IDOL: Improved Different Optimization Levels Testing for Solidity Compilers

Lantian Li, Yejian Liang, and Zhongxing Yu*
Shandong University, Qingdao, China
lilantian@mail.sdu.edu.cn, yejianliang@mail.sdu.edu.cn, zhongxing.yu@sdu.edu.cn
*corresponding author

Abstract—As blockchain technology continues to evolve and mature, smart contracts have become a key driving force behind the digitization and automation of transactions. Smart contracts greatly simplify and refine the traditional business transaction processes, and thus have had a profound impact on various industries such as finance and supply chain management. However, because smart contracts cannot be modified once deployed, any vulnerabilities or design flaws within the contract cannot be easily fixed, potentially leading to significant financial losses or even legal issues. The compiler, as a critical component in the development process, directly affects the quality and security of smart contracts. This paper innovatively proposes a method, known as the Improved Different Optimization Levels (IDOL), for testing the Solidity compiler. The key idea behind IDOL is to perform reverse optimization transformations (*i.e.*, change optimized form into unoptimized form) to generate semantically equivalent variants of the smart contracts under test, aiming to maximize the opportunities to trigger the compiler's optimization logic. We conducted a preliminary evaluation of IDOL and three confirmed compiler optimization bugs have been uncovered at the time of writing.

Keywords–Compiler Testing; Smart Contract; Solidity

1. Introduction

As blockchain technology continues to evolve and mature, smart contracts have become a key driving force behind the digitization and automation of transactions. A smart contract is a distributed computing paradigm that can automatically execute, control, or record events and actions between parties [1]. The smart contract development language, such as Solidity, provides an efficient way of implementing the automatic execution of contracts. However, the immutable nature of smart contracts also introduces new challenges. Since smart contracts cannot be modified once deployed, any vulnerabilities within the contract cannot be easily fixed. Thus, ensuring the security of smart contracts is of great importance.

The smart contract compiler is a critical component in the smart contract development process, and its security directly impacts both the quality of the contracts deployed on the blockchain and the overall system security. Thus, systematic testing of the compiler is of vital importance. Compiler errors may arise at various stages of the compilation process, such as syntax analysis, semantic analysis, and code generation. Moreover, to generate efficient target code, the compiler must also account for the characteristics and optimization requirements of different hardware platforms. As such, ensuring the compiler's quality is an inherently challenging task. In particular, when using a specific program as test input, determining whether the compiler's output is correct can be difficult due to the lack of explicit specifications. That is, the "test oracle" issue is serious for compiler testing [2].

Randomized Differential Testing (RDT) [3] is an approach that combines randomized testing with differential testing to address the oracle issue. This method generates identical random inputs for multiple systems with different implementations or configurations, and compares their outputs to identify potential errors. Since RDT does not rely on predefined correct outputs, it is well-suited for complex systems that lack explicit specifications, such as compilers and interpreters. Given that many compiler errors are often triggered during the optimization phase [4], Different Optimization Levels (DOL) testing was proposed as a specialized form of RDT [5]. This method applies different optimization levels to the same test program on the same compiler, and compares the consistency of the output results to detect potential optimization errors.

To more effectively detect bugs in smart contract compilers, this paper innovatively proposes an Improved Different Optimization Levels (IDOL) testing method to test the Solidity compiler. The key idea behind IDOL is to perform reverse transformations on typical compiler optimizations to generate semantically equivalent variants of the smart contracts under test, aiming to maximize the opportunities to trigger the compiler's optimization logic. IDOL consists of three major steps. First, a set of original smart contracts are generated in bulk using Solsmith[1]. Then, for each original smart contract, semantically equivalent variants are generated by referencing typical optimization strategies and specific optimization rules of the Solidity compiler. Finally, each variant is compiled and executed under different optimization levels, and the output results are compared. If differences are found in the execution results at different optimization levels, it indicates that there is a bug in at least one optimization path. We conducted a preliminary evaluation of IDOL and three confirmed compiler optimization bugs have been uncovered at the time of writing.

2. The IDOL Approach

The `--optimize-runs` parameter in the Solidity compiler is used to specify the expected execution frequency of each code segment throughout the contract's lifecycle, thereby guiding the optimizer to balance between deployment size and runtime cost. Smaller values typically generate bytecode that is smaller in size but less efficient in execution, making it suitable

[1]https://github.com/logicseek/Solsmith

```
for (int i = 0; i <
      n; i++) {
   x = y + z;
   a[i] = 6 * i + x
        * x;
}
```

(a) Original Code

```
x = y + z;
t1 = x * x;
for (int i = 0; i < n
     ; i++) {
   a[i] = 6 * i + t1;
}
```

(b) Optimized Code

Figure 1: Loop-invariant Code Motion Optimization.

```
int i, a[100];
i = 0;
while (i < 100) {
   a[i] = 0;
   i++;
}
```

(a) Original Code

```
int i, a[100];
i = 0;
if (i < 100) {
   do {
      a[i] = 0;
      i++;
   } while (i < 100);
}
```

(b) Optimized Code

Figure 2: Loop Inversion Optimization.

for low-frequency invocation scenarios. Larger values instead prioritize execution efficiency, making them more appropriate for high-frequency invocation logic. Given the large number of optimization parameters available, to simplify the experiments and cover typical optimization scenarios, we selected three representative configurations for testing: optimization disabled (`--optimize=false`), optimization enabled with `runs=1`, and optimization enabled with `runs=200`.

