

Dynamic Mutant Subsumption Analysis using LittleDarwin

Ali Parsai
University of Antwerp
Antwerp, Belgium
ali.parsai@uantwerpen.be

Serge Demeyer
University of Antwerp
Antwerp, Belgium
serge.demeyer@uantwerpen.be

ABSTRACT

Many academic studies in the field of software testing rely on mutation testing to use as their comparison criteria. However, recent studies have shown that redundant mutants have a significant effect on the accuracy of their results. One solution to this problem is to use mutant subsumption to detect redundant mutants. Therefore, in order to facilitate research in this field, a mutation testing tool that is capable of detecting redundant mutants is needed. In this paper, we describe how we improved our tool, LittleDarwin, to fulfill this requirement.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging**;

KEYWORDS

Software Testing, Mutation Testing, Mutant Subsumption, Dynamic Mutant Subsumption

ACM Reference Format:

Ali Parsai and Serge Demeyer. 2017. Dynamic Mutant Subsumption Analysis using LittleDarwin. In *Proceedings of 8th International Workshop on Automated Software Testing, Paderborn, Germany, September 4-5, 2017 (A-TEST'17)*, 4 pages.

<https://doi.org/10.1145/3121245.3121249>

1 INTRODUCTION

Many academic studies on fault detection need to assess the quality of their technique using seeded faults. One of the widely-used systematic ways to seed simulated faults into the programs is mutation testing [5]. Mutation testing is the process of injecting faults into software (i.e. creating a mutant), and counting the number of these faults that make at least one test fail (i.e. kill the mutant). The process of creating a mutant consists of applying a predefined transformation on the code (i.e. mutation operator) that converts a version of the code under test into a faulty version. It has been shown that mutation testing is an appropriate method to simulate real faults and perform comparative analysis on testing techniques [2, 3, 10].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

A-TEST'17, September 4-5, 2017, Paderborn, Germany

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5155-3/17/09...\$15.00

<https://doi.org/10.1145/3121245.3121249>

There has been many studies to optimize the process of mutation testing by following the maxim *{do faster, do smarter, do fewer}* [17]. In particular, *do fewer* aims to reduce the number of produced mutants. There are several techniques that implement this logic (e.g. selective mutation [13, 14, 16], and mutant sampling [20, 23–25]). However, only recently the academics began to investigate the threats to validity the redundant mutants introduce in software testing experiments [18]. Papadakis et al. demonstrate that the existence of redundant mutants introduces a significant threat by “artificially inflating the apparent ability of a test technique to detect faults” [18].

One of the recent solutions to alleviate this problem is to use mutant subsumption [1]. Mutant *A* truly subsumes mutant *B* if and only if all inputs that kill *A* also kill *B*. This means that mutant *B* is redundant, since killing *A* is sufficient to know that *B* is also killed. It is possible to provide a more accurate analysis of a testing experiment by determining and discarding the redundant mutants. However, it is often impossible to check mutants for every possible input to the program in practice. Therefore, as a compromise, dynamic mutant subsumption is used instead [1]. Mutant *A* dynamically subsumes mutant *B* with regards to test set *T* if and only if there exists at least one test that kills *A*, and every test that kills *A* also kills *B*. Given the fact that mutant subsumption only recently has been at the center of attention, there are no mature tools that can perform dynamic mutant subsumption analysis on real-life Java programs. This, however, is necessary to facilitate further research on the topic. Therefore we aim to fill this void by developing such tool.

We used LittleDarwin¹ mutation testing framework to implement the features needed to perform dynamic mutant subsumption analysis. LittleDarwin is an extensible and easy to deploy mutation testing tool for Java programs [22]. LittleDarwin has been used previously in several other studies [20, 21], and it is shown to be capable of analyzing large and complicated Java software systems [19].

The rest of the paper is organized as follows: In Section 2, background information about mutation testing is provided. In Section 3, the current state of the art is discussed. In Section 4, we provide details on how LittleDarwin can help performing dynamic mutant subsumption analysis. Finally, we present our conclusions in Section 5.

