practically useful, we should also keep the efficiency in mind.

As argued by D’Antoni et. al [15], the classical definition of NFA, whose transition labels are a character in the alphabet, is not suitable for practical implementation, and they propose s-NFAs, which allow transition labels to carry a set of characters, to address this limitation. In this section, we present our formalization of s-NFAs in Isabelle.

Definition 5.1 (Abstract s-NFAs). An s-NFA is a 5-tuple: $\mathcal{A} = (Q, \Sigma, \Delta, I, F)$, where Q is a finite set of states, Σ is an alphabet (Σ may be infinite), $\Delta \subseteq Q \times 2^\Sigma \times Q$ is a finite set of transition relations, $I \subseteq Q$ is a set of initial states, and $F \subseteq Q$ is a set of accepting states.

For a transition (q, α, q') in an s-NFA, the implementation of α can use various different data structures, such as intervals for string solvers and Binary Decision Diagrams (BDD) for symbolic model checkers. Definition 5.1 of s-NFAs is called abstract s-NFAs because the transition labels are denoted as sets instead of using concrete data structures, like intervals and BDDs.

In order to make our formalization of s-NFAs reusable for different implementation of transition labels, we exploit Isabelle’s refinement framework and divide the formalization of s-NFAs in two levels: (1) the abstract level (introduced in this section) and (2) the implementation level (introduced in the following section). When the correctness of the algorithms at the abstract level is proven, the implementation level does not need to re-prove it, only the refinement relations between these two levels must be shown.

At the abstract level, transition labels are modeled as a set as shown in Definition 5.1. At the implementation level, the sets of characters are refined to concrete data structures, such as intervals or BDDs.

Our formalization of s-NFAs extends an existing classical NFA formalization developed by Tuerk et al. [40]; this formalization is part of the verified CAVA model checker [26]. We extend it to support s-NFAs and further operations, such as the concatenation operation of two NFAs. Figure 5 shows the Isabelle formalization of s-NFAs. An s-NFA is defined by a record with four elements: (1) Q is the set of states, (2) I is the set of initial states, (3) F is the set of accepting states, and (4) Δ is the set of transitions with transition labels defined by sets. The definition **well-formed** is a predicate to check whether an NFA satisfies (1) the states occurring in the sets of initial and accepting states are also in the set $(Q \ \mathcal{A})$ of states. In Isabelle, the value of a field `fd` in a record `rd` can be

```
record NFA = Q :: "'q set"
           Δ :: "('q*'a set*'q) set"
           I :: "'q set"
           F :: "'q set"
```

definition well-formed where

```
well-formed  $\mathcal{A} = (I \ \mathcal{A}) \subseteq (Q \ \mathcal{A}) \wedge (F \ \mathcal{A}) \subseteq (Q \ \mathcal{A})$ 
            $\wedge$  finite  $(Q \ \mathcal{A}) \wedge$  finite  $\Delta \wedge$ 
            $(\forall (q, \alpha, q') \in \Delta, q \in (Q \ \mathcal{A}) \wedge q' \in (Q \ \mathcal{A}))$ 
```

fun reachable w Δ q q' where

```
reachable [] Δ q q' = (q=q')
| reachable (a # w) Δ q q' =
   $(\exists q_i \alpha. (q, \alpha, q_i) \in \Delta$ 
     $\wedge a \in \alpha \wedge$  reachable w Δ  $q_i \ q')$ 
```

definition NFA_accept where

```
NFA_accept w  $\mathcal{A} = (\exists q \ q'. q \in (I \ \mathcal{A}) \wedge q' \in (F \ \mathcal{A})$ 
            $\wedge$  reachable w  $(\Delta \ \mathcal{A}) \ q \ q')$ 
```

definition \mathcal{L} where

```
 $\mathcal{L} \ \mathcal{A} = \{w. \text{accept } w \ \mathcal{A}\}$ 
```

Figure 5. Isabelle formalization of abstract symbolic NFA

extracted by `(fd rd)`. Therefore, $(Q \ \mathcal{A})$ denotes the value of the field Q in \mathcal{A} . (2) The states in the set of transitions should also be in $(Q \ \mathcal{A})$. (3) $(Q \ \mathcal{A})$ is a finite set. (4) Δ is finite.

For a word $w = a_1 a_2 \dots a_n$ and two states q and q' , the function **reachable** checks whether there exists a path $q_0, \alpha_1, q_1, \dots, q_{n-1}, \alpha_n, q_n$ in the NFA \mathcal{A} such that $q = q_0 \wedge q' = q_n$, $(q_{i-1}, \alpha_i, q_i) \in (\Delta \ \mathcal{A})$, $1 \leq i \leq n$ and $a_j \in \alpha_j$, $1 \leq j \leq n$. The definition **NFA_accept** checks whether a word is accepted by an NFA and the definition of \mathcal{L} is the language of an NFA.

Based on this definition of NFAs, a collection of automata operations are defined and their correctness lemmas are proven. Here we only present the operation of the concatenation shown in Figure 6.

NFA_concat_basic constructs the concatenation of two NFAs. The correctness of this definition relies on the fact that the sets of states in the two NFAs are disjoint. In the definition of **NFA_concat_basic**, the set of transitions of the concatenation is the union of (1) $(\Delta \ \mathcal{A}_1)$, (2) $(\Delta \ \mathcal{A}_2)$, and (3) $\{(q, \alpha, q'') \mid \exists q'. (q, \alpha, q') \in (\Delta \ \mathcal{A}_1) \wedge q' \in (F \ \mathcal{A}_1) \wedge q'' \in (I \ \mathcal{A}_2)\}$. The transitions in (3) concatenate the language of \mathcal{A}_1 to the language of \mathcal{A}_2 .

