



# StringFuzz: A Fuzzer for String Solvers

Dmitry Blotsky<sup>1</sup>(✉), Federico Mora<sup>2</sup>, Murphy Berzish<sup>1</sup>, Yunhui Zheng<sup>3</sup>,  
Ifaz Kabir<sup>1</sup>, and Vijay Ganesh<sup>1</sup>

<sup>1</sup> University of Waterloo, Waterloo, Canada  
{dblotsky, vganesh}@uwaterloo.ca

<sup>2</sup> University of Toronto, Toronto, Canada  
fmora@cs.toronto.edu

<sup>3</sup> IBM T.J. Watson Research Center,  
Yorktown Heights, USA



**Abstract.** In this paper, we introduce StringFuzz: a modular SMT-LIB problem instance transformer and generator for string solvers. We supply a repository of instances generated by StringFuzz in SMT-LIB 2.0/2.5 format. We systematically compare Z3str3, CVC4, Z3str2, and Norn on groups of such instances, and identify those that are particularly challenging for some solvers. We briefly explain our observations and show how StringFuzz helped discover causes of performance degradations in Z3str3.

## 1 Introduction

In recent years, many algorithms for solving string constraints have been developed and implemented in SMT solvers such as Norn [6], CVC4 [12], and Z3 (e.g., Z3str2 [13] and Z3str3 [7]). To validate and benchmark these solvers, their developers have relied on hand-crafted input suites [1, 4, 5] or real-world examples from a limited set of industrial applications [2, 11]. These test suites have helped developers identify implementation defects and develop more sophisticated solving heuristics. Unfortunately, as more features are added to solvers, these benchmarks often remain stagnant, leaving increasing functionality untested. As such, there is an acute need for a more robust, inexpensive, and automatic way of generating benchmarks to test the correctness and performance of SMT solvers.

Fuzzing has been used to test all kinds of software including SAT solvers [10]. Inspired by the utility of fuzzers, we introduce StringFuzz and describe its value as an exploratory testing tool. We demonstrate its efficacy by presenting limitations it helped discover in leading string solvers. To the best of our knowledge, StringFuzz is the only tool aimed at automatic generation of string constraints. StringFuzz can be used to mutate or transform existing benchmarks, as well as randomly generate structured instances. These instances can be scaled with respect to a variety of parameters, e.g., length of string constants, depth of concatenations (concats) and regular expressions (regexes), number of variables, number of length constraints, and many more.

## Contributions

1. <sup>1</sup> **The StringFuzz tool:** In Sect. 2, we describe a modular fuzzer that can transform and generate SMT-LIB 2.0/2.5 string and regex instances. Scaling inputs (e.g., long string constants, deep concatenations) are particularly useful in identifying asymptotic behaviors in solvers, and StringFuzz has many options to generate them. We briefly document StringFuzz’s components and modular architecture. We provide example use cases to demonstrate its utility as an exploratory solver testing tool.
2. **A repository of SMT-LIB 2.0/2.5 instances:** We present a repository of SMT-LIB 2.0/2.5 string and regex instance suites that we generated using StringFuzz in Sect. 3. This repository consists of two categories: one with new instances generated by StringFuzz (**generated**); and another with transformed instances generated from a small suite of industrial benchmarks (**transformed**).
3. **Experimental Results and Analysis:** We compare the performance of Z3str3, CVC4, Z3str2, and Norn on the StringFuzz suites *Concats-Balanced*, *Concats-Big*, *Concats-Extracts-Small*, and *Different-Prefix* in Sect. 4. We highlight these suites because they make some solvers perform poorly, but not others. We analyze our experimental results, and pinpoint algorithmic limitations in Z3str3 that cause poor performance.

## 2 StringFuzz

**Implementation and Architecture.** StringFuzz is implemented as a Python package, and comes with several executables to generate, transform, and analyze SMT-LIB 2.0/2.5 string and regex instances. Its components are implemented as UNIX “filters” to enable easy integration with other tools (including themselves). For example, the outputs of generators can be piped into transformers, and transformers can be chained to produce a stream of tuned inputs to a solver. StringFuzz is composed of the following tools:

### `stringfuzzg`

This tool generates SMT-LIB instances. It supports several generators and options that specify its output. Details can be found in Table 1a.

### `stringfuzzx`

This tool transforms SMT-LIB instances. It supports several transformers and options that specify its output and input, which are explained in Table 1b. Note that transformers *Translate* and *Reverse* also preserve satisfiability under certain conditions.