2 BACKGROUND

The idea of mutation testing was first mentioned by Lipton, and later developed by DeMillo, Lipton and Sayward [5]. The first implementation of a mutation testing tool was done by Timothy Budd in 1980 [4]. Mutation testing is performed as follows: First, a faulty

¹<https://littledarwin.parsai.net/>

```

int multiply(int a, int b) {
  int i, result = 0;
  if (b < 0) { // M3: b >= 0
    a *= -1; // M0: a /= -1 // M1: a *= +1
    b *= -1; // M7: b /= -1 // M4: b *= +1
  }

  for (i = 0; i < b; i++) // M6: i-- // M5: i >= b
    result += a; // M2: result -= a

  return result;
}

```

Figure 1: An Example Mutated Method

version of the software is created by introducing faults into the system (*Mutation*). This is done by applying a known transformation (*Mutation Operator*) on a certain part of the code. After generating the faulty version of the software (*Mutant*), it is passed onto the test suite. If there is an error or failure during the execution of the test suite, the mutant is marked as killed (*Killed Mutant*). If all tests pass, it means that the test suite could not catch the fault, and the mutant has survived (*Survived Mutant*) [9].

If the output of a mutant for all possible input values is the same as the original program, it is called an *equivalent mutant*. It is not possible to create a test case that passes for the original program and fails for an equivalent mutant, because the equivalent mutant is indistinguishable from the original program. This makes the creation of equivalent mutants undesirable, and leads to false positives during mutation testing. In general, detection of equivalent mutants is undecidable due to the halting problem [15]. Manual inspection of all mutants is the only way of filtering all equivalent mutants, which is impractical in real projects due to the amount of work it requires. Therefore, the common practice within today's state-of-the-art is to take precautions to generate as few equivalent mutants as possible, and accept equivalent mutants as a threat to validity (accepting a false positive is less costly than removing a true positive by mistake [6]).

$$\text{Mutation Coverage} = \frac{\text{Number of killed mutants}}{\text{Number of all non-equivalent mutants}} \quad (1)$$


Mutation testing allows software engineers to monitor the fault detection capability of a test suite by means of mutation coverage (see Equation 1) [9]. A test suite is said to achieve *full mutation test adequacy* whenever it can kill all the non-equivalent mutants, thus reaching a mutation coverage of 100%. Such test suite is called a *mutation-adequate test suite*.

3 STATE OF THE ART

Mutant subsumption is defined as the relationship between two non-equivalent mutants A and B in which A subsumes B if and only if all inputs that kill A is guaranteed to kill B [12]. The subsumption relationship for faults has been defined by Kuhn in 1999 [11], but its use for mutation testing has been popularized by Jia et al. for creating hard to kill higher-order mutants [8]. Later on, Ammann et al. tackled the theoretical side of mutant subsumption [1]. In their paper, Ammann et al. define *dynamic mutant subsumption*, which redefines the relationship using test cases. Mutant A dynamically

Table 1: Range of Input Values that Kill Mutants of the Example Mutated Method (left), DMSG for the Example Mutated Method (right)

Mutants	Range of a	Range of b
M0	\emptyset	\emptyset
M1	$(-\infty, \infty) - \{0\}$	$(-\infty, 0)$
M2	$(-\infty, \infty) - \{0\}$	$(-\infty, \infty) - \{0\}$
M3	$(-\infty, \infty) - \{0\}$	$(-\infty, \infty) - \{0\}$
M4	$(-\infty, \infty) - \{0\}$	$(-\infty, 0)$
M5	$(-\infty, \infty) - \{0\}$	$(-\infty, \infty)$
M6	$(-\infty, \infty) - \{0\}$	$(-\infty, \infty) - \{0\}$
M7	\emptyset	\emptyset



subsumes Mutant B if and only if (i) A is killed, and (ii) every test that kills A also kills B. Kurtz et al. [12] use the notion of dynamic mutant subsumption graph (DMSG) to visualize the concept of dynamic mutant subsumption. Each node in a DMSG represents a set of all mutants that are mutually subsuming. Edges in a DMSG represent the dynamic subsumption relationship between the nodes. They introduce the concept of static mutant subsumption graph, which is a result of determining the subsumption relationship between mutants using static analysis techniques.