The set of the initial states in the concatenation is computed by first checking whether there is a state that is in both the sets of initial and accepting states of \mathcal{A}_1 . If there exists such a state then the initial states in the concatenation should include the initial states of \mathcal{A}_2 , as \mathcal{A}_1 accepts the

definition `NFA_concat_basic` **where**
`NFA_concat_basic` $\mathcal{A}_1 \mathcal{A}_2 \equiv$ (
 $Q = (Q \mathcal{A}_1) \cup (Q \mathcal{A}_2),$
 $\Delta = (\Delta \mathcal{A}_1) \cup (\Delta \mathcal{A}_2) \cup$
 $\{(q, \alpha, q'') \mid \exists q'. (q, \alpha, q') \in (\Delta \mathcal{A}_1) \wedge$
 $q' \in (F \mathcal{A}_1) \wedge q'' \in (I \mathcal{A}_2)\},$
 $I = \text{if } ((I \mathcal{A}_1) \cap (F \mathcal{A}_1) = \emptyset) \text{ then } (I \mathcal{A}_1)$
 $\text{ else } ((I \mathcal{A}_1) \cup (I \mathcal{A}_2)),$
 $F = (F \mathcal{A}_2)$
 $)$

definition `NFA_concat` **where**
`NFA_concat` $\mathcal{A}_1 \mathcal{A}_2 f_1 f_2 =$
`remove_unreachable_states`
 $(\text{NFA_concat_basic}$
 $(\text{NFA_rename } f_1 \mathcal{A}_1)(\text{NFA_rename } f_2 \mathcal{A}_2))$

Figure 6. The definition of concatenation operation

empty word, otherwise the initial states of the concatenation are \mathcal{A}_1 's initial states.

The definition of `NFA_concat` firstly renames the states in both NFAs to ensure their sets of states disjoint. In the term “`NFA_rename f A`”, f is a renaming function, for instance, f can be “ $\lambda q.(q, 1)$ ”, which renames any state q to $(q, 1)$. “`NFA_rename f A`” renames all states in $(Q \mathcal{A})$ by the function f . Correspondingly, all states in transitions, initial states, and accepting states are also renamed by f . Renaming states makes it easier to ensure the sets of states of two NFAs are disjoint. The function `remove_unreachable_states` removes all states in an NFA that are not reachable from the initial states of the NFA. This enables us to check the language emptiness of an NFA by only checking the emptiness of its set of accepting states. At the implementation level, the algorithm also only generates reachable states for NFAs. But with `remove_unreachable_states` used in `NFA_concat` at the abstract level, the refinement relation between these two levels can be specified as the NFA *isomorphism* of the two output concatenations of automata at the two levels, instead of language equivalence, which is more challenge to prove.

In order to ensure the correctness of this operation, we need to prove the following lemma in Isabelle.

Lemma 5.2 (Correctness of `NFA_concat`).
fixes $\mathcal{A}_1 \mathcal{A}_2 f_1 f_2$
assumes 1. **well-formed** \mathcal{A}_1
2. **well-formed** \mathcal{A}_2
3. $(\text{image } f_1 (Q \mathcal{A}_1)) \cap (\text{image } f_2 (Q \mathcal{A}_2)) = \emptyset$
4. $\text{inj_on } f_1 (Q \mathcal{A}_1) \wedge \text{inj_on } f_2 (Q \mathcal{A}_2)$
shows $\mathcal{L}(\text{NFA_concat } \mathcal{A}_1 \mathcal{A}_2 f_1 f_2) =$
 $\{w_1 @ w_2. w_1 \in \mathcal{L}(\mathcal{A}_1) \wedge w_2 \in \mathcal{L}(\mathcal{A}_2)\}$

The assumptions **well-formed** \mathcal{A}_1 and **well-formed** \mathcal{A}_2 require the input NFAs to be well-formed. In addition, the assumption $(\text{image } f_1 (Q \mathcal{A}_1)) \cap (\text{image } f_2 (Q \mathcal{A}_2)) = \emptyset$ requires that after renaming the states in both \mathcal{A}_1 and \mathcal{A}_2 , the new sets of states of the two NFAs are disjoint. The term “**image f S**” applies the function f to the elements in the set S and returns a new set $\{f e. e \in S\}$. The premises **inj_on** $f_1 (Q \mathcal{A}_1)$ and **inj_on** $f_2 (Q \mathcal{A}_2)$ require that the functions f_1 and f_2 are injective over the sets $(Q \mathcal{A}_1)$ and $(Q \mathcal{A}_2)$, respectively. The conclusion in “**shows**” specifies that the language of the concatenation of \mathcal{A}_1 and \mathcal{A}_2 equals the set of the concatenations of the words in $\mathcal{L}(\mathcal{A}_1)$ and the words in $\mathcal{L}(\mathcal{A}_2)$.

Moreover, some other important functions are defined, for instance:

- The product of two NFAs: `NFA_product` $\mathcal{A}_1 \mathcal{A}_2$, which computes the product of \mathcal{A}_1 and \mathcal{A}_2 . We proved the theorem: $\mathcal{L}(\text{NFA_product } \mathcal{A}_1 \mathcal{A}_2) = \mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$.
- Checking the isomorphism of two NFAs: `NFA_isomorphism` $\mathcal{A}_1 \mathcal{A}_2 \triangleq (\exists f. \text{inj_on } f (Q \mathcal{A}_1) \wedge \text{NFA_rename } f \mathcal{A}_1 = \mathcal{A}_2)$. Some lemmas are proven for it, such as, **well-formed** $\mathcal{A}_1 \wedge \text{well-formed } \mathcal{A}_2 \wedge \text{NFA_isomorphism } \mathcal{A}_1 \mathcal{A}_2 \longrightarrow \mathcal{L}(\mathcal{A}_1) = \mathcal{L}(\mathcal{A}_2)$.

Section 4 and this section present the abstract level algorithms for the forward-propagation and the NFA operations, respectively. In order to make CertiStr efficient and practically useful, we need to use some efficiently implemented data structures, such as red-black-trees (RBT) and hash maps. This will be introduced in the following section.

6 Implementation-Level Algorithms

In this section, we will first present the relations between abstract concepts, like sets and maps, and the implementation data structures in Isabelle (Subsection 6.1), which is the prerequisite to understand the implementation-level algorithm in Subsection 6.2.

6.1 Isabelle Collections Framework

Isabelle Collections Framework (ICF) [27] provides an efficient, extensible, and machine checked collections framework. The framework features the use of data refinement techniques [25] to refine an abstract specification (using high-level concepts like sets) to a more concrete implementation (using collection data structures, like RBT and hashmaps). The code-generator of Isabelle can be used to generate efficient code.