### `stringstats`

This tool takes an SMT-LIB instance as input and outputs its properties: the number of variables/literals, the max/median syntactic depth of expressions, the max/median literal length, etc.

---

<sup>1</sup> All source code, problem suites, and supplementary material referenced in this paper are available at the StringFuzz website [3].

**Table 1.** StringFuzz built-in (a) generators and (b) transformers.(a) `stringfuzzg` built-in generators.

Name	Generates instances that have ...
<i>Concat</i>	Long concats and optional random extracts.
<i>Lengths</i>	Many variables (and their concats) with length constraints.
<i>Overlaps</i>	An expression of the form $A.X = X.B$ .
<i>Equality</i>	An equality among concats, each with variables or constants.
<i>Regex</i>	Regexes of varying complexity.
<i>Random-Text</i>	Totally random ASCII text.
<i>Random-AST</i>	Random string and regex constraints.

(b) `stringfuzzx` built-in transformers.

Name	The transformer ...
<i>Fuzz</i>	Replaces literals and operators with similar ones.
<i>Graft</i>	Randomly swaps non-leaf nodes with leaf nodes.
<i>Multiply</i> <sup>a</sup>	Multiplies integers and repeats strings by N.
<i>Nop</i>	Does nothing (can translate between SMT-LIB 2.0/2.5).
<i>Reverse</i> <sup>b</sup>	Reverses all string literals and concat arguments.
<i>Rotate</i>	Rotates compatible nodes in syntax tree.
<i>Translate</i> <sup>b</sup>	Permutates the alphabet.
<i>Unprintable</i>	Replaces characters in literals with unprintable ones.

<sup>a</sup>Can guarantee satisfiable output instances from satisfiable input instances [3].

<sup>b</sup>Can guarantee input and output instances will be equisatisfiable [3].

We organized StringFuzz to be easily extended. To show this, we note that while the whole project contains 3,183 lines of code, it takes an average of 45 lines of code to create a transformer. StringFuzz can be installed from source, or from the Python PIP package repository.

**Regex Generating Capabilities.** StringFuzz can generate and transform instances with regex constraints. For example, the command “`stringfuzzg regex -r 2 -d 1 -t 1 -M 3 -X 10`” produces this instance:

```
(set-logic QF_S)
(declare-fun var0 () String)
(assert (str.in.re var0 (re.+ (str.to.re "R5"))))
(assert (str.in.re var0 (re.+ (str.to.re "!PC"))))
(assert (<= 3 (str.len var0)))
(assert (<= (str.len var0) 10))
(check-sat)
```

Each instance is a set of one or more regex constraints on a single variable, with optional maximum and minimum length constraints. Each regex constraint is a concatenation (`re.++` in SMT-LIB string syntax) of regex terms:

$$(re.++ T1 (re.++ T2 \dots (re.++ T_{n-1} T_n )))$$

and each term  $T_i$  is recursively defined as any one of: repetition ( $\mathbf{re.}^*$ ), Kleene star ( $\mathbf{re.}^+$ ), union ( $\mathbf{re.union}$ ), or a character literal. Nested operators are nested up to a specified (using the `--depth` flag) depth of recursion. Terms at depth 0 are regex constants. Below are 3 example regexes (in regex, not SMT-LIB, syntax) of depth 2 that can be produced this way:

$$((a|b)|(cc)^+) \quad ((ddd)^*)^+ \quad ((ee)^+ | ((fff)^*))^+$$

**Equisatisfiable String Transformations.** StringFuzz can also transform problem instances. This is done by manipulating parsed syntax trees. By default most of the built-in transformers only guarantee well-formedness, however, some can even guarantee equisatisfiability. Table 1b lists the built-in transformers and notes these guarantees.

**Example Use Case.** In Sect. 3 we use StringFuzz to generate benchmark suites in a batch mode. We can also use StringFuzz for on-line exploratory debugging. For example, the script below repeatedly feeds random StringFuzz instances to CVC4 until the solver produces an error:

```
while stringfuzzg -r random-ast -m \
  | tee instance.smt25 | cvc4 --lang smt2.5 --tlimit=5000 --strings-exp; do
  sleep 0
done
```

### 3 Instance Suites

In this section, we describe the benchmark suites we generated with StringFuzz, and on which we conducted our experimental evaluation. Table 2a lists instances that were generated by `stringfuzzg`. Table 2b lists instances derived from existing seed instances by iteratively applying `stringfuzzx`. Every transformed instance is named according to its seed and the transformations it undertook. For example, `z3-regex-1-fuzz-graft.smt2` was transformed by applying *Fuzz* and then *Graft* to `z3-regex-1.smt2`.

