

# Refactoring-based Benchmarks Generation

June 19, 2023

## Problem description

Benchmarks programs are required in a variety of contexts: for testing compilers, static analyzers, or program verifiers, as data sets for machine learning algorithms, for evaluating program repair techniques, etc. However, systematically generating programs that are *simple* (i.e., easy to understand) but also *expressive* (i.e., include various language features) is very challenging.

Csmith [5], a widely-used generator for random C programs, relies on complex static analyses and run time checks to ensure that the programs do not have undefined and unspecified behavior. Enumeration-based approaches [6] produce all the C programs with  $k$  variables that match a small, syntactic skeleton. Techniques for synthesizing Java programs either start from user-provided generators [1] or require an Alloy [3] model of the constraints which should be satisfied by the programs [4].

To overcome some of these limitations, recent work [2] automatically constructs incrementally complex Java programs that throw null pointer exception. The approach uses simple, manually written seed programs (when executed, these always raise the exception) and then applies syntactic transformations that preserve the faulty behavior for at least one set of inputs. However, extending this technique to generate benchmarks with other properties requires expert knowledge in designing the transformations. We believe that this process can be automated by constructing the transformations using code refactoring.

*The goal of this project is to develop a generic framework for synthesizing diverse but increasingly complex benchmarks programs that fulfill a given set of properties. The benchmarks will be constructed by applying (reverse) refactoring techniques. Our framework should also provide theoretical guarantees (similar to [6]) about the space of the generated benchmarks, with respect to a property-based notion of equivalence between benchmarks that we will define<sup>1</sup>.*

---

<sup>1</sup>Hurmuz defines the equivalence between two benchmarks programs with respect to a static analysis [2], while Zhang et. al [6] include in the generated benchmarks suite only those programs which are not  $\alpha$ -equivalent ( $\alpha$  is a variable renaming transformation).

```

class Benchmark {
  int foo() {
    return 0;
  }
}

```

Figure 1: Simple Java program with one instruction per method.

## Example

Let us assume that our goal is to automatically generate syntactically-correct Java benchmarks in which all the methods have maximum two instructions. That is, the desired set of properties is:

$$P = \{language = Java, \neg errors, \#instructions/method \leq 2\} \quad (2).$$

To synthesize them, we can start with methods with no parameters and exactly one, type correct return instruction, as shown in Figure 1. We can then increase the number of instructions by replacing the returned constant (0 in Figure 1) with an expression that evaluates to the same value. A possible result of this transformation is shown in Figure 2 and has the inverse effect of the refactoring for inlining temporary variables.

## Approach

We will first design a domain specific language to express properties about the benchmarks (e.g., as those in (2)). We will then automatically synthesize the simplest (shortest) programs that fulfill all the required properties. Note that if they exist, these programs do not have to be unique. We can use a compiler to check if they are syntactically correct.

Next, we will collect various code refactorings that, when applied directly or reversed on the simplest programs, preserve each of the properties. We will then apply these refactorings in a systematic way (potentially multiple times), to obtain increasingly complex but *non-redundant* benchmarks. The notation of

```

class Benchmark {
  int foo() {
    int tmp = 0;
    return tmp;
  }
}

```

Figure 2: The program from Figure 1, transformed by introducing a temporary variable.

```
class Benchmark {
  int zero () {
    return 0;
  }
}
```

Figure 3: Potentially redundant Java program with one instruction per method, obtained from the program from Figure 1 by renaming the method.

redundancy is application-dependent. If the benchmarks are used, for example, for testing the Java compiler, then the program from Figure 3 is equivalent with the one from Figure 1, and thus it is redundant. However, if the goal of the benchmarks is to validate an algorithm for inferring method names, then both programs should be kept, as they represent a positive and a negative example.

## Prerequisites

The student is expected to have good programming skills. Prior experience with refactoring techniques is a plus.

## Opportunities

The student will have the chance to learn about state-of-the-art test case generation techniques.

## Contact

Alexandra Bugariu: [bugariua@mpi-sws.org](mailto:bugariua@mpi-sws.org)

## References

- [1] Brett Daniel, Danny Dig, Kely Garcia, and Darko Marinov. Automated testing of refactoring engines. In *Proceedings of the the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*, ESEC-FSE '07, pages 185–194, New York, NY, USA, 2007. Association for Computing Machinery.
- [2] Madalina Hurmuz. Automatically generating Java benchmarks with known errors. Master’s thesis, ETH Zürich, 2022.
- [3] Daniel Jackson. *Software Abstractions: Logic, Language, and Analysis*. The MIT Press, 2012.

- [4] Gustavo Soares, Rohit Gheyi, and Tiago Massoni. Automated behavioral testing of refactoring engines. *IEEE Transactions on Software Engineering*, 39(2):147–162, 2013.
- [5] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. Finding and understanding bugs in c compilers. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '11*, pages 283–294, New York, NY, USA, 2011. Association for Computing Machinery.
- [6] Qirun Zhang, Chengnian Sun, and Zhendong Su. Skeletal program enumeration for rigorous compiler testing. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017*, pages 347–361, New York, NY, USA, 2017. Association for Computing Machinery.