The concrete data structures implement a collection of interfaces that mimic the operations of abstract concepts. We list the following RBT implementation of set interfaces that will be used to present the implementation-level algorithms:

- `RBT.empty`, generates an RBT representation of the empty set.
- `RBT.mem e D`, checks whether the element e is in the set represented by RBT D .

```

record NFA_rbt =
  Q :: "'q RBT"
  Δ :: "('q * 'a Interval * 'q) RBT"
  I :: "'q RBT"
  F :: "'q RBT"

```

Figure 7. NFA data structure definition using RBT

- **RBT.insert** $e D$, inserts the element e into the set represented by RBT D .
- **RBT.inter** $D_1 D_2$, computes the RBT representation of the set intersection of D_1 and D_2 . D_1 and D_2 are both RBT represented sets.
- **RBT.union** $D_1 D_2$, computes the RBT representation of the set union of D_1 and D_2 . D_1 and D_2 are both RBT represented sets.
- **RBT.iterator** $D v f$, where f is of the type $t_1 \Rightarrow t_2 \Rightarrow t_2, v$ is of the type t_2 , and the elements in D are of the type t_1 . This interface implementation iteratively applies the function f to the elements in D with the initial value v . For instance, the first iteration of the computation of “**RBT.iterator** $\{e_1, e_2\} v f$ ” selects an element from $\{e_1, e_2\}$, assume e_1 is selected, and then applies f to e_1 and v , i.e., “ $f e_1 v$ ”. Assume the value of “ $f e_1 v$ ” is v' . The second iteration applies f to e_2 and v' . Now all elements in the set are traversed, “**RBT.iterator** $\{e_1, e_2\} v f$ ” finishes with the return value of “ $f e_2 v'$ ”.
- **RBT.to_list** D translates a set represented by D to a list.
- **RBT.α** D translates RBT represented set D to the set in the abstract concept.

6.2 Implementation-Level Algorithm for NFA_concat

In this subsection, we present the implementation-level algorithm by the NFA operation **NFA_concat**.

In order to formalize implementation-level algorithms, the first step is to consider the data refinement. At the abstract level, we use sets to store states, transitions, initial states, and accepting states. At the implementation level, we use RBTs to replace sets to store these elements of NFAs. The type of NFAs using RBT data structures is specified in Figure 7. Note that, here we only give a simplified version of the RBT NFA definition to make it easy to understand. Indeed, RBT has more type arguments.

Moreover, for each transition (q, α, q') , the label α is represented by a set at the abstract level. We use intervals to replace sets to store α . An interval $[n_1, n_2]$ is a pair of elements in a totally ordered set. The semantics of an interval is defined by a set as **semi** $[n_1, n_2] = \{n \mid n_1 \leq n \leq n_2\}$.

In order to support NFA operation implementation, some operations for intervals are implemented. We list the following 3 interval operations here.

- Computing the intersection of two intervals: **intersectionI** $[I_1, I_2] [I'_1, I'_2] \triangleq [\max(I_1, I'_1), \min(I_2, I'_2)]$. The term $\max(I_1, I'_1)$ returns the bigger one of I_1 and I'_1 , and $\min(I_2, I'_2)$ returns the smaller one of I_2 and I'_2 .
- Checking the non-emptiness of the interval $[I_1, I_2]$: **emptyI** $[I_1, I_2] \triangleq I_1 \leq I_2$.
- Checking the membership of an element e and the interval $[l_1, l_2]$: **memI** $e [l_1, l_2] \triangleq l_1 \leq e \leq l_2$.

Algorithm 3 The algorithm idea of NFA concatenation

```

1: procedure Construct_Trans( $S_{\Delta 1}, S_{\Delta 2}, S_{I_2}, S_{F_1}$ )
2:    $S_D \leftarrow$  RBT.union  $S_{\Delta 1} S_{\Delta 2}$ ;
3:   RBT.iterator  $S_{\Delta 1} S_D$ 
4:      $(\lambda(q, a, q'') S.$ 
5:       if (RBT.mem  $q'' S_{F_1}$ ) then
6:         RBT.iterator  $S_{I_2} S$ 
7:            $(\lambda q' S'. \text{RBT.insert}(q, a, q') S')$ 
8:         end if);
9:   return  $S_D$ ;
10: end procedure
11:
12: procedure NFA_Concate_Impl( $\mathcal{A}_1, \mathcal{A}_2, f_1, f_2$ )
13:    $\triangleright \mathcal{A}_1, \mathcal{A}_2$  are of type NFA_rbt
14:    $\mathcal{A}'_1 \leftarrow$  NFA_Rename_Impl  $f_1 \mathcal{A}_1$ ;
15:    $\mathcal{A}'_2 \leftarrow$  NFA_Rename_Impl  $f_2 \mathcal{A}_2$ ;
16:   if RBT.inter  $(I \mathcal{A}'_1) (F \mathcal{A}'_1) \neq$  RBT.empty then
17:      $S_I \leftarrow$  RBT.union  $(I \mathcal{A}'_1) (I \mathcal{A}'_2)$ 
18:   else
19:      $S_I \leftarrow (I \mathcal{A}'_1)$ 
20:   end if;
21:    $S_Q \leftarrow S_I$ ;  $S_{\Delta} \leftarrow$  RBT.empty;  $wl \leftarrow$  RBT.to_list  $S_Q$ ;
22:    $S'_{\Delta} \leftarrow$  Construct_Trans $(\Delta \mathcal{A}'_1, \Delta \mathcal{A}'_2, I \mathcal{A}'_2, F \mathcal{A}'_1)$ ;
23:   while  $wl \neq []$  do  $\triangleright []$  denotes empty list
24:      $q_s \leftarrow wl.\text{rm\_first}$ ;
25:      $\triangleright$  “rm_first”: removes and returns the first
26:       element in  $wl$ 
27:     RBT.iterator  $S'_{\Delta} S_{\Delta}$ 
28:        $(\lambda(q, \alpha, q') S.$ 
29:         if  $q = q_s \wedge (\text{emptyI } \alpha)$  then
30:           if  $\neg(\text{RBT.mem } q' S_Q)$  then
31:             RBT.insert  $q' S_Q$ ;
32:              $wl \leftarrow wl@[q']$ ;
33:           end if;
34:           RBT.insert  $(q, \alpha, q') S$ ;
35:         end if;
36:       )
37:   end while
38:   return  $(S_Q, S_{\Delta}, S_I, \text{RBT.inter } S_Q (F \mathcal{A}'_2))$ 
39: end procedure