Figure 1 shows a Java method and its set of mutants. This method takes a and b as input, and returns $a \times b$ as output. To do this, a is added b times. If b is negative, both a and b are negated so that b becomes positive. Table 1 shows the range of input values that kills each mutant. As the table shows, M0 and M7 are equivalent mutants, since the change they introduce does not impact the program semantically. M1 and M4 are killed by the same range of inputs. The same holds true for M2, M3, and M6. It can be seen that $\{M1, M4\}$ truly subsume $\{M2, M3, M6\}$, since any input that kills M1 or M4, also kills M2, M3, and M6; however, the opposite does not hold. Also, $\{M2, M3, M6\}$ truly subsume $\{M5\}$ for the same reason. Using a test suite that includes a test case from each of the input ranges in Table 1, it is possible to draw the DMSG for this method.

The main purpose behind the use of mutant subsumption is to reliably detect redundant mutants, which create multiple threats to the validity of mutation testing [18]. This is often done by determining the dynamic subsumption relationship among a set of mutants, and keeping only those that are not subsumed by any other mutant. In our example, keeping only M1 (or M4) suffices, since it subsumes all the other mutants.

4 DYNAMIC MUTANT SUBSUMPTION ANALYSIS WITH LITLEDARWIN

Figure 3 shows the input and output of LittleDarwin's dynamic mutant subsumption (DMS) component. To facilitate dynamic mutant subsumption analysis in LittleDarwin, we retain all the output provided by the build system for each mutant. As a result, we can parse this output and extract useful information, e.g. which test cases kill a particular mutant. LittleDarwin's DMS component can then use this information to determine dynamic subsumption relationship between each mutant pair. This component then outputs the results in two different ways: (i) the dynamic mutant subsumption

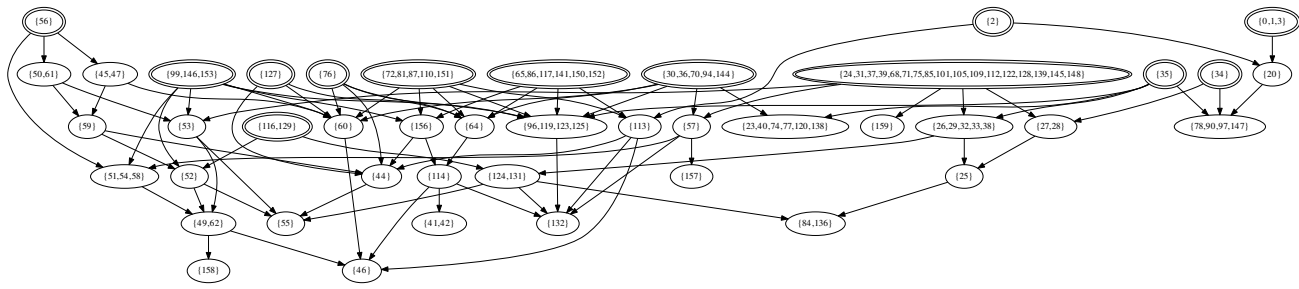


Figure 2: Dynamic Mutant Subsumption Graph for JTerminal

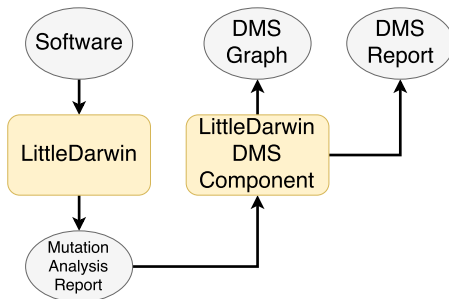


Figure 3: Dynamic Mutant Subsumption Component I/O

```

int nextCharacter = reader.read();
if (nextCharacter == -1) {
    // M56: nextCharacter == +1
    // M45: nextCharacter != -1
    buffer.append((char) character);
    break;
} else if (nextCharacter == MULTI_CSI[1]) {
    introducedControlSequence = true;
} else {
    text.append((char) character);
    text.append((char) nextCharacter);
}
    
```

Figure 4: Mutants 45 and 56 of JTerminal

graph, to visualize the subsumption relation, and (ii) a detailed report is generated in CSV² format that contains all the information processed by the DMS component. For each mutant, mutant ID, mutant path, source path, mutated line number, whether it is a subsuming mutant, number of failed tests, the mutants it subsumes, the mutants that it is subsumed by, and the mutants that are mutually subsuming with it are provided in this report. Since LittleDarwin is a Java mutation testing framework, the application of the DMS component is also restricted to Java programs.