The *Amazon* category contains 472 instances derived from two seeds supplied by our industrial collaborators. The *Regex* category is seeded by the Z3str2 regex test suite [4], which contains 42 instances. Through cumulative transformations we expanded the 42 seeds to 7,551 unique instances. Finally, the *Sanitizer* category is obtained from five industrial e-mail address and IPv4 sanitizers.

## 4 Experimental Results and Analysis

We generated several problem instance suites with StringFuzz that made one solver perform poorly, but not others.<sup>2</sup> They are *Concats-Balanced*, *Concats-Big*, *Concats-Extracts-Small*, and *Different-Prefix*. Figure 1 shows the suites that

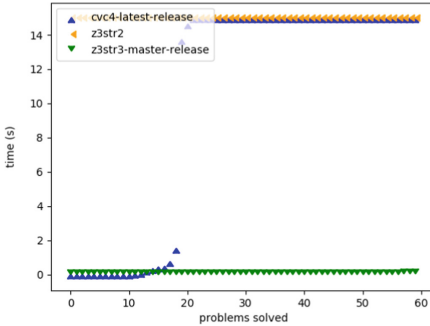

<sup>2</sup> Only the results that made one solver perform poorly and not others are presented, but results for all StringFuzz suites are available on the StringFuzz website [3].

**Table 2.** Repository of 10,258 SMT-LIB 2.0/2.5 instances.(a) `stringfuzzg`-generated instances.

Name	Instances have a ...	Quantity
<i>Concats-<math>\{Small, Big\}</math></i>	Right-heavy, deep tree of concats.	120
<i>Concats-Balanced</i>	Balanced, deep tree of concats.	100
<i>Concats-Extracts-<math>\{Small, Big\}</math></i>	Single concat tree, with character extractions.	120
<i>Lengths-<math>\{Long, Short\}</math></i>	Single, large length constraint on a variable.	200
<i>Lengths-Concats</i>	Tree of fixed-length concats of variables.	100
<i>Overlaps-<math>\{Small, Big\}</math></i>	Formula of the form $A.X = X.B$ .	80
<i>Regex-<math>\{Small, Big\}</math></i>	Complex regex membership test.	120
<i>Many-Regexes</i>	Multiple random regex membership tests.	40
<i>Regex-Deep</i>	Regex membership test with many nested operators.	45
<i>Regex-Pair</i>	Test for membership in one regex, but not another.	40
<i>Regex-Lengths</i>	Regex membership test, and a length constraint.	40
<i>Different-Prefix</i>	Equality of two deep concats with different prefixes.	60

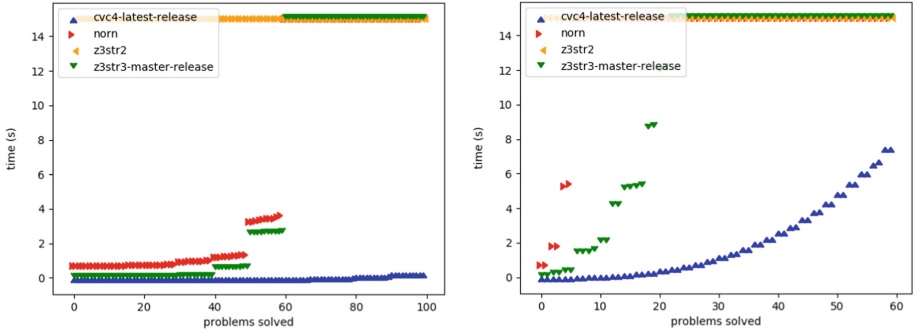
(b) `stringfuzzx`-generated instances.

Name	Seed	Quantity
<i>Amazon</i>	Two industrial regex membership instances.	472
<i>Regex</i>	Z3str2 regular expression test suite.	7,551
<i>Sanitizer</i>	Five e-mail and IPv4 sanitiser examples.	1,170

(a) Performance on *Concats-Extracts-Small*(b) Performance on *Different-Prefix***Fig. 1.** Instances hard for CVC4

were uniquely difficult for CVC4. Figure 2 shows the suites that were uniquely difficult for Z3str3. All experiments were conducted in series, each with a timeout of 15 s, on an Ubuntu Linux 16.04 computer with 32 GB of RAM and an Intel® Core™ i7-6700 CPU (3.40 GHz).