```

definition `NFA_α` **where**

$$\text{NFA}_\alpha \mathcal{A} \equiv (\mid Q = \text{RBT}.\alpha (Q \mathcal{A}),$$

$$\Delta = (\text{image } (\lambda(q, \alpha, q').(q, \text{semI } \alpha, q'))$$

$$\quad (\text{RBT}.\alpha (\Delta \mathcal{A}))),$$

$$I = \text{RBT}.\alpha (I \mathcal{A}),$$

$$F = \text{RBT}.\alpha (F \mathcal{A}))$$

definition `refine_rel` **where**
`refine_rel` $\mathcal{A} \mathcal{A}' =$
`NFA_isomorphism` (`NFA_α` \mathcal{A}') \mathcal{A}

lemma `refine_rel_lang`:
`well-formed` $\mathcal{A} \wedge \text{refine_rel } \mathcal{A} \mathcal{A}'$
 $\implies \mathcal{L}(\mathcal{A}) = \mathcal{L}(\text{NFA}_\alpha \mathcal{A}')$

Figure 8. Refinement relation formalization

Algorithm 3 shows the basic idea of the `NFA_concat` implementation (the procedure `NFA_Concat_Impl`). The texts after “▷” are comments. `NFA_Rename_Impl` is the implementation of the NFA renaming function `NFA_Rename` at the abstract level.

The sub-procedure `Construct_Trans` generates the set of transitions for the concatenation of the two automata, which contains the transitions in both $(\Delta \mathcal{A}_1)$ and $(\Delta \mathcal{A}_2)$, and the transitions that concatenate the two automata (cf. Figure 6 for the abstract algorithm of the concatenation).

In Algorithm 3, Line 16-20 computes the initial states of the concatenation. Line 21 initializes S_Q , S_Δ , and wl , where S_Q stores the states of the concatenation NFA, S_Δ stores the transitions of the concatenation NFA, wl is a list to mimic a queue that stores the states to be expanded. Line 23 to 37 is the while loop for expanding the reachable states and transitions. Finally Line 38 returns the concatenation of the two NFAs. This algorithm constructs the concatenation of two NFAs with only reachable states, which means that we can check the emptiness of the concatenation NFA by only checking the emptiness of its set of accepting states.

In order to ensure the correctness of the implementation for concatenation, we need to prove the refinement relation between `NFA_concat` and `NFA_Concate_Impl`.

In Figure 8, we formalize the refinement relations between the two NFAs at the abstract level and the implementation level. The function `NFA_α` translates an implementation-level NFA \mathcal{A} (its type is `NFA_rbt`) to an abstract-level NFA of the type NFA. The definition `refine_rel` defines the refinement relation between an abstract-level NFA \mathcal{A} and an implementation NFA \mathcal{A}' , which requires the isomorphism between \mathcal{A} and $(\text{NFA}_\alpha \mathcal{A}')$

Lemma 6.1 specifies that the concatenation implementation correctly refines the abstract concatenation definition.

Lemma 6.1 (Correctness of `NFA_Concate_Impl`).

fixes $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}'_1, \mathcal{A}'_2, f_1, f_2, f_3, f_4$

assumes

1. `refine_rel` $\mathcal{A}_1 \mathcal{A}'_1 \wedge \text{refine_rel } \mathcal{A}_2 \mathcal{A}'_2$
2. `inj_on` $f_1 (Q \mathcal{A}_1) \wedge \text{inj_on } f_2 (Q \mathcal{A}_2)$
3. $(\text{image } f_1 (Q \mathcal{A}_1)) \cap (\text{image } f_2 (Q \mathcal{A}_2)) = \emptyset$
4. `inj_on` $f_3 (\text{RBT}.\alpha (Q \mathcal{A}'_1)) \wedge \text{inj_on } f_4 (\text{RBT}.\alpha (Q \mathcal{A}'_2))$
5. $(\text{image } f_3 (\text{RBT}.\alpha (Q \mathcal{A}'_1))) \cap (\text{image } f_4 (\text{RBT}.\alpha (Q \mathcal{A}'_2))) = \emptyset$
6. $\mathcal{A} = \text{NFA_concat}(\mathcal{A}_1, \mathcal{A}_2, f_1, f_2)$
7. $\mathcal{A}' = \text{NFA_Concate_Impl}(\mathcal{A}'_1, \mathcal{A}'_2, f_3, f_4)$

shows `refine_rel` $\mathcal{A} \mathcal{A}'$

This theorem shows that the implementation-level concatenation of two NFAs is a refinement for the concatenation of the corresponding abstract-level NFAs. With Lemma 6.1 and Lemma `refine_rel_lang`, we can easily infer that the language of the concatenation at the abstract level is equal to the language of the concatenation at the implementation level.

In this section, we presented the implementation of the concatenation operation. This implementation uses intervals to store transition labels. Indeed, our formalization of s-NFAs can be reused to implement s-NFAs using other data structures, such as BDDs, bitvectors, and Boolean formulas, to store transition labels. The corresponding operations, such as `intersectionI`, `emptyI`, and `memI`, for these data structures need to be implemented at the implementation level accordingly.

Moreover, the forward-propagation algorithm and other automata operations also have their corresponding implementations and their refinement relations with the abstract algorithms have been proven.

7 Evaluation and Development Efforts

The tool CertiStr (<https://github.com/uuverifiers/ostrich/tree/CertiStr>) is developed in the Isabelle proof assistant (version 2020) and one can extract executable OCaml code from the formalization in Isabelle. To obtain a complete solver, we added a (non-certified) front-end for parsing SMT-LIB problems (Figure 1); this front-end is written in Scala, and mostly borrowed from OSTRICH [13]. In this section, we evaluate CertiStr against the Kaluza benchmark [35], the most commonly used benchmark to compare string solvers. Moreover, we also elaborate the development efforts for CertiStr.