To showcase the ability of LittleDarwin in performing dynamic mutant subsumption analysis, we use JTerminal³ as a subject. We improved the test suite of JTerminal by automatically generating test cases using EvoSuite [7]. The information about characteristics of JTerminal is shown in Table 2. The DMSG for JTerminal is depicted in Figure 2. In this figure, each number represents a single

²Comma-separated Values

³<https://www.grahamedgecombe.com/projects/jterminal>

```

@Test
public void testDoubleCsi() {
    char[] ch = { '2', '[', 's' };
    parser.parse(new String(ch));

    assertEquals(1, objects.size());
    Object obj = objects.get(0);
    assertEquals(AnsiControlSequence.class, obj.getClass());

    AnsiControlSequence seq = (AnsiControlSequence) obj;
    String[] params = seq.getParameters();
    assertEquals('s', seq.getCommand());
    assertEquals(0, params.length);
}
    
```

Figure 5: The Test that Kills Mutant 45, but Not Mutant 56

killed mutant, each node represents a group of mutants that are killed by exactly the same set of test cases, and each edge shows the dynamic subsumption relationship between each node where the node at the end is subsumed by the node at the start. The survived mutants are not shown in this figure. The double-circled nodes contain the subsuming mutant groups. In order to remove the redundant mutants, one only needs to keep one mutant from each subsuming mutant group and discard the rest.

Take M45 and M56 as an example. According to the DMSG, M56 subsumes Mutant M45. Using the CSV report, we can locate the actual mutation of the source code (Figure 4). Both M45 and M56 belong to method parse of class AnsiControlSequenceParser, and mutate the same statement on line 99. M45 acts as a negation of the conditional statement. This means that any input character (except -1) that used to trigger "else if" and "else", now trigger this branch. Since this branch contains a "break" statement, it avoids the rest of the iteration of the loop to be executed. If the input is -1, the "else" branch would be executed, which wrongfully appends -1 to "text". M56, however, changes only two special cases. If the input is +1, the "if" branch would be executed, and the current iteration breaks. If the input is -1, the same thing as M45 happens. For any other input, the program executes as it should. This means that M56 truly subsumes M45. Figure 5 shows the test case that kills M45, but not M56. The input value here is a single control sequence, which is neither -1 or +1, and therefore cannot kill M56. However, since it should have been handled by "else if" branch and M45 does not allow that, it kills M45. Hence, in Figure 2 (on the left side), we can see that M56 dynamically subsumes M45. Analysis such as this

Table 2: JTerminal Software Information

Project	Ver.	Size (LoC)		#C	TS	SC	BC	MC	#M
		Prod.	Test						
JTerminal	1.0.1	687	428	8	2	66%	56%	60.0%	160

Acronyms: Version (Ver.), Line of code (LoC), Production code (Prod.), Number of commits (#C), Team size (TS), Statement coverage (SC), Branch coverage (BC), Mutation coverage (MC), Number of Mutants (#M)

allows researchers to understand the relations between the mutants and reduce the effects of redundant mutants on their results.

5 CONCLUSION

Many academic studies in the field of software testing rely on mutation testing to use as their comparison criteria, and the existence of redundant mutants is a significant threat to their validity. We developed a component for our mutation testing tool, LittleDarwin, to facilitate the detection of redundant mutants using dynamic mutant subsumption analysis. We performed dynamic mutant subsumption analysis on a small, real-world project to demonstrate the capabilities of our tool. Using our tool, it is possible to detect and filter out redundant mutants, and help in increasing the confidence in results of experiments using mutation testing as a comparison criteria.