**Usefulness to Z3str3: A Case Study.** StringFuzz’s ability to produce scaling instances helped uncover several implementation issues and performance limitations in Z3str3. Scaling inputs can reveal issues that would normally be out of scope for unit tests or industrial benchmarks. Three different performance and

(a) Performance on *Concats-Balanced*(b) Performance on *Concats-Big***Fig. 2.** Instances hard for Z3str3

implementation bugs were identified and fixed in Z3str3 as a result of testing with the StringFuzz scaling suites *Lengths-Long* and *Concats-Big*.

StringFuzz also helped identify a number of performance-related issues and opportunities for new heuristics in Z3str3. For example, by examining Z3str3’s execution traces on the instances in the *Concats-Big* suite we discovered a potential new heuristic. In particular, Z3str3 does not make full use of the solving context (e.g. some terms are empty strings) to simplify the concatenations of a long list of string terms before trying to reason about the equivalences among sub-terms. Z3str3 therefore introduces a large number of unnecessary intermediate variables and propagations.

## 5 Related Work

Many solver developers create their own test suites to validate their solvers [1, 4, 5]. Several popular instance suites are also publicly available for solver testing and benchmarking, such as the Kaluza [2] and Kausler [11] suites. There are likewise several fuzzers and instance generators currently available, but none of them can generate or transform string and regex instances. For example, the FuzzSMT [9] tool generates SMT-LIB instances with bit-vectors and arrays, but does not support strings or regexes. The SMTpp [8] tool pre-processes and simplifies instances, but does not generate new ones or fuzz existing ones.

## References

1. CVC4 regression test suite. <https://github.com/CVC4/CVC4/tree/master/test/regress>
2. Kaluza benchmark suite. <http://webblaze.cs.berkeley.edu/2010/kaluza/>
3. Stringfuzz source code, benchmark suites, and supplemental material. <http://stringfuzz.dmitryblotsky.com>

4. Z3str2 test suite. <https://github.com/z3str/Z3-str/tree/master/tests>
5. Z3str3 test scripts. <https://github.com/Z3Prover/z3/tree/master/src/test>
6. Abdulla, P.A., et al.: Norn: an SMT solver for string constraints. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 462–469. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-21690-4\\_29](https://doi.org/10.1007/978-3-319-21690-4_29)
7. Berzish, M., Ganesh, V., Zheng, Y.: Z3str3: a string solver with theory-aware heuristics. In: Stewart, D., Weissenbacher, G., (eds.), 2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, 2–6 October 2017, pp. 55–59. IEEE (2017)
8. Bonichon, R., Déharbe, D., Dobal, P., Tavares, C.: SMTpp: preprocessors and analyzers for SMT-LIB. In: Proceedings of the 13th International Workshop on Satisfiability Modulo Theories, SMT 2015 (2015)
9. Brummayer, R., Biere, A.: Fuzzing and delta-debugging SMT solvers. In: Proceedings of the 7th International Workshop on Satisfiability Modulo Theories, SMT 2009, pp. 1–5. ACM, New York, NY, USA (2009)
10. Brummayer, R., Lonsing, F., Biere, A.: Automated testing and debugging of SAT and QBF solvers. In: Strichman, O., Szeider, S. (eds.) SAT 2010. LNCS, vol. 6175, pp. 44–57. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-14186-7\\_6](https://doi.org/10.1007/978-3-642-14186-7_6)
11. Kausler, S., Sherman, E.: Evaluation of string constraint solvers in the context of symbolic execution. In: Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering, ASE 2014, pp. 259–270. ACM, New York, NY, USA (2014)
12. Liang, T., Reynolds, A., Tinelli, C., Barrett, C., Deters, M.: A DPLL( $T$ ) theory solver for a theory of strings and regular expressions. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 646–662. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08867-9\\_43](https://doi.org/10.1007/978-3-319-08867-9_43)
13. Zheng, Y., Zhang, X., Ganesh, V.: Z3-str: a Z3-based string solver for web application analysis. In: Meyer, B., Baresi, L., Mezini, M., (eds.) Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, ESEC/FSE 2013, Saint Petersburg, Russian Federation, 18–26 August 2013, pp. 114–124. ACM (2013)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