7.1 The Effectiveness and Efficiency

The Kaluza benchmark contains 47284 tests, among which 38043 tests (80.4%) are in the string constraint fragment of CertiStr. CertiStr supports string constraints with word equations of concatenation, regular constraints, and monadic length. The benchmark is classified into 4 groups according to the results of the Kaluza string solver: (1) sat/small, (2) sat/big, (3) unsat/small, and (4) unsat/big.

Table 1. Experimental results of the Kaluza benchmark

	total tests	sat	unknown	unsat	solved%	avg. time(s)	timeout	CVC4
sat_small	19634	19302	332	0	98.3%	< 0.01	0	< 0.01
sat_big	774	521	253	0	67.3%	0.84	0	0.04
unsat_small	8775	7376	824	575	91%	0.58	0	< 0.01
unsat_big	8860	2502	4880	1478	45%	8.01	728	0.5
total	38043	29700	6289	2054	83.5%	1.87	728	0.11

Table 2. Experimental results for state and transition explosion

No. Concat	11	12	13	14	15	16
states	2048	4096	8192	16384	32768	65536
transitions	177147	531441	1594323	4782969	14348907	43046721
time (s)	0.43	1.29	4.10	12.37	50.59	168.28

Note that the benchmark classification is not in all cases accurate, in the sense that there are some files in unsat/big and unsat/small that are satisfiable, owing to known incorrect results that were produced by the original Kaluza string solver [28, 35].

We ran CertiStr over these 38043 tests on an AMD Opteron 2220 SE machine, running 64-bit Linux and Java 1.8 (for running the Scala front-end of CertiStr). The results are shown in Table 1. The column “total tests” denotes the total number of tests in a group. The columns “sat”, “unknown”, “unsat” denote the numbers of tests for which CertiStr returns *sat*, *unknown*, and *unsat*, respectively. The column “solved%” denotes the percentage of the tests for which the string solver returns *sat* or *unsat*. The columns “avg.time(s)” and “timeout” denote the average time for running each test in CertiStr and the number of tests that time out, respectively. The column “CVC4” denotes the average execution time of the state of the art string solver CVC4, which can efficiently solve all the string constraints without timeout. The time limit for solving each test is 60 seconds.

From Table 1, we can conclude that in total, 83.5% of the tests can be solved by CertiStr with the result *sat* or *unsat*. The remaining of the tests are unknown or timeout. The groups *sat_small* and *unsat_small* have the higher solved percentages, more than 90%, compared with *sat_big* and *unsat_big*. This is easy to understand as small tests have a higher probability to satisfy the tree property. The worst group is *unsat_big*. CertiStr can solve 45% tests in this group. The reason is that in this group, there are a lot of tests with more than 100 concatenation constraints, which significantly increases their probability of violating the tree property. For the groups *unsat_big* and *unsat_small*, CertiStr detects 2502 and 7376 tests, respectively, that are indeed satisfiable.

Moreover, in *unsat_big*, there are 728 tests that time out. After analyzing these tests, we found that a variable with a lot of concatenation constraints (i.e., the variable appears many

times on the left-hand sides of the concatenation constraints) can easily yield s-NFA state and transition explosion during the forward-propagation. In order to evaluate such a state explosion problem, we set a test of the form in Example 7.1:

Example 7.1.

$$x = x_1 + x_1 \wedge x = x_2 + x_2 \wedge x = x_3 + x_3 \wedge \dots$$

It is a test case with only the variable x on the left-hand side. On the right-hand side, there are concatenations of the form $x_i + x_i$. No regular constraints are in the test.

We ran CertiStr over the test case by increasingly adding more concatenations for the variable x . Table 2 shows the results for solving the test case. Each column contains (1) the number of concatenation constraints (*No. Concat*), (2) the size of the automaton of the variable x after executing the forward-propagation (The numbers of *states* and *transitions* of the automaton), and (3) the execution time for running the forward-propagation over the test (*time*).

From the table we can conclude that after adding 16 concatenation constraints for the variable x , the automaton generated for x contains 65536 states and 43046721 transitions. Solving this test takes 168.28s.

Now consider the 728 tests that time out. These tests have similar constraints, in which some variables have a lot of concatenation constraints. During the forward-propagation, state and transition explosion happens for these tests, therefore they cannot be solved in 60s.

Compared with CVC4, which is not automata-based and incorporates more optimizations over its string theory decision procedure, CertiStr is less efficient. We can also optimize CertiStr further in the future. For instance, we can optimize the language intersection operation $\Sigma^* \cap \mathcal{L}(\mathcal{A})$ to $\mathcal{L}(\mathcal{A})$, and the language concatenation $\Sigma^* + \Sigma^*$ to Σ^* . We can also implement NFA minimization algorithms to further improve its efficiency. These optimizations can avoid the state explosion problem to some extent. However, integrating such

optimizations in a verified solver is challenging, as more cases need to be proven correct.

7.2 The Efforts of Developing CertiStr

In this subsection, we discuss the efforts of developing CertiStr. Table 3 shows the efforts. The rows “Abs Automata Lib” and “Imp Automata Lib” denote abstract level and implementation level automata libraries, respectively. The rows “Abs Forward Prop” and “Imp Forward Prop” denote abstract level and implementation level forward-propagations, respectively. The column “Loc” denotes the lines of Isabelle code for each module. The column “Terms” denotes the number of definitions, functions, locales, and classes. The column “Theorems” denotes the number of theorems and lemmas in a module. The development needs around one person-year efforts.

Table 3. Development effort of certified string solver

Module	Loc	Terms	Theorems
Abs Automata Lib	4484	100	254
Imp Automata Lib	8498	270	203
Abs Forward Prop	6850	37	39
Imp Forward Prop	2175	41	18
Total	22007	448	514

We now discuss the challenges in developing CertiStr.

The challenge in the automata library is the usage of sets as transition labels. Classical NFAs require only the equality comparison for transition labels. But for s-NFAs, we have more operations for labels. For instance, we require the operations **semiI**, **intersectionI**, **nemptyI**, **memI** for intervals. These operations make it harder to prove the correctness of the automata library.