REFERENCES

- [1] P. Ammann, M. E. Delamaro, and J. Offutt. 2014. Establishing Theoretical Minimal Sets of Mutants. In *2014 IEEE Seventh International Conference on Software Testing, Verification and Validation*. 21–30. DOI: <https://doi.org/10.1109/ICST.2014.13>
- [2] James H. Andrews, Lionel C. Briand, and Yvan Labiche. 2005. Is Mutation an Appropriate Tool for Testing Experiments?. In *Proc. ICSE 2005 (27th international conference on software engineering) (ICSE '05)*. ACM, New York, NY, USA, 402–411. DOI: <https://doi.org/10.1145/1062455.1062530>
- [3] James H. Andrews, Lionel C. Briand, Yvan Labiche, and Akbar Siami Namin. 2006. Using Mutation Analysis for Assessing and Comparing Testing Coverage Criteria. *IEEE Transactions on Software Engineering* 32, 8 (aug 2006), 608–624. DOI: <https://doi.org/10.1109/tse.2006.83>
- [4] Timothy Alan Budd. 1980. *Mutation Analysis of Program Test Data*. Ph.D. Dissertation. Yale University, New Haven, CT, USA. AAI8025191.
- [5] Richard A. DeMillo, Richard J. Lipton, and F. G. Sayward. 1978. Hints on Test Data Selection: Help for the Practicing Programmer. *Computer* 11, 4 (apr 1978), 34–41. DOI: <https://doi.org/10.1109/C-M.1978.218136>
- [6] Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern Recognition Letters* 27, 8 (jun 2006), 861–874. DOI: <https://doi.org/10.1016/j.patrec.2005.10.010> ROC Analysis in Pattern Recognition.
- [7] Gordon Fraser and Andrea Arcuri. 2017. EvoSuite at the SBST 2017 Tool Competition. In *10th International Workshop on Search-Based Software Testing (SBST'17) at ICSE'17*. 39–42.
- [8] Yue Jia and Mark Harman. 2008. Constructing Subtle Faults Using Higher Order Mutation Testing. In *Proc. SCAM 2008 (Eighth IEEE International Working Conference on Source Code Analysis and Manipulation)*. Institute of Electrical & Electronics Engineers (IEEE), 249–258. DOI: <https://doi.org/10.1109/scam.2008.36>
- [9] Yue Jia and Mark Harman. 2011. An Analysis and Survey of the Development of Mutation Testing. *IEEE Transactions on Software Engineering* 37, 5 (sep 2011), 649–678. DOI: <https://doi.org/10.1109/TSE.2010.62>
- [10] René Just, Darioush Jalali, Laura Inozemtseva, Michael D. Ernst, Reid Holmes, and Gordon Fraser. 2014. Are Mutants a Valid Substitute for Real Faults in Software Testing?. In *Proc. FSE 2014 (Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering) (FSE 2014)*. ACM, New York, NY, USA, 654–665. DOI: <https://doi.org/10.1145/2635868.2635929>
- [11] D. Richard Kuhn. 1999. Fault Classes and Error Detection Capability of Specification-based Testing. *ACM Trans. Softw. Eng. Methodol.* 8, 4 (Oct. 1999), 411–424. DOI: <https://doi.org/10.1145/322993.322996>
- [12] B. Kurtz, P. Ammann, and J. Offutt. 2015. Static analysis of mutant subsumption. In *Software Testing, Verification and Validation Workshops (ICSTW), 2015 IEEE Eighth International Conference on*. 1–10. DOI: <https://doi.org/10.1109/ICSTW.2015.7107454>
- [13] Aditya P. Mathur. 1991. Performance, effectiveness, and reliability issues in software testing. In *Proc. COMPSAC 1991 (The Fifteenth Annual International Computer Software & Applications Conference)*. Institute of Electrical & Electronics Engineers (IEEE), 604–605. DOI: <https://doi.org/10.1109/cmpsac.1991.170248>
- [14] A. Jefferson Offutt, Ammei Lee, Gregg Rothermel, Roland H. Untch, and Christian Zapf. 1996. An Experimental Determination of Sufficient Mutant Operators. *ACM Transactions on Software Engineering Methodology* 5, 2 (April 1996), 99–118. DOI: <https://doi.org/10.1145/227607.227610>