Unlike the conventional approach which directly applies different optimization levels to the generated test programs, our approach first performs semantically equivalent transformations on the test programs to make it easier to trigger potential errors in the compiler's optimization process. We then proceed with testing under different optimization levels. The equivalent transformations arise from compiler optimizations, currently including loop invariant code motion, loop inversion, and several optimization strategies specific to the Solidity compiler. *By performing reverse optimization transformations, we can induce the compiler to execute the corresponding optimization processes without altering the program semantics, thus increasing the likelihood of bug triggering*. Below we introduce two of the considered compiler optimization strategies.

Loop-invariant Code Motion. It moves loop-invariant statements or expressions outside the loop body without changing the semantics of the program. The two operations `x = y + z` and `x * x` in Figure 1a are loop invariants and will be moved out of the loop body during compilation optimization, thus transforming into the form of Figure 1b.

Loop Inversion. This optimization transformation replaces the `while` loop with an `if` block containing a `do..while` loop. Although the optimized code in Figure 2b may appear more complex than the original code in Figure 2a, it may actually run faster. This is because modern CPUs use instruction pipelines and any jump in the code can cause a pipeline stall, which degrades performance.

Note that the IDOL transformations are reverse, *i.e.*, they convert the optimized form back to the unoptimized form.

3. Preliminary Evaluations and Results

The IDOL testing process is illustrated in Figure 3. First, a series of test programs are generated using the Solsmith tool. Then, a mutation process is applied to these test programs to produce corresponding equivalent variants. These equivalent variants next are given as inputs to the Solidity compiler (*i.e.,* solc) and sloc compiles them at different optimization levels. Finally, for each variant, the output results at different optimization levels are compared. If discrepancies are found, it suggests that there is a bug in at least one optimization path. We conducted a preliminary evaluation of IDOL using 160,000 test programs generated by Solsmith, and have identified three confirmed compiler optimization bugs. As a comparison, we also tried the original DOL approach using the 160,000 generated test programs, but these three bugs are not detected.

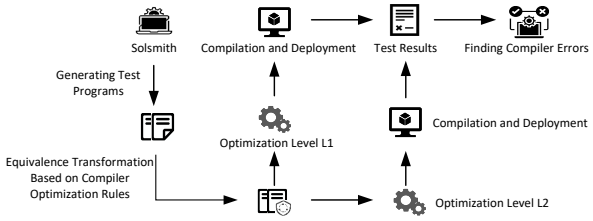

Figure 3: The IDOL Testing Process

Optimizer Keccak Cache Bug. For Keccak-256 hash values with the same memory content but different sizes, the compiler incorrectly treats them as equal, causing the bytecode optimizer to mistakenly reuse previously computed hash values.

Traditional Code Generation Pipeline Bug. The cause is that the traditional code pipeline does not evaluate complex expressions, preventing the contract from executing these expressions, which in turn leads to incorrect contract behavior.

FullInliner Non-Expression Split Parameter Evaluation Order Bug. When the compiler performs FullInliner optimization, it fails to correctly handle the splitting of complex expressions, excessively altering the original logic and resulting in an incorrect parameter evaluation order.

4. Conclusion

This paper presents an innovative IDOL testing method. The key idea behind IDOL is to perform reverse optimization transformations to generate semantically equivalent variants of the smart contracts under test, aiming to maximize the opportunities to trigger the compiler's optimization logic. An initial evaluation of IDOL uncovers three optimization bugs, demonstrating the effectiveness and potential of the approach.

References

[1] L. Li, Y. Liang, Z. Liu, and Z. Yu, "Understanding solidity event logging practices in the wild," in *ESEC/FSE*, p. 300–312, 2023.

[2] E. T. Barr *et al.*, "The oracle problem in software testing: A survey," *TSE*, vol. 41, no. 5, pp. 507–525, 2014.

[3] A. Groce *et al.*, "Randomized differential testing as a prelude to formal verification," in *ICSE*, pp. 621–631.

[4] J. Wu, J. Zheng, Z. Yang, and Z. Yu, "Compiler optimization testing based on optimization-guided equivalence transformations," in *FSE*, 2025.

[5] J. Chen *et al.*, "An empirical comparison of compiler testing techniques," in *ICSE*, pp. 180–190, 2016.