The challenge in the forward-propagation module is the proofs of Theorem 4.2, 4.3, and 4.4. These theorems need around 4000 lines of code to prove.

8 Related Works

String solvers. As introduced in Section 1, there already exist various non-certified string solvers, such as Kaluza [35], CVC4 [28], Z3 [17], Z3-str3 [4, 5], Z3-Trau [9], S3P [39], OSTRICH [11–13], SLOTH [22], and Norn [1], among many others. These solvers are intricate and support more string operations than CertiStr. Some of the solvers, such as Z3-str3 and CVC4, opted to support more string operations and settle with incomplete solvers (e.g. with no guarantee of termination) that could still solve many constraints that arise in practice. Other solvers are designed with stronger theoretical guarantees; for instance, OSTRICH is complete for the straight-line string constraint fragment [29]. CertiStr can solve the string constraints that are out of the scope of the straight-line fragment and guarantees to terminate for the constraints without cyclic concatenation dependencies, but

it will return *unknown* for constraints that do not satisfy the tree property.

Symbolic Automata. CertiStr depends on automata operations for regular expression propagation and consistency checking. An efficient implementation of automata is crucial. In comparison to finite-state automata in the classical sense, symbolic automata [15, 16] have proven to be more appropriate for applications like model checking, natural language processing, and networking. *Certified automata libraries* [7, 26] can provide trustworthiness for automata-based analysis in applications. However, to the best of our knowledge, existing certified automata libraries are based on the classical definition of automata, which makes them inefficient for practical applications with very large alphabets. Our work in this paper contributes to the development of certified symbolic automata libraries.

Paradigms toward certified constraint solvers. Two main paradigms have emerged for the verification of constraint solvers: (i) the verification of the solver itself, i.e., the development of solvers with machine-checked end-to-end correctness guarantees. For instance, Shi et al. [36] built a certified SMT quantifier-free bit-vector solver. (ii) The generation of certificates, i.e., the actual solver is not verified, but its outputs are accompanied by proof witnesses that can then be independently checked by a verified, trustworthy certificate checker. For instance, Ekici et al. [19] provide an independent and certified checker for SAT and SMT proof witnesses. In our work, we follow the first paradigm.

Security applications based on string analysis. There are various security applications of automata-based analysis and string solvers, such as detecting web security vulnerabilities. Yu et al. [43] investigated the approach to use automata theory for detecting security vulnerabilities, such as XSS, in PHP programs. Their work also depends on a forward analysis. But the forward analysis is over the control flow graphs of PHP programs. CertiStr aims to build a stand-alone string solver and the forward-propagation is only over the string constraints. Many developments of string solvers, string analysis techniques, and automata-theoretic techniques were also directly motivated by security vulnerability detection, e.g., [2, 3, 10, 13, 22, 29, 35, 37–39]. To the best of our knowledge, CertiStr is the first tool that provides certification of a string analysis method that is applicable to security vulnerability detection. We believe there is a need for further development of certified string analyzers: string solving techniques are intricate and hence error-prone, but are applied for security analysis whose correctness is of critical importance.

9 Conclusion and Future Works

In this paper, we present CertiStr, a certified string solver for the theory of concatenation and regular constraints. The backend of CertiStr is verified in Isabelle proof assistant, which provides a rigorous guarantee for the results of the

solver. We ran CertiStr over the benchmark Kaluza to show the efficacy of CertiStr.

As future works, firstly, we plan to support more string operations, such as *string replacement* and *capture groups*, which are also widely used in programming languages, such as JavaScript and PHP, increasing the applicability of CertiStr. Secondly, the front end of CertiStr still needs to be verified in Isabelle, especially the correctness of the desugaring from string constraints with monadic length functions and disjunctions to the language of CertiStr.

Acknowledgments

The authors would like to thank anonymous reviewers for their valuable comments. This research was supported in part by the ERC Starting Grant 759969 (AV-SMP), Max-Planck Fellowship, Amazon Research Award, the Swedish Research Council (VR) under grant 2018-04727, and by the Swedish Foundation for Strategic Research (SSF) under the project WebSec (Ref. RIT17-0011).