- [15] A. Jefferson Offutt and Jie Pan. 1997. Automatically detecting equivalent mutants and infeasible paths. *Software Testing, Verification and Reliability* 7, 3 (sep 1997), 165–192. DOI: [https://doi.org/10.1002/\(sici\)1099-1689\(199709\)7:3<165::aid-stvr143-3.0.co;2-u](https://doi.org/10.1002/(sici)1099-1689(199709)7:3<165::aid-stvr143-3.0.co;2-u)
- [16] A. Jefferson Offutt, Gregg Rothermel, and Christian Zapf. 1993. An Experimental Evaluation of Selective Mutation. In *Proc. ICSE 1993 (15th international conference on Software engineering) (ICSE '93)*. IEEE Computer Society Press, Los Alamitos, CA, USA, 100–107. <http://dl.acm.org/citation.cfm?id=257572.257597>
- [17] A. Jefferson Offutt and Roland H. Untch. 2001. Mutation 2000: Uniting the Orthogonal. In *Mutation Testing for the New Century*, W. Eric Wong (Ed.). The Springer International Series on Advances in Database Systems, Vol. 24. Springer US, 34–44. DOI: https://doi.org/10.1007/978-1-4757-5939-6_7
- [18] Mike Papadakis, Christopher Henard, Mark Harman, Yue Jia, and Yves Le Traon. 2016. Threats to the Validity of Mutation-based Test Assessment. In *Proceedings of the 25th International Symposium on Software Testing and Analysis (ISSTA 2016)*. ACM, New York, NY, USA, 354–365. DOI: <https://doi.org/10.1145/2931037.2931040>
- [19] Ali Parsai. 2015. *Mutation Analysis: An Industrial Experiment*. Master's thesis. University of Antwerp.
- [20] Ali Parsai, Alessandro Murgia, and Serge Demeyer. 2016. Evaluating Random Mutant Selection at Class-level in Projects with Non-adequate Test Suites. In *Proc. EASE 2016 (20th International Conference on Evaluation and Assessment in Software Engineering) (EASE '16)*. ACM, New York, NY, USA, Article 11, 10 pages. DOI: <https://doi.org/10.1145/2915970.2915992>
- [21] Ali Parsai, Alessandro Murgia, and Serge Demeyer. 2016. A Model to Estimate First-Order Mutation Coverage from Higher-Order Mutation Coverage. In *Proc. QRS 2016 (IEEE International Conference on Software Quality, Reliability and Security)*. Institute of Electrical and Electronics Engineers (IEEE), 365–373. DOI: <https://doi.org/10.1109/QRS.2016.48>
- [22] Ali Parsai, Alessandro Murgia, and Serge Demeyer. 2017. LittleDarwin: a Feature-Rich and Extensible Mutation Testing Framework for Large and Complex Java Systems. In *Proc. FSEN 2017 (Fundamentals of Software Engineering)*.
- [23] W. Eric Wong and Aditya P. Mathur. 1995. Reducing the cost of mutation testing: An empirical study. *Journal of Systems and Software* 31, 3 (1995), 185–196. DOI: [https://doi.org/10.1016/0164-1212\(94\)00098-0](https://doi.org/10.1016/0164-1212(94)00098-0)
- [24] Lingming Zhang, Milos Gligoric, Darko Marinov, and Sarfraz Khurshid. 2013. Operator-based and random mutant selection: Better together. In *Proc. ASE 2013 (28th IEEE/ACM International Conference on Automated Software Engineering)*. Institute of Electrical & Electronics Engineers (IEEE), 92–102. DOI: <https://doi.org/10.1109/ASE.2013.6693070>
- [25] Lu Zhang, Shan-Shan Hou, Jun-Jue Hu, Tao Xie, and Hong Mei. 2010. Is Operator-based Mutant Selection Superior to Random Mutant Selection?. In *Proc. ICSE 2010 vol. 1 (Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1) (ICSE '10)*. ACM, New York, NY, USA, 435–444. DOI: <https://doi.org/10.1145/1806799.1806863>