References

- [1] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Yu-Fang Chen, Lukás Holik, Ahmed Rezzine, Philipp Rümmer, and Jari Stenman. 2014. String Constraints for Verification. In *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 (Lecture Notes in Computer Science, Vol. 8559)*, Armin Biere and Roderick Bloem (Eds.). Springer, 150–166. https://doi.org/10.1007/978-3-319-08867-9_10
- [2] Roberto Amadini, Mak Andron, Graeme Gange, Peter Schachte, Harald Søndergaard, and Peter J. Stuckey. 2019. Constraint Programming for Dynamic Symbolic Execution of JavaScript. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 16th International Conference, CPAIOR 2019, Thessaloniki, Greece, June 4-7, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11494)*, Louis-Martin Rousseau and Kostas Stergiou (Eds.). Springer, 1–19. https://doi.org/10.1007/978-3-030-19212-9_1
- [3] John Backes, Pauline Bolignano, Byron Cook, Catherine Dodge, Andrew Gacek, Kasper Søe Luckow, Neha Rungta, Oksana Tkachuk, and Carsten Varming. 2018. Semantic-based Automated Reasoning for AWS Access Policies using SMT. In *2018 Formal Methods in Computer Aided Design, FMCAD 2018, Austin, TX, USA, October 30 - November 2, 2018*, Nikolaj Bjørner and Arie Gurfinkel (Eds.). IEEE, 1–9. <https://doi.org/10.23919/FMCAD.2018.8602994>
- [4] Murphy Berzish, Vijay Ganesh, and Yunhui Zheng. 2017. Z3str3: A string solver with theory-aware heuristics. In *2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, October 2-6, 2017*, Daryl Stewart and Georg Weissenbacher (Eds.). IEEE, 55–59. <https://doi.org/10.23919/FMCAD.2017.8102241>
- [5] Murphy Berzish, Mitja Kulczynski, Federico Mora, Florin Manea, Joel D. Day, Dirk Nowotka, and Vijay Ganesh. 2021. An SMT Solver for Regular Expressions and Linear Arithmetic over String Length. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12760)*, Alexandra Silva and K. Rustan M. Leino (Eds.). Springer, 289–312. https://doi.org/10.1007/978-3-030-81688-9_14
- [6] Dmitry Blotsky, Federico Mora, Murphy Berzish, Yunhui Zheng, Ifaz Kabir, and Vijay Ganesh. 2018. StringFuzz: A Fuzzer for String Solvers. In *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 10982)*, Hana Chockler and Georg Weissenbacher (Eds.). Springer, 45–51. https://doi.org/10.1007/978-3-319-96142-2_6
- [7] Julian Brunner. 2017. Transition Systems and Automata Isabelle Library. Arch. Formal Proofs. https://www.isa-afp.org/entries/Transition_Systems_and_Automata.html
- [8] Alexandra Bugariu and Peter Müller. 2020. Automatically testing string solvers. In *ICSE '20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*, Gregg Rothermel and Doo-Hwan Bae (Eds.). ACM, 1459–1470. <https://doi.org/10.1145/3377811.3380398>
- [9] Diep Bui and contributors. 2019. Z3-Trau. <https://github.com/diepbp/z3-trau>.
- [10] Taolue Chen, Yan Chen, Matthew Hague, Anthony W. Lin, and Zhilin Wu. 2018. What is decidable about string constraints with the ReplaceAll function. *Proc. ACM Program. Lang.* 2, POPL (2018), 3:1–3:29. <https://doi.org/10.1145/3158091>
- [11] Taolue Chen, Alejandro Flores-Lamas, Matthew Hague, Zhilei Han, Denghang Hu, Shuanglong Kan, Anthony W. Lin, Philipp Ruemmer, and Zhilin Wu. 2022. Solving String Constraints with Regex-Dependent Functions through Transducers with Priorities and Variables. *Proc. ACM Program. Lang.* 6, POPL (2022).
- [12] Taolue Chen, Matthew Hague, Jinlong He, Denghang Hu, Anthony Widjaja Lin, Philipp Rümmer, and Zhilin Wu. 2020. A Decision Procedure for Path Feasibility of String Manipulating Programs with Integer Data Type. In *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*. 325–342. https://doi.org/10.1007/978-3-030-59152-6_18
- [13] Taolue Chen, Matthew Hague, Anthony W. Lin, Philipp Rümmer, and Zhilin Wu. 2019. Decision procedures for path feasibility of string-manipulating programs with complex operations. *Proc. ACM Program. Lang.* 3, POPL (2019), 49:1–49:30. <https://doi.org/10.1145/3290362>
- [14] Lucas C. Cordeiro, Pascal Kesseli, Daniel Kroening, Peter Schrammel, and Marek Trtik. 2018. JBMC: A Bounded Model Checking Tool for Verifying Java Bytecode. In *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10981)*, Hana Chockler and Georg Weissenbacher (Eds.). Springer, 183–190. https://doi.org/10.1007/978-3-319-96145-3_10
- [15] Loris D'Antoni and Margus Veanes. 2017. The Power of Symbolic Automata and Transducers. In *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10426)*, Rupak Majumdar and Viktor Kuncak (Eds.). Springer, 47–67. https://doi.org/10.1007/978-3-319-63387-9_3
- [16] Loris D'Antoni and Margus Veanes. 2021. Automata modulo theories. *Commun. ACM* 64, 5 (2021), 86–95. <https://doi.org/10.1145/3419404>
- [17] Leonardo Mendonça de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *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 (Lecture Notes in Computer Science, Vol. 4963)*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24
- [18] Volker Diekert. 2002. Makanin's Algorithm. In *Algebraic Combinatorics on Words*, M. Lothaire (Ed.). Encyclopedia of Mathematics and its Applications, Vol. 90. Cambridge University Press, Chapter 12, 387–442.
- [19] Burak Ekici, Guy Katz, Chantal Keller, Alain Mebsout, Andrew J. Reynolds, and Cesare Tinelli. 2016. Extending SMTCoa, a Certified Checker for SMT (Extended Abstract). In *Proceedings First International*

- Workshop on Hammers for Type Theories, HaTT@IJCAR 2016, Coimbra, Portugal, July 1, 2016 (EPTCS, Vol. 210)*, Jasmin Christian Blanchette and Cezary Kaliszyk (Eds.). 21–29. <https://doi.org/10.4204/EPTCS.210.5>
- [20] Claudio Gutiérrez. 1998. Solving Equations in Strings: On Makanin’s Algorithm. In *LATIN ’98: Theoretical Informatics, Third Latin American Symposium, Campinas, Brazil, April, 20-24, 1998, Proceedings (Lecture Notes in Computer Science, Vol. 1380)*, Claudio L. Lucchesi and Arnaldo V. Moura (Eds.). Springer, 358–373. <https://doi.org/10.1007/BFb0054336>
- [21] Matthew Hague, Anthony W. Lin, Philipp Rümmer, and Zhilin Wu. 2020. Monadic Decomposition in Integer Linear Arithmetic. In *Automated Reasoning - 10th International Joint Conference, IJCAR 2020, Paris, France, July 1-4, 2020, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12166)*, Nicolas Peltier and Viorica Sofronie-Stokkermans (Eds.). Springer, 122–140. https://doi.org/10.1007/978-3-030-51074-9_8
- [22] Lukás Holík, Petr Janku, Anthony W. Lin, Philipp Rümmer, and Tomáš Vojnar. 2018. String constraints with concatenation and transducers solved efficiently. *Proc. ACM Program. Lang.* 2, POPL (2018), 4:1–4:32. <https://doi.org/10.1145/3158092>
- [23] Artur Jez. 2016. Recompression: A Simple and Powerful Technique for Word Equations. *J. ACM* 63, 1 (2016), 4:1–4:51. <https://doi.org/10.1145/2743014>
- [24] Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. 2002. MONA Implementation Secrets. *Int. J. Found. Comput. Sci.* 13, 4 (2002), 571–586. <https://doi.org/10.1142/S012905410200128X>
- [25] Peter Lammich. 2013. Automatic Data Refinement. In *Interactive Theorem Proving - 4th International Conference, ITP 2013, Rennes, France, July 22-26, 2013, Proceedings (Lecture Notes in Computer Science, Vol. 7998)*, Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie (Eds.). Springer, 84–99. https://doi.org/10.1007/978-3-642-39634-2_9
- [26] Peter Lammich. 2014. The CAVA Automata Isabelle Library. Arch. Formal Proofs. https://www.isa-afp.org/entries/CAVA_Automata.html.
- [27] Peter Lammich and Andreas Lochbihler. 2010. The Isabelle Collections Framework. In *Interactive Theorem Proving, First International Conference, ITP 2010, Edinburgh, UK, July 11-14, 2010, Proceedings (Lecture Notes in Computer Science, Vol. 6172)*, Matt Kaufmann and Lawrence C. Paulson (Eds.). Springer, 339–354. https://doi.org/10.1007/978-3-642-14052-5_24
- [28] Tianyi Liang, Andrew Reynolds, Cesare Tinelli, Clark W. Barrett, and Morgan Deters. 2014. A DPLL(T) Theory Solver for a Theory of Strings and Regular Expressions. In *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 (Lecture Notes in Computer Science, Vol. 8559)*, Armin Biere and Roderick Bloem (Eds.). Springer, 646–662. https://doi.org/10.1007/978-3-319-08867-9_43
- [29] Anthony Widjaja Lin and Pablo Barceló. 2016. String solving with word equations and transducers: towards a logic for analysing mutation XSS. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM, 123–136. <https://doi.org/10.1145/2837614.2837641>
- [30] Gennady S Makanin. 1977. The problem of solvability of equations in a free semigroup. *Sbornik: Mathematics* 32, 2 (1977), 129–198.
- [31] Muhammad Numair Mansur, Maria Christakis, Valentin Wüstholtz, and Fuyuan Zhang. 2020. Detecting critical bugs in SMT solvers using blackbox mutational fuzzing. In *ESEC/FSE ’20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*, Prem Devanbu, Myra B. Cohen, and Thomas Zimmermann (Eds.). ACM, 701–712. <https://doi.org/10.1145/3368089.3409763>
- [32] Yasuhiko Minamide. 2005. Static approximation of dynamically generated Web pages. In *Proceedings of the 14th international conference on World Wide Web, WWW 2005, Chiba, Japan, May 10-14, 2005*, Allan Ellis and Tatsuya Hagino (Eds.). ACM, 432–441. <https://doi.org/10.1145/1060745.1060809>
- [33] Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*. Vol. 2283. Springer Science & Business Media.
- [34] Yannic Noller, Corina S. Pasareanu, Aymeric Fromherz, Xuan-Bach Dinh Le, and Willem Visser. 2019. Symbolic Pathfinder for SV-COMP - (Competition Contribution). In *Tools and Algorithms for the Construction and Analysis of Systems - 25 Years of TACAS: TOOLympics, Held as Part of ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 11429)*, Dirk Beyer, Marieke Huisman, Fabrice Kordon, and Bernhard Steffen (Eds.). Springer, 239–243. https://doi.org/10.1007/978-3-030-17502-3_21
- [35] Prateek Saxena, Devdatta Akhawe, Steve Hanna, Feng Mao, Stephen McCamant, and Dawn Song. 2010. A Symbolic Execution Framework for JavaScript. In *31st IEEE Symposium on Security and Privacy, S&P 2010, 16-19 May 2010, Berkeley/Oakland, California, USA*. IEEE Computer Society, 513–528. <https://doi.org/10.1109/SP.2010.38>
- [36] Xiaomu Shi, Yu-Fu Fu, Jiayang Liu, Ming-Hsien Tsai, Bow-Yaw Wang, and Bo-Yin Yang. 2021. CoqQFBV: A Scalable Certified SMT Quantifier-Free Bit-Vector Solver. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12760)*, Alexandra Silva and K. Rustan M. Leino (Eds.). Springer, 149–171. https://doi.org/10.1007/978-3-030-81688-9_7
- [37] Caleb Stanford, Margus Veanes, and Nikolaj Bjørner. 2021. Symbolic Boolean derivatives for efficiently solving extended regular expression constraints. In *PLDI ’21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, Stephen N. Freund and Eran Yahav (Eds.). ACM, 620–635. <https://doi.org/10.1145/3453483.3454066>
- [38] Minh-Thai Trinh, Duc-Hiep Chu, and Joxan Jaffar. 2014. S3: A Symbolic String Solver for Vulnerability Detection in Web Applications. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM, 1232–1243. <https://doi.org/10.1145/2660267.2660372>
- [39] Minh-Thai Trinh, Duc-Hiep Chu, and Joxan Jaffar. 2016. Progressive Reasoning over Recursively-Defined Strings. In *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9779)*, Swarat Chaudhuri and Azadeh Farzan (Eds.). Springer, 218–240. https://doi.org/10.1007/978-3-319-41528-4_12
- [40] Thomas Tuerk. 2012. A Formalisation of Finite Automata in Isabelle / HOL. <https://www.thomas-tuerk.de/assets/talks/cava.pdf>.
- [41] Margus Veanes, Pieter Hooimeijer, Benjamin Livshits, David Molnar, and Nikolaj Bjørner. 2012. Symbolic finite state transducers: algorithms and applications. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, John Field and Michael Hicks (Eds.). ACM, 137–150. <https://doi.org/10.1145/2103656.2103674>
- [42] Hung-En Wang, Tzung-Lin Tsai, Chun-Han Lin, Fang Yu, and Jie-Hong R. Jiang. 2016. String Analysis via Automata Manipulation with Logic Circuit Representation. In *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9779)*, Swarat Chaudhuri and Azadeh Farzan (Eds.). Springer, 241–260. https://doi.org/10.1007/978-3-319-41528-4_13
- [43] Fang Yu, Muath Alkhalaf, Tevfik Bultan, and Oscar H. Ibarra. 2014. Automata-based symbolic string analysis for vulnerability detection. *Formal Methods Syst. Des.* 44, 1 (2014), 44–70. <https://doi.org/10.1007/s10703-013-0189-1>